

Chang, Myong-hun; Harrington, Joseph E.

Working Paper

Agent-based models of organizations

Working Paper, No. 515

Provided in Cooperation with:

Department of Economics, The Johns Hopkins University

Suggested Citation: Chang, Myong-hun; Harrington, Joseph E. (2004) : Agent-based models of organizations, Working Paper, No. 515, The Johns Hopkins University, Department of Economics, Baltimore, MD

This Version is available at:

<https://hdl.handle.net/10419/71984>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Agent-Based Models of Organizations
Handbook of Computational Economics II:
*Agent-Based Computational Economics**
Kenneth L. Judd and Leigh Tesfatsion, editors

Myong-Hun Chang
Department of Economics
Cleveland State University
Cleveland, OH 44115
216-687-4523, -9206 (Fax)
m.chang@csuohio.edu
academic.csuohio.edu/changm

Joseph E. Harrington, Jr.
Department of Economics
Johns Hopkins University
Baltimore, MD 21218
410-516-7615, -7600 (Fax)
joe.harrington@jhu.edu
www.econ.jhu/people/harrington

March 24, 2005

Contents

1	Introduction	3
2	How to model an organization	6
2.1	Agents	7
2.2	Organizations	8
2.3	Environments	10
2.4	Implementation of an Agent-Based Model of an Organization	11
2.5	How Does Agent-Based Computational Economics Differ from Neo-classical Economics?	12
3	Search and Learning	14
3.1	Modelling Search	16
3.2	Organizational Search with Units Solving Similar Problems	20
3.3	Organizational Search with Units Solving Different Problems	25
3.4	Evolving an Organizational Structure	30
3.5	What Do We Learn from a Computational Agent-Based Approach?	32
4	Information Processing	34
4.1	Generic Properties of Information Processing Networks	35
4.2	Organizations with Adaptive/Learning Agents	37
4.3	Adaptation and Evolution of Organizational Structure	41
4.4	Summary	43
5	Effort, Norms, and Endogenous Hierarchies	44
5.1	Effort and the <i>Commons</i> Problem in Organizations ¹	44
5.2	Organizational Norms	47
5.3	Growing an Organization	49
6	Critique of the Past and Directions for the Future	52

*The authors gratefully acknowledge the comments of Rich Burton, Josh Epstein, Vijay Mathur, Scott Page, Jan Rivkin, Nicolaj Siggelkow, Leigh Tesfatsion, and two anonymous economic theorists.

¹The issues addressed in this section are closely related to the concerns of Janssen and Ostrom (2005).

1 Introduction

An organization is a collection of agents that interact and produce some form of output. Formal organizations - such as corporations and governments - are typically constructed for an explicit purpose though this purpose needn't be shared by all organizational members. An entrepreneur who creates a firm may do so in order to generate personal wealth but the worker she hires may have very different goals. As opposed to more amorphous collections of agents such as friendship networks and societies at large, organizations have a formal structure to them (though informal structures typically emerge as well) with the prototypical example being a corporation's organizational chart. This structure serves to define lines of communication and the distribution of decision-making. Organizations are also distinguished by their well-defined boundaries as reflected in a clear delineation as to who is and who is not a member. This boundary serves to make organizations a natural unit of selection; for example, corporations are formed and liquidated though they can also morph into something different through activities like mergers.

The primary task of organization theory is to understand how organizations behave and to identify and describe the determinants of organizational performance.² To take an *agent-based* approach means not having to assign an objective to an organization and instead modelling the agents that comprise it with explicit attention to how decisions are made and how the interaction of these decisions produce organizational output. The smallest decision-making unit is then required to be smaller than the organization itself. The anthropomorphic view associated with the theory of the firm - firms are profit-maximizers - is not an agent-based model. Though neo-classical economics has many agent-based models of organizations, including agency theory and team theory, these models are generally quite restrictive in terms of the assumptions placed on agent behavior, the number and heterogeneity of agents, the richness of the interaction among agents, and the features of the environment. These restrictions are forced upon scholars by virtue of the limited power of analytical methods. To derive universal results ("proving" them) requires limiting the size of one's universe (the class of models). While some structures are relatively simple in their real form (for example, many auctions), organizations are inherently complex; they are their own brand of society, plagued with conflicting interests while dealing with multi-faceted problems amidst a coevolving environment. Proving universal results is only achieved at the cost of severely restricting the richness of the setting.

A *computational* agent-based model uses the power of computing to "solve" a model. A model is written down, parameter values are specified, random variables are realized, and, according to agents' behavioral rules, agent output is produced. Organizational output comes from the specified mapping from the environment and agents' actions into the output space. At the end of vast CPU time, the simulation output can yield results that are rich and insightful but ultimately are a collection of examples, perhaps many examples - thousands of periods, hundreds of runs, dozens of

²A thoughtful statement as to what is an organization and what organization theory is about can be found in Aldrich (1999).

parameter configurations - but still noticeably finite. In deploying numerical methods, the presumption is that the model is unsolvable by the human mind (in practice, not necessarily in principle). If the use of computing power is not to reflect laziness or ineptness on the modeler's part, a computational agent-based model must then have some minimum level of complexity - whether due to agent heterogeneity, the structure of interactions among agents, a poorly behaved environment, dynamics, or some other feature. A legitimate computational agent-based model is then not simply one that is solved by a computer but rather one for which it is necessary that it be solved by a computer.³

Organization theory is traditionally of two varieties: i) broad, institutionally rich, and vague while using informal arguments articulated in a narrative; and ii) narrow, simplistic, and mathematically precise while using formal logic articulated in a set of assumptions, a statement of a theorem, and a proof. The appeal of computational modelling is that it achieves middle ground in that it has the precision of (ii) and the ability to handle a rich set of features as with (i). It trades off the universality of results of (ii) for a richer model while maintaining rigor and formality. This trade-off is generally judged to be a good one when it comes to modelling a complex entity such as an organization.

In writing this chapter, the hope is to describe to the reader the central research questions addressed, synthesize the models and methods deployed, and survey some of the new insight being delivered. Given the incipency of this literature, what we will not provide is a coherent set of results because such has not yet emerged. Work on computational agent-based models of organizations is very much in the exploratory phase with highly varied approaches to pursuing a broad range of questions. Our objective is to introduce it to economists in a methodological manner and provide a broader and less idiosyncratic perspective to those who are already engaging in this type of work.

Before launching into specific models, let us offer a quick review of some of the questions addressed by research so that these can frame the reader's mind. What are the determinants of organizational behavior and performance? How does organizational structure influence performance? How do the skills and traits of agents matter and how do they interact with structure? What determines whether more skilled agents and a more decentralized structure are complements or substitutes? What is the proper balance of exploration and exploitation? How do the characteristics of the environment - including its stability, complexity, and competitiveness - influence the appropriate allocation of authority and information? How is behavior and performance influenced when an organization is coevolving with other organizations from which it can learn? Can an organization evolve its way to a better structure?

Related Literatures There are a number of closely related literatures that will not be covered here. A more complete treatment of agent-based models of organizations

³Not all computational models of organizations are agent-based; Carroll and Harrison (1998) being an example. Their formulation begins not with a specification of agents but rather a system of equations describing hiring, socialization, and turnover.

would discuss the extensive literature in neoclassical economics on organizations.⁴ These models are rich in their modelling of incentives but mired in poverty when it comes to modelling agent heterogeneity, the cognitive limitations of agents, organizational structure, and the coevolving nature of a population of organizations. A second related literature is on networks, for implicit in any non-degenerate model of an organization is a network which describes how agents communicate and influence each other. As there are two other chapters in this handbook devoted to networks [Vriend (2005), Wilhite (2005)], we will generally exclude such work other than that which is specifically designed to understand organizations. Thirdly is work on distributed artificial intelligence which develops better ways to solve problems through the distribution of tasks among agents.⁵ While some of these models have something meaningful to say regarding the questions of this chapter, the ultimate objective is quite different. For example, as the objective is developing more efficient solutions rather than better explaining organizations, it is common to assume agents' goals coincide with the organizational goal. Finally, there is the line of work best referred to as organizational engineering. This research develops a relatively literal description of an organization which can then be calibrated and simulated to provide quantitative answers to policy questions. As a result, the models are not designed to provide qualitative insight and have different objectives from the work that is reviewed here. At the risk of unfairly over-generalizing, organizational engineering models are designed for prediction, not explanation.

As to other review articles, the Introduction to Lomi and Larsen (2001) offers a most enlightening historical perspective that draws on many scholarly antecedents. The review article closest to what we are doing here is Carley and Gasser (1999) though they give emphasis to organizational engineering. Sorensen (2002) provides a nice review of organizational models based on the *NK*-approach (Kauffman, 1993) and cellular automata. One of the best papers that discusses the general topic of complexity and how it relates to issues in organization theory is Carroll and Burton (2000). Collections of papers dealing with computational organization theory (not just agent-based modelling) include Baum and Singh (1994), Carley and Prietula (1994), Cohen and Sproull (1996), Prietula, Carley, and Gasser (1998), and Lomi and Larsen (2001). Also see Baum (2002) for general work on organizations with several entries dealing with computational modelling.

Roadmap and a Guide for Neoclassical Economists A synthesis of the central features of computational agent-based models of organizations is provided in Section 2. The literature itself is partitioned according to the basic task with which an organization is faced. Section 3 focuses on models for which organizations search and learn; it represents the most well-developed body of work. Section 4 looks into modelling the processing of information which is, roughly speaking, a production function for organizational decision-making. Thus far, models are a bit mechanical

⁴Holmstrom and Tirole (1989), Milgrom and Roberts (1992), and Prendergast (1999) offer good general treatments.

⁵See, for example, Durfee (1999) and Mackie-Mason and Wellman (2005).

and the literature is not as developed. While these two research strands make up the bulk of the literature, other issues are tackled and Section 5 describes the best work on some of the more important organizational issues not covered in Sections 3 and 4. A critical appraisal is provided in Section 6 where we also identify some lines for future work.

For the neoclassical economist largely unfamiliar with computational agent-based modelling, we recommend focusing on Sections 2 and 3. Section 2 introduces many concepts and elements of this modelling approach and, in its final subsection, contrasts neoclassical and computational agent-based models and suggests why a neoclassical economist should be interested in these methods. The search and learning literature reviewed in Section 3 is the closest in style to that conducted by neoclassical economists and, in addition, we take the opportunity to begin synthesizing existing results and contrasting the associated insight with what one would get using a neoclassical approach. We ask: What do we learn from the computational agent-based approach that we would not have learned from using the neoclassical approach?

2 How to model an organization

How can intelligence emerge from nonintelligence? To answer that, we'll show that you can build a mind from many little parts, each mindless by itself. ... These we'll call agents. Each mental agent by itself can only do some simple thing that needs no mind or thought at all. Yet when we join these agents in societies - in certain very special ways - this leads to true intelligence. [Marvin Minsky, *The Society of Mind* (1986), p. 17]

The typical neoclassical description of a firm - the organization that has drawn the most attention within economics - is as a profit-maximizing entity. Being a single-agent formulation, it represents a rather uninteresting model of an organization.⁶ Similarly, there are models in the agent-based literature, such as the early work of Levinthal (1997), that model an organization as a single agent adaptively learning. However, to be a meaningful agent-based model of an organization, an agent must be "smaller" than the organization itself. But then, how does organizational behavior emerge from a collection of agents making choices? Just as Marvin Minsky asks how mindless components can form a mind and produce intelligence, we ask how agents - representing human actors - can form an organization and produce output beyond the capacity of any individual agent.

This section is divided into five parts. The first part reviews the concept of an agent. An agent represents the smallest decision-making unit of an organization. Next we turn to examining the various dimensions of an organization; what transforms a collection of agents into an organization? The third section describes the environment into which an organization is placed and the task with which it is presented. The

⁶Though no economist would see the theory of the firm as a model of an organization, this misses the point. The theory of the firm is used to make predictions about corporations which *are* organizations.

fourth section offers a brief discussion on computationally implementing an agent-based organizational model. In the final section, this approach is contrasted with the more standard approach in economics.

2.1 Agents

There are many definitions of an “agent” in the agent-based literature. An agent is said to be purposeful, autonomous, adaptive, and so on. While these terms serve to convey a sense of what the researcher is after, they only shift the question of “what is an agent?” to “what does it mean to be purposeful? autonomous? adaptive?” Perhaps the best we can do is to describe our intent - what is this thing called an agent suppose to represent? - and what we actually do - how is an agent instantiated? In almost all models of organizations, an agent represents a flesh-and-blood human.⁷ Being purposeful may mean adjusting behavior to improve some measure of well-being; being autonomous may mean choosing actions even if they are in conflict with an organizational goal; being adaptive may mean modifying behavior in response to past experiences. Though the terms are vague, the way in which they are implemented has substantive content.

The neoclassical approach in economics to modelling agents takes preferences and beliefs as primitives. Typically, an agent is endowed with a utility function and, given beliefs over that which is unknown to the agent, acts to maximize expected utility or, in an intertemporal setting, the expectation of the discounted sum of utility. When an agent is making choices in a multi-agent context and what is best depends on what others do - and this certainly describes an organization - this approach is augmented with the (Bayes-Nash) equilibrium assumption that each agent understands how other agents behave. This doesn’t necessarily mean that agents know exactly what others will do but they do know other agents’ decision rules - how private information maps into actions. Agents have complete understanding though may lack complete information.

In contrast to the assumption of a hyper-rational agent, it is standard in the computational agent-based literature to assume agents are boundedly rational. The most concise statement of this modelling approach is that agents engage in adaptive search subject to various cognitive constraints (and not just informational constraints). These models may continue to deploy the optimization framework though assuming myopic optimization and that beliefs are empirically-based rather than the product of understanding what is optimal behavior for others. Agents observe but do not necessarily theorize. For example, a common specification is that an agent engages in hill-climbing as it adopts a new alternative when doing so yields higher current performance than the previously selected action (Chang and Harrington, 2000). Or the optimization framework may be entirely discarded as preferences and beliefs are replaced with behavioral rules cast as primitives. For example, in information processing models, an agent receives data and is endowed with a rule that converts

⁷This needn’t be the case for, in actual organizations, agents can be software such as expert systems or automated bidding rules at auctions.

it into a lower-dimensional message sent to the next agent in line (Carley, 1992; Barr and Saraceno, 2002).

Within this bounded rationality framework, models often provide a parameter by which one can “tune” the cognitive skills of an agent. When rules adapt to experience, a key parameter is how much experience an agent has as well as the size of memory (Carley, 1992). In the context of information dissemination, the likelihood that an agent observes an innovation reflects a level of skill (DeCanio and Watkins, 1998). For hill-climbing algorithms, agents may only evaluate alternatives imperfectly - less skilled agents may have noisier evaluations (Chang and Harrington, 1997) - or are constrained in the set of alternatives - more skilled agents are able to consider options in a wider neighborhood around their current practices (Kollman et al., 2000). A novel and promising approach is to assume that an agent has a “model” of how actions map into performance but where the model is of lower dimensionality than reality (Gavetti and Levinthal, 2000).

2.2 Organizations

Let us now turn to the issue of what transforms a collection of agents into an organization. Our discussion is organized along three questions. Who comprises an organization? How are agents connected to produce organizational output? And, how are agents motivated?

An organization is comprised of multiple agents and indeed one common question in the literature is how the number of agents influences organizational performance. But more than pure numbers is relevant, especially when agents are heterogeneous. There is an architecture to organizations, which we will elaborate upon momentarily, which raises questions of how agents are distributed across various units and how agents are matched to tasks. Given the often significant role to agent heterogeneity in computational agent-based models, it is surprising that there is little research exploring how agents with different skills are distributed across the different levels of an organization. This is an area begging for work.

Organizational structure is another one of those terms that has defied a common definition. A broad but useful one refers to it as “those aspects of the pattern of behavior in an organization that are relatively stable and that change only slowly” (March and Simon, 1958, p. 170). Under the rubric of organizational structure, we will place three dimensions. First, there is the allocation of information. This refers to how information moves between the environment and the organization - which agents receive data from the environment - and how it moves within the organization - who reports to whom. This may have a fairly stable component to it, as might be described by the rules of communication laid out in an organizational chart. Such well-defined flows are a common feature of information processing models (Miller, 2001). However, just as people create dirt paths in a park by veering from the sidewalk, information can flow outside of mandated channels. There is then an endogenous feature to how information is distributed.⁸ For example, Chang and

⁸This type of model is more fully explored in Vriend (2005).

Harrington (2000) allow an upper level manager to observe a new practice and then decide whether to communicate it to lower level agents.

A second element to organizational structure is the allocation of authority - who makes the decisions - associated with which are two critical facets: modularity and decentralization. An organization may have to perform many sub-tasks in solving a problem and a key structural issue is how these sub-tasks are combined into distinct modules which are then re-integrated to produce an organizational solution. The degree to which a problem can be efficaciously modularized depends on the nature of the task (what is referred to as decomposability, which we discuss later). Two classic structures that represent alternative modular forms are the *M-form* - where all of the sub-tasks associated with a particular product line are combined - and the *U-form* - where all similar sub-tasks are combined (for example, the marketing divisions for all product lines are in the same module). With this allocation of tasks, there is still the issue of which agents ultimately make the decisions. In the context of a hierarchy - which describes most organizations - to what degree is authority centralized in higher levels? Is authority matched with who has the best information? Here we are referring to formal authority which, as noted by Aghion and Tirole (1997), may differ from real authority. If an agent with decision-making authority relies heavily on the information provided by other agents then the real authority (or power) may lie with those providing the information. The allocation of information and real authority are thus intertwined.

A third element of structure is the least well-defined: organizational norms and culture. Though there are probably as many definitions of culture as scholars who have sought to define it, we'll put forth the one of Sathe (1985): "Culture is a set of shared assumptions regarding how the world works (beliefs) and what ideals are desirable (values)." Agent behavior is somehow influenced by an organization's past and this past is embodied in what is called norms or culture. Of particular interest is modelling the associated feedback dynamic - norms, being determined by past behavior, influence current behavior which then serves to define future norms. This is a driving force in March (1991).

The final element to organizations to be covered is agent motivation. Agents may be modelled as having preferences - for example, they desire income and dislike exerting effort - but how that translates into behavior depends on an organization's incentive scheme for rewarding and punishing.⁹ The compensation scheme for corporate managers may drive them to seek higher organizational profit, while the scheme for division managers may be tied to division profit (so as to induce high effort) which can then create a conflict of interests. Conflict may also arise when an organization uses promotion or bonuses based on relative performance to encourage effort.¹⁰ An important element to any conception of an organization is the degree of such conflicts and how it varies within and across levels. By contrast, models of distributed prob-

⁹Many computational agent-based models of organizations are not explicit about the form of the incentive scheme but, if one makes standard assumptions about agents' preferences, there is often an obvious implicit specification.

¹⁰Though these forces haven't been modelled in the agent-based literature, there has been some computational work elsewhere (Harrington, 1998, 1999).

lem solving in AI assume agents have a coherence of goals. More realistic models of organizations recognize that conflict of interest is an endemic feature of actual organizations.

An organization has an output - say, a set of practices - and delivers some measure of performance. Performance may be measured by profit (or some analogous criterion) or may involve specifying a particular target (for example, the global optimum) and then measuring performance by the frequency with which an organization reaches it or, if eventually it'll always reach it, the average time it takes to do so. While most organizations are designed with a particular objective in mind (the objective of the entrepreneur), it doesn't follow that organizational behavior is consistent with that or any other objective. An organization's members may have different goals than those of the entrepreneur. Fortunately, an agent-based model needn't answer the dicey question of "what is an organization's utility function" as it is sufficient to instantiate agents and let organizational behavior emerge from the interaction of agents amongst themselves and with the environment. By building an organization from the ground-up, we can avoid taking an anthropomorphic view to complex entities such as organizations.

2.3 Environments

An organization resides in an environment and is faced with a problem (or task) and constraints to be faced in trying to solve it. The problem may be choosing a political platform, if it is a political party, or producing and selling a product, if it is a firm. Problems vary in terms of their difficulty. A problem may be more difficult because it requires more information. It may be more difficult because there are interactions between various choice variables which makes it less likely that one can search, dimension by dimension, for a multi-dimensional solution. Relatedly, it may be more difficult because directed search is infeasible or ineffective. Knowing where to go from one's current position to achieve higher performance can greatly ease search. Such directed search may be infeasible because there is no metric on the solution space; there is no notion of two solutions being close. Even if there is a metric, the relationship between performance and actions may not be well-behaved in that the components of the gradient may quickly change sign and admit many optima. This means that hill-climbing algorithms can get stuck on lousy local optima and it isn't clear where to look for better ones.¹¹

Related to the issue of difficulty is the extent to which a problem is decomposable. A problem is said to be decomposable if there exists a way in which to partition it into sub-problems such that the concatenation of the solutions to the sub-problems is a solution to the original problem. Such problems are easier and quicker to solve as it means solving a collection of simpler (lower dimensional) problems in parallel. Furthermore, how a problem decomposes suggests a "natural" organizational structure, an issue explored in Ethiraj and Levinthal (2002).

¹¹Page (1996) provides a rigorous investigation into what it means for a problem to be more difficult from the perspective of search algorithms.

An organization's problem may also have a dynamic component to it. In solving a single task in real time, the best solution may evolve with changes in the environment. A less stable environment makes the problem more difficult as the organization is pursuing a moving target. Or an organization may face a series of problems. Is the same problem being faced repeatedly or are the problems distinct and, if so, how are they related? As long as the problems retain some similarity, the solution to one will provide clues for another, thereby creating opportunities to learn.

A more distinctive feature of an organization's environment is the presence of other entities that are also solving problems; there may be a pool of organizations coevolving. Other organizations may influence an organization's current performance - consider a setting in which they compete (Barr and Saraceno, 2003) - or influence future performance when they can learn from each other (Miller, 2001) or exchange personnel (Axtell, 1999). There may be other adapting agents such as consumers - Chang and Harrington (2003) allow consumers to search at the same time that firms are adapting their practices - or lobbyists (if the organizations are governments). In providing an endogenous source of change in an organization's environment, coevolution can provide rich and non-trivial dynamics.

2.4 Implementation of an Agent-Based Model of an Organization

Having identified many of the components that go into an agent-based model of an organization, how does one implement it computationally? As space constraints prevent a comprehensive answer, let us focus on two broad and essential elements to implementation: *agent processes* and *super-agent processes*. In a computational model, an agent is instantiated as a mapping from inputs into outputs. Input includes information from outside the organization (from customers, input suppliers, competitors, etc.) and information from inside the organization (subordinates, peers, superiors); it takes the form of processed information, new ideas, actions. The ensuing output may be a concrete action or a message to other agents. The important point is that many of the elements of an organization - communication network, hierarchy, incentive schemes, and the like - are embedded in an agent's mapping. When one writes a code that specifies that agent i observes some data and makes a recommendation to agent h who, after also receiving a recommendation from agent j , chooses between these alternatives, one is making assumptions about the allocation of information (i and j receive information from the environment while h does not) and the allocation of authority (h has authority while i and j do not). The particular form of this mapping similarly depends on organizational features such as the form of compensation and norms (peer pressure, standard operating practices, etc.) as well as agent-specific traits including preferences, beliefs, and cognitive skills. This mapping may evolve over time due to learning but also because the identity of an agent changes as a result of personnel turnover. In sum, an organization is implicit in the modelling of agents' mappings. What emerges from the interaction among agents and the environment is organizational behavior.

Lying on top of these agent processes are super-agent processes which systematically influence an organization but are not embodied in agents' mappings. Super-

agent processes are commonly used to endogenize organizational structure. This may mean using a genetic algorithm defined over a population of organizations which creates new organizational designs and weeds out poorly performing ones (Miller, 2001).¹² Or one might model the adaptation of organizational design through a hill-climbing algorithm (Ethiraj and Levinthal, 2002) or simulated annealing (Carley and Svoboda, 1996). These super-agent processes provide a black box mechanism to substitute for modelling the agents who actually make these decisions. For example, a CEO typically decides on organizational structure, creditors decide whether to force an organization to exit, and entrepreneurs decide whether to create a new organization. As modelling all agents is often too daunting a task, super-agent processes represent a parsimonious way in which to encompass these other forces.

2.5 How Does Agent-Based Computational Economics Differ from Neoclassical Economics?

The objective of this section is to summarize the essential differences between agent-based computational economics (or ACE) and neoclassical economics (or NCE). In so doing, we will argue why economists ought to be interested in ACE.

The first essential difference is that agent behavior is characterized by *adaptive search* in ACE, which departs from the assumption in NCE that agent behavior is optimal (for some preferences and beliefs). In short, NCE describes “what is best,” while ACE describes “what is better.” With ACE, learning is based more on experience than understanding, more on retrospection rather than foresight. Furthermore, imperfections to agent behavior are modelled very differently. With NCE, imperfections are due to incomplete information. Consistent with the bounded rationality approach, cognitive limitations are central to ACE which means that what information is possessed may not be fully processed. This distinction between optimal behavior and adaptive search has a considerable impact on the logic of the model and the ensuing insight that is produced. This will come out in Section 3 when we examine a particular class of ACE models.

The next two distinctive elements of ACE emanate from the methods used in solving the model. Results are proven with NCE, while they are numerically derived for a particular parameter specification with ACE. Computational implementation has implications for both modelling and analysis.

The second essential difference is the way in which agents’ environments are modelled. The forte of expert NCE modelers is constructing a well-behaved environment in the sense of, for example, having a unique optimum or equilibrium and allowing comparative statics to be signed. In other words, building a plausible model that can be mentally solved. With ACE, there are much fewer constraints of this sort

¹²The role of selection is particularly interesting because part of what makes a collection of agents an “organization” is that it is a unit of selection. Corporations are created and fail; governments are put in power and overthrown. By comparison, general societies are more amorphous and thus less natural a unit of selection. Indeed, conquerors can be assimilated in which case which society has really prevailed? While the same might be said of firms - consider hostile takeovers - it is not as compelling.

since the model is solved numerically. This allows for *complex environments* which are richer with more descriptive realism. Without as many modelling constraints, a researcher is more apt to be able to make the primitive assumptions thought to be most appropriate and let the environment be what it will be. Complexity is not shunned but rather embraced when it is a property of the environment that actual agents and organizations face. In short, ACE allows for richer environments than does the NCE approach and, furthermore, makes complexity a trait of the environment whose role is to be explored. Indeed, research reveals that qualitative results can vary significantly with environmental complexity.

The third essential difference is in the mode of analysis. Dynamic models in the NCE tradition typically focus on the long-run, whether a steady-state or a stationary distribution. Behavior is characterized when all has settled - the environment has calmed down (in actuality or in expectation) and the system has converged to some form of equilibrium. A primary virtue of the ACE approach is that, by running simulations, it can describe *medium-run dynamics*. By medium-run dynamics we mean that some learning and adaptation has taken place but the system is not close to stabilizing. Not only are medium-run dynamics important if one wants to understand the transitional impact of various policies but, if convergence to equilibrium is slow (or if there is convergence at all), it may be the time scale of greatest relevance.

These three identifying traits of ACE - adaptive search with cognitive constraints, complex environments, and medium-run dynamics - are quite complementary in that a complex environment makes optimal behavior more problematic an assumption and, furthermore, it is more appropriate to describe the system using medium-run dynamics rather than a long-run equilibrium.

In light of these unique features, economists should be interested in ACE because it offers a new set of modelling and analytical tools which, in many instances, are quite complementary to that of NCE. First, a computational agent-based approach can be used when the environment is inherently complex and poorly behaved (multiple optima, non-quasi-concave, coevolution among agents, etc.) so that analytical methods are likely to fail and the assumption of game-theoretic equilibrium is particularly problematic as a characterization of behavior. Rather than making heroic assumptions on behavior and the environment in order to ensure the model can be mentally solved, one can use ACE methods. Second, ACE can characterize medium-run dynamics, a long neglected element of NCE in spite of its importance. Third, ACE methods can be used to explore not just traditional NCE issues - such as the role of organizational structure - but also previously ignored issues such as the role of environmental complexity and the cognitive limitations of organizational members. Complexity may differ across economic settings because of the production process and the extent of complementarities among an organization's activities. For example, greater connectedness among agents due to innovations in information technology may mean a better global optimum but also a more complex environment in terms of more local optima. Cognitive limitations may differ across organizations because of education, training, and how effectively an organization "selects" smarter people. Also, the extent to which cognitive skills matter will vary across positions within an

organization; such skills are less important for tasks that can be routinized and more important for those that are continually subject to novelty. These new tools and issues are capable of providing new insight into organizations, as we'll show in this chapter.

Free your mind. [Morpheus to Neo from *The Matrix*.]

A challenge to a neoclassical economist in reading this chapter will be the unorthodox logic of these models. The optimization framework produces a certain logic which can be quite distinct from that due to adaptive search. The canonical ACE environment is one in which an agent is searching on a landscape with multiple optima. Changes - such as with respect to organizational structure - may actually result in a lower global optimum but nevertheless enhance performance because search may not always find the global optimum or medium-run dynamics may generally not be near the global optimum. For example, a change which throws the organization into the basin of attraction for a better optimum can enhance performance even though it may be deleterious in the short-run. The logic of these models rests not just on how the landscape is affected in terms of its highest point but on a broader range of landscape properties which impact how search is conducted. With NCE, what matters are the set of optima or equilibria; with ACE, the entire landscape can matter because at issue is how likely adaptive search can take an organization from one point of the space to another. The path matters and not just the destination.

... a straight line may be the shortest distance between two points, but it is by no means the most interesting. [The Doctor from "The Time Monster" episode of *Doctor Who*.]

3 Search and Learning

[T]he assumption that business behavior is ideally rational and prompt, and also that in principle it is the same with all firms, works tolerably well only within the precincts of tried experience and familiar motive. It breaks down as soon as we leave those precincts and allow the business community under study to be faced by - not simply new situations, which also occur as soon as external factors unexpectedly intrude - but by new possibilities of business action which are as yet untried and about which the most complete command of routine teaches nothing. [Joseph A. Schumpeter, *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process* (1939), p. 98]

In this section, we take the perspective that a primary task of an organization is to constantly search for and adopt routines that improve (though do not necessarily maximize) performance. This *search-and-learn* perspective of a firm, as an alternative to the neoclassical approach, is central to the evolutionary theory of the firm where firms are "modeled as simply having, at any given time, certain capabilities and decision rules [which are] modified as a result of both deliberate problem-solving efforts and random events." (Nelson and Winter, 1982, p. 4)¹³

¹³One of the earliest computational papers on organizational search is Levinthal and March (1981).

As formulated by Schumpeter and Nelson and Winter, a firm is represented by a single agent - an entrepreneur carrying out search and making performance-enhancing adoption decisions for the entire enterprise. The agent-based approach to modelling organizations takes this one step further. It recognizes that the bounded rationality on the part of a single decision maker, faced with a large and complex routine space, makes an organizational search strategy utilizing multiple agents compelling. The main objective of this research program is understanding how a firm's performance is influenced by the way in which parallel search is carried out among multiple agents.¹⁴ This typically takes the form of managers of various departments independently searching for better routines. Furthermore, if we make the reasonable assumption that there is no single individual who is instantaneously and costlessly informed of all new knowledge in the organization, it then becomes crucial for effective organizational decision-making that there be collocation of the uncovered information and the right to act on that information. This collocation may occur at the top, thereby requiring knowledge to be pulled up the hierarchy, or at lower-level units, thereby requiring decision rights to be pushed down (Jensen and Meckling, 1995).

As Hayek (1945) stated so forcefully, the assumption of bounded rationality puts an upper limit on the effectiveness with which the central authority can process and act on the large set of information sent up by an organization's lower levels. Pushing against this limit are two beneficial roles that the centralized authority structure may play in formal organizations. First, it can act as a conduit for *knowledge transfer*. Depending on the circumstances surrounding the local units, a piece of information uncovered by one may prove to be of value to other units. The global exploitation of a local discovery realizes an immediate static gain - as a useful routine is shared - but, as we will later explain, there may also be dynamic implications in that mutual learning can influence what units adopt in the future. While an informal social learning mechanism may be capable of facilitating these static and dynamic processes, upper level management can have an important role to play in this regard. Second, centralization can help disparate units to work together by providing *coordination*. To the extent that an action taken by one unit may interact with the productivity of various actions of other units, superior organizational performance may require upper management to intervene and constrain the choices made by these units. Our discussion will focus on how various organizational forms influence these aspects of multi-unit search.

This section is organized as follows. Section 3.1 begins with a description of how an organization's search space is modelled. There are two general approaches: the *NK model* (which is imported from biology) and the *economic model* (which is built upon economic primitives). We then briefly discuss the cognitive requirements for a search unit exploring such landscapes as well as their implications for multi-agent search. The relevant literature is then divided into two broad classes. One class

¹⁴Burton and Obel (1980) is one of the pioneering efforts in using a computational model to understand the effect of organizational form. The authors compare the *M-form* and *U-form* as a function of the degree of decomposability in production technology; see Section 2.2 for definitions of these structural forms. Their model anticipated many of the crucial elements of organizational modeling considered in more recent papers reviewed in this chapter.

has all of the units of an organization engaged in similar operations and striving to solve similar (though perhaps not identical) problems. This is covered in Section 3.2. Examples include retail chains and multi-plant manufacturers. The second class, which is reviewed in Section 3.3, has the organizational problem segmented into distinct and dissimilar sub-problems which are distributed among the units who separately engage in search. The typical *U*-form organization is an example. The evolution of organizational designs is covered in Section 3.4. Finally, Section 3.5 distills some of the new insight and contrasts it with what a neoclassical economic approach delivers.

3.1 Modelling Search

Two approaches have been taken in modelling the search space faced by an organization. One approach is to assume the space of routines, over which an organization is searching, is a highly structured space; typically, it is a subset of Euclidean space with a metric that allows one to measure how “close” two routines are. Given this search space, a mapping from it to the real line is constructed which assigns performance to each routine. How this mapping is constructed varies significantly between the *NK* model and the economic model. A second approach involves less structure as its primitive is a probability distribution over the performance (say, profit) attached to an idea. Examples utilizing this approach are March (1991) and Chang and Harrington (1997). As the dominant approach is the first one, we will focus exclusively on it with the exception of discussing March (1991) in Section 5.2 due to its unique analysis of the evolution of organizational norms.

Agent-based models of organizational search characterize an organization by a fixed number of attributes. The search space for an organizational unit, frequently called a *landscape*, is defined on Euclidean space in which each attribute of an organization is represented by a dimension of the space and a final dimension indicates the performance of the organization. The organization’s attributes are indexed by the set $S \equiv \{1, 2, \dots, N\}$. For each attribute, there exists a fixed number of possible options which we will refer to as “practices” and which Nelson and Winter (1982) call “routines.” The practice of the organization in attribute $j \in S$ takes values in a non-empty set $Z_j \subseteq \mathfrak{R}$, where \mathfrak{R} is the set of all real numbers. Letting $A \equiv Z_1 \times \dots \times Z_N$, a vector defined in A then completely describes the organization’s practices. There is a metric $d : A \times A \rightarrow \mathfrak{R}$ which measures how “close” practices are to one another. Finally, to each vector of practices, there corresponds a level of performance for the organization as described by $v : A \rightarrow \mathfrak{R}$. The search spaces in the *NK* model and the economic model, to be discussed below, are two special cases of this general model.

A key factor in the organization’s search process is the exact shape of the landscape. Figure 1 shows two possible search landscapes for an organization which has two attributes with 15 possible practices for each. Figure 1(a) captures a smooth landscape having a unique local (and thereby global) optimum, while 1(b) captures a rugged landscape with many local optima. The shape of the landscape is typically determined by the way the organization’s various attributes interact with one another. How the interaction pattern affects the extent of ruggedness is discussed

below for both the NK model and the economic model.

NK Model Even though the NK model was initially conceived by Kauffman (1993) for understanding biological systems, it has been extensively applied in many other domains including computational organization theory. An organization is conceptualized as a system of activities. It makes decisions concerning N activities where each activity can take on two states, 0 or 1, so that, referring back to the general model, $A = \{0, 1\}^N$. A particular configuration of activity is then described by a binary vector of length N . The distance between two such vectors, $\underline{x} \equiv (x_1, \dots, x_N)$ and $\underline{y} \equiv (y_1, \dots, y_N)$, is captured by the Hamming distance:

$$d(\underline{x}, \underline{y}) = \sum_{i=1}^N |x_i - y_i|; \quad (1)$$

that is, the number of dimensions for which the vectors differ. As part of the NK model, the mapping v from the activity vector to the level of performance is a primitive. v is set to depend on the performance contributions that these activities make individually, where the contribution of each activity depends on the interactions among a subset of activities. The degree of interdependence among activities is captured by a parameter K which is the number of other activities that directly affect the contribution of a given activity. In its original formulation, these K activities are randomly selected from S for each activity.

To be more concrete, let $v_i(x_i, x_i^1, \dots, x_i^K)$ denote the contribution of activity i to the organization's performance where its dependence on activity i , x_i , and the K activities to which it is coupled, (x_i^1, \dots, x_i^K) , is made explicit. It is common to assume that the value attached to v_i is randomly drawn from $[0, 1]$ according to a uniform distribution for each possible vector $(x_i, x_i^1, \dots, x_i^K)$. The overall organizational performance is then

$$v(\underline{x}) = \left(\frac{1}{N}\right) \sum_{i=1}^N v_i(x_i, x_i^1, \dots, x_i^K). \quad (2)$$

Normalization by N enables performance comparisons when N is changed.

The interaction parameter, K , controls the difficulty of the search problem by making the value contribution of an activity dependent upon K other activities. When $K = 0$, the activities are completely independent so that changing the state of one activity does not affect the performance contribution of the remaining $N - 1$ activities. The landscape is then single-peaked so the globally optimal vector of activities is also the unique local optimum. That is, improving v_i by changing x_i must raise the organization's performance since the contribution of the other activities is unaffected by x_i . The other extreme is when $K = N - 1$ so that a change in the state of an activity changes the performance contributions of *all* other activities. This typically results in numerous local optima for $v(\cdot)$ due to the complementarity among activities. That is, changing any one of a collection of activities could lower v but simultaneously changing all activities could raise v . Kauffman (1993) shows that the number of local optima increases in K .

Rather than specify the coupled or interacting activities to be randomly selected, many organizational models using the NK framework choose the interaction pattern so as to explore how different architectures influence performance. For those purposes, it is convenient to capture the interdependencies in an *adjacency matrix* (Ghemawat and Levinthal, 2000). Figure 2 shows four such matrices for $N = 6$ in which the degree of interdependence as well as the exact structure of the interdependence differ. If the performance contribution of the j th activity (row j) is affected by the chosen activity in the i th activity (column i) then the element in the matrix corresponding to row j and column i has an ‘x’. This is always true of the principal diagonal as the contribution of an activity depends upon the practice chosen for that activity. Figure 2(a) is an adjacency matrix for an organization in which $K = 0$ so that the activities are completely independent. Figure 2(b) is when $K = 5$ and each activity is influenced by every other activity in S . Figure 2(c) captures a special case of $K = 2$, where the interdependencies are restricted to non-overlapping strict subsets of S ; the activities in $\{1, 2, 3\}$ influence one another, while those in $\{4, 5, 6\}$ influence one another. Figure 2(d) is another case of $K = 2$, though there is no obvious systematic structure in comparison to the other matrices. This is what would be typical if the interactions were random.

Economic Model The essence of the NK model is to build a generic landscape through a random construction process. In contrast, the economic model builds it systematically from a set of economic primitives. By way of example, let us describe the specification in Chang and Harrington (2000). Consider an organization – such as a retail chain – that consists of a corporate headquarters (HQ) and $M \geq 2$ units (such as stores). In this section, we will focus on constructing the landscape for only one of the stores and defer the discussion of the overall organizational search problem. As in the NK model, there exist N activities to a store’s operation. For each activity there are R possible practices so that $A = \{1, \dots, R\}^N$. A store is then characterized by a vector of N operating practices $(z_1, \dots, z_N) \in A$, where $z_i \in \{1, \dots, R\}$ is the store’s practice for the i th dimension. These practices influence the appeal of the store to consumers. The distance between any two vectors of practices, \underline{x} and \underline{y} , is measured by Euclidean distance:

$$d(\underline{x}, \underline{y}) = \sqrt{\sum_{i=1}^N (x_i - y_i)^2}. \quad (3)$$

Each consumer has an ideal vector of store practices which is an element of $\{1, \dots, R\}^N$. The net surplus to a consumer of type $\underline{w} \equiv (w_1, w_2, \dots, w_N)$ from buying q units at a price of p from the unit is specified to be

$$\left[\Gamma - \sqrt{\sum_{i=1}^N (z_i - w_i)^2} \right]^\gamma \cdot q^\beta - pq.$$

It is assumed that $\beta \in (0, 1)$, $\gamma \geq 1$, and $\Gamma - \sqrt{\sum_{i=1}^N (z_i - w_i)^2} > 1$. Having q take

its utility-maximizing value, a consumer's demand is

$$(\beta/p)^{\frac{1}{1-\beta}} \left[\Gamma - \sqrt{\sum_{i=1}^N (z_i - w_i)^2} \right]^{\frac{\gamma}{1-\beta}}.$$

The set of consumers in a market is represented by a cdf F defined on the space of consumer types, $\{1, \dots, R\}^N$, and is allowed to vary across markets so that the environment a store varies across stores. In Chang and Harrington (2000), additional structure is placed upon F as a consumer's type is assumed to lie in $\{(1, \dots, 1), \dots, (R, \dots, R)\}$ so that it can be represented by a scalar. This captures the idea that a consumer's preferences over the various dimensions are correlated so that, for example, a consumer who prefers value 3 for dimension 1 is likely to prefer value 3 for the other dimensions. The set of consumers in the market is represented by a triangular density function defined on $\{1, \dots, R\}$.¹⁵

Using the derived demand for a consumer and specifying the optimal price of c/β , a store's profit is:

$$v(\underline{z}) \equiv \left[\left(\frac{c}{\beta} \right) - c \right] \left(\frac{\beta^2}{c} \right)^{\frac{1}{1-\beta}} \int \left[\Gamma - \sqrt{\sum_{i=1}^N (z_i - w_i)^2} \right]^{\frac{\gamma}{1-\beta}} dF(\underline{w}) \quad (4)$$

The crucial property here is that a store's profit is decreasing in the distance between its practices and those desired by its customers. For a given store, the profit function defined above then represents its performance landscape over which it searches for better combinations of practices. As in the NK model, an important property of the landscape structure is its ruggedness. Here, the number of local optima can be shown to increase in γ , the consumers' sensitivity to store practices, as well as the degree of preference complementarity (Chang and Harrington, 2004). Unlike the NK model, for which the level of complexity is directly specified by the interaction parameter K , the economic model allows the complexity in a decision problem to result from more fundamental economic primitives.

Modelling Search by a Single Agent The potential for multi-agent search to outperform single-agent search on a per-agent basis derives from its capacity to overcome the bounded rationality of individual agents through sharing and coordination in the search process. Two forms of bounded rationality stand out in these models: a lack of information about the search space (landscape), and a constraint on the considered set of alternatives to the status quo.

If a unit has full information about the mapping from practices to performance then search is irrelevant as the organization can simply identify and adopt the practice vector that corresponds to the global peak. Almost all agent-based models instead

¹⁵Chang and Harrington (2004) relax the assumption of a perfect correlation in a consumer's preferences over dimensions and examine how the degree of preference complementarity affects the relationship between an organization's structure and its performance.

take the view that agents are largely uninformed and assume the other extreme - nothing is known about the shape of the landscape - so that agents must resort to blind search using some form of hill-climbing algorithm. The myopic but adaptive search on the landscape entails considering a practice vector that is different in several dimensions from the organization's current one - the change may involve as few as one and as many as all dimensions.

This forces us to confront the second form of bounded rationality. To what extent is the organization capable of considering different changes? Is it capable of contemplating a major change in its operation which involves changing practices in all dimensions? Or, is it constrained to considering only minor modifications? The ruggedness of the landscape - which is determined by K in the NK model and partially determined by γ in the economic model - turns out to affect the efficacy of search. When the landscape is smooth and single-peaked, constraining the breadth of change that an organization considers has no impact on the optimum eventually attained - as any hill-climbing algorithm will find the global optimum - though it will influence the speed of convergence and thus intermediate-run performance. This form of bounded rationality does make a difference, however, when the landscape is rugged. While an organization capable of carrying out transformations involving all dimensions will still eventually attain the global optimum,¹⁶ an organization which is only capable of considering changes involving a small subset of the dimensions may become trapped on an inferior local optimum.

Central to the search-and-learn perspective of organizational theory is this dynamic interaction between a boundedly rational search unit and the structure of its search space, which serves to restrict the set of search paths and outcomes that the unit is capable of achieving. The organization as a multi-agent search mechanism can overcome such restrictions through the sharing of their discoveries and internal coordination.¹⁷

3.2 Organizational Search with Units Solving Similar Problems

Examples of organizations in which various units are solving similar problems include retail chains, multi-plant manufacturers, and manufacturers producing a line of related products. Such a situation is modelled by endowing each unit with a performance landscape over which it searches. All of the landscapes are defined over the same space of activities and similarity between units' problems is reflected in the similarity of the landscape, that is, how activities map into performance. Given that units are searching over similar landscapes, the possibility of knowledge transfer among units is significant. The main organizational issue here is how inter-unit

¹⁶It should be noted, however, that it may take a very long time for the organization to find such a global optimum by chance when N is relatively large.

¹⁷Both Levinthal (1997) and Rivkin (2000) consider the impact of this interaction on the Darwinian selection process in a population of firms climbing an NK landscape. Levinthal (1997) examines how successful firms with tightly coupled systems (high K) find adaptation difficult in the face of environmental change, while Rivkin (2000) allows imitation among firms and focuses on how tight coupling protects a successful firm from potential imitators. It should be noted that both papers are restricted to single-agent models of an organization.

learning can be promoted through an appropriate organizational structure.

Kollman, Miller, and Page (2000) Recognizing the possibility of multiple searches as the central benefit from decentralization, Kollman, Miller, and Page (2000), hereafter KMP, consider four factors affecting the magnitude of this benefit: 1) difficulty of the problem; 2) sophistication in search; 3) heterogeneity among unit preferences; and 4) organizational size. Extending the NK model into the multi-unit organizational setting, KMP endow each unit with an NK search space which is common for all units (including the central authority).¹⁸ Search involves myopic hill-climbing on a fixed landscape. The objective is to investigate the efficacy of a multi-unit organization in searching for solutions in parallel by exploiting units' search capacities and combining the revealed information to the benefit of the entire organization.

Four types of organizational forms (or search rules) are considered: 1) full centralization in which search is carried out solely by the central authority and the best policy found is mandated for all units – hence, this is equivalent to single agent search; 2) full decentralization in which each unit searches independently and makes its own adoption decision (so that there is no inter-unit spillovers of knowledge); 3) partial decentralization with “best adoption” which means that each unit searches on its own but, after a fixed number of search periods, the central authority mandates the best policy discovered; and 4) partial decentralization with “incremental adoption” which means that each unit searches on its own for a given length of time and then the central authority forces the units to change policies incrementally (attribute by attribute) toward the best known current policy so that, ultimately, all units have the same policy.¹⁹ The potential trade-off between centralization and decentralization is that the former may draw from a better distribution while the latter has multiple units searching. Under each of these organizational rules, KMP examines the impact on the organization's performance of the four previously mentioned factors.

A focus of their analysis is to understand the relationship between the complexity of the environment - measured by K in the NK formulation - and the cognitive constraints of the organization's units which are represented by the maximum number of dimensions, denoted z , along which a new idea can depart from the status quo policy. To begin, the benefits from decentralization are always positive under the best adoption rule when the units are as capable as the central authority. There is also an interaction between problem difficulty and the benefits of decentralized search because the greatest advantage occurs with a moderate level of difficulty. Even a single unit can do very well when problems are simple, so having more units searching in this case is of little value. When problems are very hard, each unit tends to get stuck on a local optimum of similar value (as the peaks become more numerous with

¹⁸The central authority in this setting is just another unit carrying out the search for the organization, though it may have superior search capability.

¹⁹In this case, the target policy – that is, the “best-to-date” policy – could change along the adoption process, since the organization-wide switching of unit policies takes place one attribute at a time. This is to be contrasted to the “best adoption” rule under which all units immediately adopt the best policy in its entirety, while discarding everything that they have found individually through local search.

more similar values as K increases) so once again organizational structure doesn't matter. It is when the problem is of moderate difficulty that the additional search under decentralization makes a substantive difference.

In comparing the two partial decentralization rules, KMP find that the incremental adoption rule always outperforms the best adoption rule. This is due to the fact that the units are allowed to keep in place what has worked for them, while simultaneously allowed to try what has worked elsewhere in the organization. This blending of diverse local solutions proves superior to the alternative of requiring all but one unit to discard the knowledge they accumulated. This comparative advantage of incremental adoption is found to be non-monotonic in the difficulty of the problem. When K is low, the probability of any one unit finding the global optimum is relatively high and, therefore, the advantage of incremental adoption is minimal. And, when K is high, there are many local optima which tend to be uncorrelated so that blending them together has little value and, like any random change, generally proves deleterious. In other words, the activities identified as worthwhile by one unit (that is, are at or close to a local optimum) are unlikely to be of much value to another unit that is targeting a different optimum because these different optima could be vastly distant from one another. Once again, it is for moderately difficult problems that incremental adoption does significantly better than best adoption. Finding the global optimum is then not easy, and information associated with one local optimum is still of value to units that are at another local optimum as it may allow them to move to yet better local optima.

Chang and Harrington (2000) The focus of this work is to explore the relationship between organizational structure – specifically, the degree of centralization – and firm performance. The case of a single chain with multiple local stores is analyzed in Chang and Harrington (2000). The model is then extended in Chang and Harrington (2003) to allow for competing chains and searching consumers, thereby enabling an investigation of the coevolutionary dynamics among organizations, units within an organization, and consumers in heterogeneous markets.

Chang and Harrington (2000) consider a retail chain consisting of M stores, each with a performance (profit) landscape defined by equation (4). The heterogeneity in the markets that the stores serve is captured by differences in the distributions of consumer types. Organizational profit is the simple sum of its stores' profits. While stores' landscapes may be similar, they are independent in that a choice made by one store does not affect the profit earned by another store. However, as explained below, inter-unit learning creates a dynamic and endogenous linkage among stores' search paths and profits.

Search over the profit landscape takes place through an iteration of myopic one-step hill-climbing, where a new idea is represented as a point in store practice space. In each period, each store possesses a vector of current practices and generates one idea where an idea is created by randomly selecting a dimension from $\{1, \dots, N\}$ and assigning to it a randomly selected element from $\{1, \dots, R\}$. If it is adopted then the store's practice in the specified dimension is changed to the new value.

Two organizational forms are considered in this setting. In the decentralized organization, a store manager evaluates his own idea and the ideas adopted by other stores in the current period. A store manager sequentially evaluates these ideas and adopts an idea if it raises *store* profit. Hence, each store manager searches over his store’s landscape and has the authority to implement any useful ideas. This is equivalent to KMP’s full decentralization, except that inter-unit learning is voluntary. In a centralized organization, a store manager once again generates an idea and considers whether, if adopted, it would raise store profit. If so, the idea is passed to *HQ*. If not, the idea is discarded. With this set of ideas, *HQ* sequentially evaluates them in a myopic manner, mandating a practice throughout the chain if doing so raises *chain* profit, and otherwise discarding the idea. Thus, uniformity of practices is a feature of centralization in this model. *HQ* then searches over its landscape which is based on chain profit, and it alone has the authority to implement ideas.

Measuring performance by average chain profit, the main insight of this study is that centralization can outperform decentralization. This occurs when markets are sufficiently similar, the horizon is sufficiently short, and consumer preferences are sufficiently sensitive to store practices relative to price. Given that markets are heterogeneous, the benefit of decentralization is clear – it allows each store manager to tailor practices to its market. So, how can a centralized structure generate higher profit? It turns out there is an implicit cost to decentralization. By adjusting practices to one’s own consumers in a decentralized chain, stores’ practices tend to drift apart. As a result, a new practice adopted by one store is increasingly unlikely to be compatible with the current practices of other stores. In essence, stores come to target distinct consumers (by targeting distinct local optima) and what works for one type of consumer doesn’t tend to work for another type of consumer in light of preference complementarities. Inter-store learning is then less under decentralization and this is detrimental to the rate of improvement in store practices. The virtue of a centralized structure is that it enhances inter-store learning by keeping stores close in store practice space so that they are targeting similar consumers. With these two countervailing effects, a centralized structure outperforms as long as markets are not too different. The value to enhanced inter-unit learning is greatest when stores are farther from local optima and for this reason centralization does particularly well in the short-run. In the long-run, decentralization is typically superior because the uniformity of practices under centralization prevents the global optimum being achieved since the global optimum has different practices in different markets. Finally, centralization also outperforms when consumers are sufficiently sensitive to store practices (γ is high). This result is related to the property that the ruggedness of the landscape increases in γ . As the number of local optima rises, stores in different markets (and thereby different landscapes) are more likely to share some common local optima. This enhances opportunities for inter-store learning and the analysis shows that this is best exploited by a centralized organization.²⁰

A changing environment is encompassed by allowing the population of consumer

²⁰The robustness of these results with respect to the shape of the landscape is explored in Chang and Harrington (2004).

types to shift probabilistically. Measuring performance by steady-state chain profit, centralization is more likely to outperform when market fluctuations are sufficiently large. Recall the earlier result in the static environment that centralization is favored in the short run because stores are farther away from local optima, in which case inter-store learning is especially valuable. As increased fluctuations in market environments shake the landscapes more vigorously, they act to push stores further away from local optima. Thus, a constantly fluctuating environment requires the firm to perpetually learn at a high rate, which then sets the stage for the short-term superiority of centralization to become a long-term advantage. Quite contrary to the received wisdom that volatility in markets require greater decentralization, Chang and Harrington (2000) find it is the centralized organization that is more effective in responding to change.

Chang and Harrington (2003) A more challenging issue is to consider how market structure interacts with organizational structure to influence the dynamic performance of chains. Does increased competition make centralization more or less desirable? To address this issue, Chang and Harrington (2003) modify the previous model by allowing for competition and consumer search. There are L chains and M markets with each chain having a store in each market. Within each market, there is a fixed population of consumers that engage in search by moving among stores. At any point in time, a consumer in a given market (served by L stores) has a favorite store and buys from it with probability $1 - Q$. With probability Q the consumer experiments by randomly selecting another store and buying from it. If the resulting surplus for the consumer is higher than what the consumer received most recently from the favorite store then this new store becomes the consumer's favorite store. If not, then the consumer's favorite store remains unchanged and, in the next period, the process is repeated. Q regulates the extent of experimentation. If $Q = 0$ then there is no competition as consumers are permanently loyal, while $Q = \frac{L-1}{L}$ implies no loyalty. The organizational structures are as before. A store evaluates the profit attached to adopting a new idea using its current base of consumers - those that are currently buying from it. In a centralized organization, HQ evaluates ideas using a measure of profit based on the current sets of consumers at its stores.

A key result is that centralization is more attractive when there is a larger number of competing chains and may even outperform in the long-run. The basis for this finding is an implicit increasing returns mechanism when competing organizations are coevolving with consumers. To understand this result, recall that centralization does particularly well in the short-run. Thus, early on a centralized chain is developing better practices and thereby attracting more customers than a decentralized chain. In the one-chain model, decentralization would eventually outperform, but that needn't be true when consumers are searching. This early advantage from centralization establishes a customer base which tends to include the most prevalent consumer types in the market, and it is this customer base which is used to evaluate the profitability of new ideas. A centralized organization then tends to adopt practices well-suited for the prevalent consumer types, which results in their retention and the

attraction of more of those types and which makes the chain even more inclined to adopt ideas suiting their preferences, and so forth. In this way, an early advantage of centralization - coming from enhanced inter-store learning - is fed into a feedback loop to maintain an advantage in the long run. As a result, a decentralized chain may not be able to catch up because it is adopting ideas for a less prevalent niche of consumers. In other words, the rate at which a chain climbs a landscape (by coming up with better practices for its current customers) influences the shape of its future landscape (by affecting the set of loyal customers). A centralized chain climbs its landscape faster and this results in its future landscape being more attractive. Coevolutionary dynamics among firms and consumers produce a powerful increasing returns mechanism.²¹

3.3 Organizational Search with Units Solving Different Problems

The previous section is applicable when the organization is divided into units solving similar problems such as selling a particular product line to consumers (retail chains) or producing a particular product line (multi-plant manufacturers). Such organizations are examples of the *M*-form, but let us now consider the *U*-form organization. The organization's various activities are allocated among functional departments such as Accounting, Finance, Sales, Purchasing, Production, and so on. A new practice adopted in Sales is unlikely to be applicable to the operation carried out in Finance - they are engaged in entirely different types of operations and thereby solving quite distinct problems. However, it *will* have an impact on the effectiveness of the overall operation of Finance when the value of certain financial practices depends on sales practices; that is, there is a complementarity between them. These organizational issues can be modelled by specifying the firm as a system of N activities in the context of the NK model but with the feature that these N activities are allocated to various departments for specialized search occurring in parallel. For instance, half of the activities may be put under the control of department *A* while the remaining activities may be under the control of department *B*, with each department attempting to find the optimal configuration of decisions over the activities it controls according to some evaluation criterion. As the departments are then searching over distinct non-overlapping set of activities, there is no prospect for inter-unit learning. Rather, the issue is how to structure the organization so that the gains of parallel search can be had while balancing it with the need to coordinate search in light of how these activities interact.

Rivkin and Siggelkow (2003) A long line of scholars studying complex organizations have observed that there are many interdependencies among elements of design such as allocation of decisions, incentives, and information flows. Rivkin and Siggelkow (2003) offer as one source of such interdependencies two conflicting

²¹As an example of how analytical and computational methods are complementary, this issue is explored analytically in Harrington and Chang (2005) as they consider a highly stripped-down version of Chang and Harrington (2003).

needs of a multi-unit organization that are central to the search-and-learn perspective. First, to be successful, an organization must search broadly for good actions (exploration). Second, it must also stabilize around good actions once discovered (exploitation). An effective organization balances search and stability. The authors focus on three prominent elements of organizational design in exploring how they interact to influence this delicate balance: 1) a central authority that may choose to review the proposals sent up from subordinates; 2) an incentive system that influences the degree to which managers act parochially for the good of their departments or for the good of the overall firm; and 3) the decomposition of an organization’s decisions into distinct departments. Their focus is on how these design elements interact with one another to determine organizational performance through the balancing of search and stability and how that relationship depends on the interdependent structure of activities as dictated by the problem and on the limits on the cognitive ability of managers.

Their simulation considers a hierarchy with a CEO and two subordinate managers, A and B. The firm engages in multi-agent search which takes place on performance landscapes generated by the NK model. An organization has $N = 6$ decision attributes and part of its design is how they are allocated among the two managers. Manager A has responsibility for a subset S_A of these attributes and manager B for the complementary subset S_B . In each period, each subordinate manager reconsiders the actions assigned to its attributes by comparing the current configuration to some fixed number α of alternatives, so that α reflects the cognitive capacity of a subordinate manager. These α alternatives are ranked by a manager on the basis of an evaluation criterion which is a weighted average of the performance of his department and of the other department.

Initially, it is supposed that $S_A = \{1, 2, 3\}$ and $S_B = \{4, 5, 6\}$. Denoting by $\delta \in [0, 1]$ the degree to which Manager A cares about the other department’s performance, the evaluation criterion for Managers A and B, respectively, are

$$v^A = \frac{v_1 + v_2 + v_3 + \delta(v_4 + v_5 + v_6)}{6}, \quad v^B = \frac{\delta(v_1 + v_2 + v_3) + v_4 + v_5 + v_6}{6},$$

where recall that v_i is the contribution of the i th activity to total organization performance. If $\delta = 0$, a manager only cares about his own department, while if $\delta = 1$ he cares about firm profit. δ then controls the degree to which managers’ incentives are aligned with those of the organization.

Finally, the form of vertical hierarchy and the ability of the CEO affect the organizational search process. From the status quo and the α alternatives, a manager sends up the best P proposals to the CEO where “best” is according to the manager’s preferences. There are two types of CEO’s: rubberstamping (decentralization) and active (centralization). The first type always approves all proposals sent up by both managers so that, effectively, an organization with a rubberstamping CEO is decentralized since the real authority lies with the department managers. The active CEO, on the other hand, selects β proposals from all combinations of the submitted proposals and implements the one that generates the highest firm profit (so $\delta = 1$ for

the active CEO). Thus, β captures the cognitive capacity of the CEO. Since an active CEO has the final authority, we will refer to this as the *centralized organization*.

In sum, there are five different factors that affect the organizational search process and, consequently, performance: the grouping of activities into departments, the amount of information sent up to senior management (P), the allocation of authority (centralization/active CEO vs. decentralization/rubberstamping CEO), the extent to which managers care about firm as opposed to department performance (δ), and the cognitive abilities of the department managers (α) and the CEO (β).

There is found to be a significant interaction between the allocation of authority and the complexity of the environment (as measured by K). When the complexity is low ($K = 0$), the benefit of centralization is non-existent since the lack of interdependencies means there is no need for coordination while, at the same time, there is a cost due to slower adaptation. In such a case, short-run performance is lower under centralization. When complexity is moderate, centralizing authority in the CEO is shown to enhance performance as the interdependence among activities makes coordination critical. But then for highly complex environments (high K), it is better to push authority back down to the managers. Centralization suffers from the problem that an active CEO is always moving the organization to points of higher firm profit and, when K is high, there are many bad local optima. As a result, the organization is typically getting stuck at a point of low performance. In essence, centralization results in excessive stability. In contrast, a decentralized organization - by giving authority to department managers who care more about their own department's profit - may periodically result in organizational performance deteriorating which, when it causes movement into a basin of attraction for a better optimum, can enhance long-run performance. This weakness to centralization can not be mitigated by increasing the skill of the CEO (as measured by β), but only by increasing the information flow, P . In sum, centralization is undesirable when interactions are pervasive and the CEO gets little information from below.

The skills and incentives of the subordinate managers have some subtle and surprising effects. In a complex environment, highly skilled managers can be harmful in a decentralized organization. By considering alternatives that are far away from their current position, a highly skilled manager may undermine the improvement efforts of other managers. The organization can suffer from excessive instability as it dances around the landscape without making much progress. Centralizing authority in the CEO provides useful coordination. An active CEO and skilled department managers are then complements, not substitutes. Managerial incentives that are more closely aligned with the interests of the firm are complementary to centralization as well. When managers are parochial (low δ), many of their suggestions are turned down in a centralized organization because the CEO uses a different criterion in evaluating them. Hence, the organization doesn't make much progress. This is contrary to the usual argument which is that, if managers have the right incentives, why does one need an active CEO? Here, the problem is that departmental managers have partial information and control and one needs the coordination that centralization delivers.

The above results are obtained for landscapes created using the usual random

interaction NK model. Rivkin and Siggelkow (2003) also considers the interdependence between decomposition and the allocation of authority. With decomposable interactions - as represented by the block-diagonal adjacency matrix in Figure 2(c) - centralizing authority is irrelevant since department managers are solving independent problems. There is no need for coordination. Superior performance can, however, come from the combination of imperfect decomposition - there is some interdependence across departments - and an active, well-informed CEO. For instance, given a block-diagonal matrix (Figure 2(c)), performance is higher when an active CEO is combined with $(S_A, S_B) = (\{1, 2, 6\}, \{3, 4, 5\})$ than with $(S_A, S_B) = (\{1, 2, 3\}, \{4, 5, 6\})$. At work is the balancing of search and stability. Some overlap expands the range of search as each manager proposes options that change the landscape faced by another department. This may serve to move the organization to a different basin and, in some cases, result in it homing in on a superior local optimum.

Siggelkow and Levinthal (2003) Using a model similar to the preceding one, Siggelkow and Levinthal (2003) examine the division of task and specialized search under three different organizational forms: centralization, decentralization, and reintegration. In the centralized firm, decisions are made only at the level of the firm as a whole, whereas a decentralized organization is disaggregated into a number of departments in which decisions are made independently. A reintegrated organization initially has a decentralized structure and then switches to centralization after a fixed number of periods (typically, 25 periods). A key variable is the degree and pattern of interactions among various activities as specified by an adjacency matrix. The decision problem for the organization is decomposable if the activities can be grouped so that all interactions are contained within each group and thus there are no cross-group interdependencies. The block-diagonal adjacency matrix in Figure 2(c) is a decomposable system. On the other hand, the decision problem is non-decomposable if there is no way to group the activities so as to eliminate all cross-group interdependencies; see, for example, the matrix in Figure 2(d).

The simulation entails creating 10,000 landscapes using the NK model with $N = 6$. The three organizational forms are compared in terms of their performance (averaged over the 10,000 landscapes) under conditions of both non-decomposability and decomposability of the decision problems. Firms carry out myopic local search and they only consider changing one activity at a time. The centralized firm evaluates an idea on the basis of firm profit: $v = \frac{v_1+v_2+v_3+v_4+v_5+v_6}{6}$. The decentralized firm is assumed to have two departments, A and B , with department A controlling activities $\{1, 2, 3\}$ and department B controlling $\{4, 5, 6\}$. In each period, each department comes up with an idea which it then evaluates on the basis of the profit contribution of those activities that are under its exclusive control. This means that the evaluation criteria used by departments A and B are $v^A = \frac{v_1+v_2+v_3}{3}$ and $v^B = \frac{v_4+v_5+v_6}{3}$, respectively. In evaluating an idea, a department takes the other department's current choices as given.

In a decomposable environment with a block-diagonal interaction structure (Fig-

ure 2(c)), they find that the decentralized firm outperforms the centralized firm in the short-run. This result is directly due to the asymmetric number of draws that are allowed under these two forms: the decentralized firm gets two draws per period (one for each department), while the centralized firm gets only one. As there is no interaction between the activities of the two departments, there is no mitigating benefit from centralization. The average levels of performance under these two forms do converge in the long run, however. The reintegrated firm's performance is nearly identical to that of the decentralized firm.

The results are quite different when the organization searches in a non-decomposable environment. Assuming a random interaction structure with $K = 2$, the advantage of having more draws under decentralization is offset by the coordination benefit attained under centralization due to the presence of cross-departmental interdependencies. More interesting is the performance of the reintegrated firm. Prior to reintegration, the performance is, of course, the same as that of a decentralized firm. After the departments are integrated, performance not only improves but it eventually outperforms the centralized firm. The problem with the organization when it is centralized is that it is apt to get stuck early on at an inferior local optimum, similar to the active CEO structure in Rivkin and Siggelkow (2003). This is less likely with the reintegrated firm as it is initially decentralized. Once centralization occurs, it is more likely to be in the basin of a better optimum which it can take advantage of now that coordination can occur. The lesson is that superior performance may be had by a temporal blending of different organizational forms.

Those simulations assume the organization starts its search from a random point on the landscape. An alternative exercise is to suppose there is an environmental shock after the firms have achieved some steady-state. Siggelkow and Levinthal (2003) position a firm at Hamming distance d from the global optimum - implying that the firms were at the global optimum *ex ante* and then were thrown off it by a shock of magnitude d . In this setting, the question is how effectively a firm can *climb back* to the global optimum. Centralization outperforms reintegration for sufficiently low values of d , while reintegration outperforms centralization for sufficiently high values of d . The appropriate organizational form then depends on the size of the shock. The intuition is that a centralized firm has a relatively high probability of getting locked onto nearby local optima which makes it less suitable for large shocks but quite desirable for small shocks since the firm is likely to start in the basin of attraction for a good optimum (recall that the firm started at the global optimum). By comparison, reintegration initially pursues a decentralized form and thus can better handle large shocks. The general lesson is that an organization should be centralized at a steady-state but should temporarily decentralize when there is a large change in its environment.

The preceding results suggest that there may be merit to grouping activities so that there is some cross-departmental interdependence even when the decision problem is decomposable. Suppose the interaction structure is characterized by the adjacency matrix in Figure 2(c). An obvious grouping of activities would be to have department A in charge of $\{1, 2, 3\}$ and department B in charge of $\{4, 5, 6\}$, thereby

eliminating any interaction between the activities controlled by these two managers. However, such a structure underperforms one which is eventually of that form but during the early periods has A controlling $\{1, 4, 5\}$ and B controlling $\{2, 3, 6\}$. Quite interestingly, the temporarily scrambled firm is superior to the “ideally” decomposed firm because cross-departmental interdependence avoids excessive stability.

3.4 Evolving an Organizational Structure

Thus far the focus has been on comparing the performance of different elements of organizational design. This begs the question of whether upper level management of an organization, which is endowed with a sub-optimal design, can effectively alter design elements so as to achieve a superior structure. What makes this a non-trivial problem is the presence of interdependence among component tasks, which is representative of any complex system, be it social, biological, or technological. The significance of this problem is well illustrated by Herbert Simon in the context of organizations:

The basic idea is that the several components in any complex system will perform particular subfunctions that contribute to the overall function. ... To design such a complex structure, one powerful technique is to discover viable ways of decomposing it into semi-independent components corresponding to its many functional parts. The design of each component can then be carried out with some degree of independence of the design of others... There is no reason to expect that the decomposition of the complete design into functional components will be unique... Much of classical organization theory in fact was concerned precisely with this issue of alternative decompositions of a collection of interrelated tasks. [Herbert A. Simon, *The Sciences of the Artificial* (1996), p. 128]

In a decomposable system such as the one in Figure 2(c), the obvious division of tasks would entail assigning activities $\{1, 2, 3\}$ to one department and $\{4, 5, 6\}$ to another. As there is no interdependence between the sets of activities of these two departments, the optimal solution they arrive at independently will form the optimal solution for the entire organization. Alternatively, systems may have inherent “near decomposability” where they can be decomposed into a collection of subsystems with the property that the components within a subsystem interact more strongly than the components belonging to different subsystems, but with a certain degree of interdependence remaining between the subsystems. In such situations, the problem solvers facing computational constraints will be motivated to decompose the problem into subproblems in order to benefit from parallel processing, while recognizing that the problem may not be decomposable.

Ethiraj and Levinthal (2002) Define an organization’s “true architecture” to be a description giving the correct number of the organization’s modules and a correct assignment of functions to the respective modules as dictated by the characteristics of

the problem. Ethiraj and Levinthal (2002) set out to identify the relationship between two key design elements - decomposability and hierarchy - and an organization's ability to *discover* its true architecture.

They consider the following four structural types: 1) hierarchical and nearly decomposable; 2) non-hierarchical and nearly decomposable; 3) hierarchical and non-decomposable; and 4) non-hierarchical and non-decomposable. Figure 3 presents the adjacency matrices of the systems that belong to each one of these categories when $N = 9$ and there are three non-overlapping modules labelled a , b , and c . Figure 3(a) is nearly decomposable and hierarchical as $b1$ in module b is influenced by $a3$ in module a and $c1$ in module c is affected by $b3$ in module b but module c does not influence modules b or a and module b does not influence module a . Hence, the inter-module interdependencies are unidirectional. Figure 3(b) is nearly decomposable and non-hierarchical in that modules a and b are mutually interdependent (through $b1$ and $a3$), while modules b and c are mutually interdependent (through $c1$ and $b3$). Figure 3(c) is a non-decomposable but hierarchical system as there is a tight coupling between modules in that all components of modules b (c) are influenced by all components of module a (b) and are unidirectional. Finally, a non-decomposable and non-hierarchical system is captured in Figure 3(d), where all modules are tightly and mutually coupled with one another. For each of these four structures, search for the true architecture occurs through three operations: splitting, combining, and re-allocation. *Splitting* of modules involves breaking up existing departments into two or more new departments. *Combining* is the opposite of splitting in that it involves integrating two or more departments. *Re-allocation* is when the organization reassigns functions from one unit to another.

Suppose the module designer observes the presence or absence of interactions among attributes within the module as the result of a change in an attribute. All attributes for which their contribution to performance is unaffected by this change are identified as not belonging to the module that includes the original attribute. All such attributes are either transferred to a randomly chosen different module (if they constitute less than half the total number of attributes in the current module) or are split into a new module (otherwise). If the change of the given attribute does not affect any other attributes within the module, then the attribute is viewed as not belonging to that module. In this case it is transferred to another randomly chosen module. In each period, the module designers also consider combining each module with another module by randomly selecting two modules and evaluating the impact of attribute changes in both modules. The modules are combined if changes in each module affect the other and remain separate otherwise.

Each module engages in one-step offline search based on local module performance. This occurs in parallel. When considering a population of systems in order to explore recombination of systems or substitution of modules, they select two systems at random and then select two functionally equivalent modules at random for recombination. The lower performing module is replaced with the higher performing module. Finally, in the multi-systems analysis, the selection mechanism used is the roulette wheel algorithm, where the probability that a system is selected equals

its performance level divided by the sum of the performance of all systems in the population at that time.

Consider a system with N attributes for which the true architecture has M modules with each module having an equal number of attributes. The initial design of the system is random and thus is likely to have the wrong number of modules, modules with the wrong attributes, and modules with different numbers of attributes. The performance measure is the number of periods it takes for the system to converge to the true architecture.

The simulation exercise is based on 100 experiments, where each experiment involves a randomly selected landscape and initial design and entails each of the four archetypes being run. The simulations show that an organization always discovers the true structure when the system is hierarchical, even when it is non-decomposable. But when it is non-hierarchical, an organization never manages to reach a stable state. The violation of both principles - hierarchy and decomposability - is seriously detrimental to discovering the right structure. These results suggest that the search rule for discovering the true system structure is robust when there is a strong interaction within modules and there is a hierarchical precedence structure underlying between-module interactions.

3.5 What Do We Learn from a Computational Agent-Based Approach?

The primary issue explored in agent-based models of organizational search and learning is the role of organizational structure and, more specifically, how a centralizing authority can influence performance by coordinating certain activities. In this section, we want to review what we've learned about when an organization should be centralized, highlight the role played by the unique features of ACE models identified in Section 2, contrast this insight with what a NCE analysis would produce, and make the case for ACE.

One important insight is that decentralization can be advantageous even when complementarities suggest that coordination is valuable. Consider an organization in which there are interdependencies across units. If each unit evaluates a new practice based upon what it generates in terms of unit performance, then decentralized search can lead to lower organizational performance due to externalities across units. A NCE analysis would suggest that centralization is beneficial because it internalizes these externalities by evaluating the impact of a new practice in terms of organizational profit. In contrast, Rivkin and Siggelkow (2003) show using an ACE model that centralization can perform worse because it results in excessive coordination. Once a centralized organization is in the basin of attraction of a particular local optimum, it steadily marches towards it and, as a result, it never learns whether there are other more attractive optima. Under decentralization, individual units - each of which is engaging in hill-climbing using the unit's performance - can inadvertently result in organizational performance declining. Though detrimental in the short-run, it may serve to throw the organization into the basin of a different and potentially better optimum. Put differently, the high level of coordination achieved under centralization

leads to excessive stability. Though stability is desirable once a good optimum is reached, it can be harmful while learning because it closes off alternatives. In the context of adaptive search - as opposed to optimal selection of organizational practices - coordination can be excessive. Second best arguments are rampant within ACE models and this is one example - the limitations of adaptive search may mean that fully internalizing externalities across agents can be detrimental, an intuition quite contrary to what would emerge from an NCE analysis.

A second important insight is that centrally mandated uniform practices can be valuable even when units face heterogeneous environments. Consider an organization in which there are no interdependencies across units. Each unit is in a different environment and organizational performance is the simple sum of the units' performances. An NCE analysis would suggest that decentralization is preferable as it allows practices to be tailored to the environment. However, Chang and Harrington (2000) show that a decentralized organization creates dynamic externalities related to knowledge transfer which impact adaptive search. Since units are solving similar problems, what one learns and adopts may prove useful to other units. Under decentralization, units fail to internalize the following externality: when a unit adopts a new practice that moves them away from other units, those other units can expect to learn less from it. A centralized organization serves a coordinating function by keeping units' practices close to one another, and this enhances knowledge transfer. Note that this result is produced by medium-run dynamics. In the long run the organization will typically achieve its global optimum and, since the global optimum is lower when constrained to uniform practices, decentralization outperforms in the long-run.

A unique feature of ACE models mentioned in Section 2 is the complexity of the environment, and this indeed played a central role in the preceding analysis. Complexity is measured by the ruggedness of the landscape. A more rugged landscape means more optima, in which case it becomes easier to get stuck on poor optima. Chang and Harrington (2000) show that a more complex environment makes knowledge transfer more important as it is more difficult for a unit, learning on its own, to succeed. This implies that more complexity means centralization is more likely to be preferred. In Kollman, Miller, and Page (2000), organizational form matters only when environments are moderately complex. In Rivkin and Siggelkow (2003), the potential advantage to the enhanced coordination from centralization increases with complexity (which is associated with more interdependencies) but the chances of getting stuck at a bad optimum also increases with complexity. When the environment is moderately complex, the first effect dominates so that centralization performs better but, when the environment is very complex, the second effect dominates so decentralization outperforms.

The above discussion reveals that ACE delivers different insight than would an NCE analysis. Furthermore, in reviewing NCE research on organizations, the forces at work are quite distinct. In one class of NCE models, organizational structure affects the incentives of lower-level agents to produce useful information for higher levels. In Aghion and Tirole (1997), decentralization promotes lower-level agents' incentives to invest in acquiring information - as their decision is less likely to be

overruled (and there is little value to investing in information if the information doesn't make a difference) - but at the cost of them pursuing their own interests which are distinct from the interests of the organization. In Dessein (2002), the problem is that lower-level agents may distort the information that they pass along to higher levels. A second class of models focuses on how organizational structure influences monitoring, wages, and the incentives for agents to work hard. In Qian (1994), a more hierarchical organization (which means more levels and each manager has fewer agents to monitor below him) enhances monitoring and lower wages but is less productive. Maskin, Qian, and Xu (2000) compare the *M*-form and *U*-form with respect to their productivity in monitoring when pay is based on relative performance. The emphasis, the forces, and the insight of these organizational models are then quite different from ACE organizational models. While ultimately these alternative approaches may compete, thus far their analyses are complementary.

4 Information Processing

Economists have also often failed to relate administrative coordination to the theory of the firm. For example, far more economies result from the careful coordination of flow through the processes of production and distribution than from increasing the size of producing or distributing units in terms of capital facilities or number of workers. Any theory of the firm that defines the enterprise merely as a factory or even a number of factories, and therefore fails to take into account the role of administrative coordination, is far removed from reality. [Alfred Chandler, *The Visible Hand: The Managerial Revolution in American Business* (1977), p. 490]

As reviewed in the previous section, search and learning models of organizations have agents receive new ideas, evaluate them, and then decide what to do – whether to discard them, pass them along to a superior, or implement them (depending on the allocation of authority). An implicit assumption is that the evaluation process is costless and instantaneous. This is a striking departure from reality. It can take resources, time, and expertise to evaluate new information and then make a decision. This section considers the costs of processing information. An organization takes input from the environment (“data”) and performs operations on it prior to making a decision. Information processing is costly because, for example, it requires hiring agents and it imposes delay in reaching a decision under the constraint of avoiding information overload. Though all models of organization involve information processing to some degree, we have reserved this terminology for those models where the cost of processing is explicitly modelled and is a primary force determining organizational performance.

The organization is faced with a task which, if it were to be handled by a single agent, would translate into long delays and inaccuracies due to processing and memory constraints. A more efficacious structure involves distributed problem-solving - multiple agents solving sub-problems and then putting these sub-solutions together

to produce a solution for the original problem. We'll address the following questions: What is the best size and structure of an organization? What is the best way in which to allocate sub-problems, organize information flows, and more broadly connect agents so as to lead to fast and accurate solutions? Should the organization be "flat" so that many agents are handling data? Should it be decentralized like a team or centralized like a hierarchy? How many levels should the hierarchy have and should communication channels cut across levels so high-level personnel connect with many levels? In addressing these questions, research has considered two sets of factors: first, the characteristics of agents with respect to their cognitive skills and accumulated knowledge; and second, the characteristics of the environment in terms of its complexity, stability, and decomposability.

We begin in Section 4.1 with the canonical model of an information processing organization and an exploration of its generic properties - properties that hold for most networks, not just optimal ones. The impact of organizational structure on performance when agents have the capacity to learn is investigated in Section 4.2, while organizational design endogenously evolves in the models reviewed in Section 4.3. We conclude with a critical discussion in Section 4.4.

4.1 Generic Properties of Information Processing Networks

Radner (1993) describes the canonical information processing problem faced by an organization.²² The organization is a network of agents (or information processors or nodes) which are endowed with a fixed ability to process incoming data and a limited capacity for doing so. For example, data might be a series of integers, the processor has the ability to multiply them together, and its capacity limits it to handling seven numbers. The architecture defines how information is distributed and tasks are assigned. In this canonical model, information enters the lowest level where it is processed and sent through the network for further processing. Once processing is completed, an output (that is, an organizational decision) emerges. The basic line of inquiry investigates the relationship between size and structure of the network and performance, which is measured by the speed with which a problem is solved. More nodes in the network (or more agents in the organization) provide more processing power - which may be particularly critical when agents have limited capacity - but at the potential cost of more delay as information has to traverse a longer path. Under certain conditions, it is shown that the most efficient network is a particular type of hierarchy.

Miller (2001) This canonical problem is explored in Miller (2001) with an eye to learning generic properties of networks. He considers randomly generated networks with the hope of identifying "order for free" without the expense of optimality. The organization faces a series of associative (and thereby decomposable) problems. The organization receives data in the form of a series of integers and the task is to generate

²²Also see Van Zandt (1999) and, for early work on modelling an organization as a network, Dow (1990).

their sum. Each agent has the ability to sum two numbers. With this class of problems, and given the assumptions placed on agents, accuracy is assured and the performance of an organization is measured by the delay in generating a solution. As the associative nature of the problem means that the sequence with which it is solved is irrelevant, such problems are ripe for distributed problem solving.

An organization is a network of nodes with each node being a processor and representing an agent. Figure 4 shows all of the possible (non-redundant) networks associated with five bits of information, (a, b, c, d, e) , where each bit is handled by exactly one agent. For example, a single-agent organization has all five bits coming into that agent who must progressively sum them by adding a to b , then adding the solution to c , and so forth, until the solution is derived after four operations and four periods. In comparison, there is a three-agent network (denoted #12) in which one agent sums three bits, another sums two bits, and a third sums the sub-solutions. The first two agents are referred to as child agents to the last one, who is the parent agent. Note that this network takes fewer periods to derive a solution but at the cost of more agents.

Faced with a sequence of problems, an agent is not allowed to work on the next problem until its output is retrieved by the next agent in the network. An agent can be in one of three states: i) inactive; ii) active and unfinished; and iii) active and finished, in which case it can, if called upon, convey its solution to its parent agent. An agent must decide on what problem to work, whether any sub-solutions from child agents can be incorporated, and whether more processing is required on the current problem. When an inactive agent is activated, it either tries to draw a child agent's solution or data from the queue. An agent remains active until processing is completed and the sub-solution is taken by the parent agent.

For the purpose of identifying generic properties, Miller considers random networks constructed as follows. A number of nodes is randomly chosen from between 1 and 50. The organization is iteratively constructed starting with a single node to which a child is added. One of those nodes is randomly selected and a child is added to it. This continues until the network has the specified number of nodes. Finally, all terminal nodes are connected to the data queue and an interior node is connected with probability $1/2$.

To explore the significance of synchronization of agents in distributively solving a problem, Miller (2001) compares the performance of networks where nodes are randomly activated with one in which there is "ordered firing" so that child nodes are activated before parent nodes. Some interesting properties arise when exploring how performance is related to organizational size, as measured by the number of nodes. When firing is synchronized, performance mildly increases with size while, with random firing, performance appears to be maximized at an intermediate number of nodes. This suggests that, to sustain larger organizations, synchronization among agents is critical. Also noteworthy is that the variation in performance across random networks is greater for small organizations. The possible explanation is that they are more susceptible to bad design causing bottlenecks, which creates delay as agents wait for sub-solutions from other agents. In contrast, the denser web of connections

when there are more nodes allows information to flow more freely, which serves to make the particular architecture less important.

4.2 Organizations with Adaptive/Learning Agents

Now consider an organization that faces what Carley (1992) calls a quasi-repetitive task. In each period a problem arises as an *iid* selection from a finite set, which provides two types of opportunities for the organization to learn. First, if the cardinality of the set of problems is not large relative to the number of periods, the organization is likely to face a problem repeatedly so they can learn from past mistakes. Second, the problems may be related, in which case the solution to one problem provides information pertinent to solving other problems. The challenge is to learn the latent function generating the problems. For an organization to take advantage of these opportunities, agents must be endowed with a capacity to learn. Exploring how the ability to learn influences the relationship between organizational structure and performance is a central issue.

Carley (1992) Suppose an organization faces a sequence of binary classification problems. For example, suppose that a new project arrives each period and the organization has to decide whether it is *profitable* or *unprofitable*. It receives information on the project that takes the form of an element of $\{0, 1\}^N$. There is a true (fixed and deterministic) latent mapping from $\{0, 1\}^N$ into $\{\textit{profitable}, \textit{unprofitable}\}$ which assigns the status of *profitable* when a majority of the bits take the value 1. Each drawn problem assigns equal probability to a bit being a 0 or 1. Based on the information, the organization must decide whether or not to conclude it is *profitable*.

In contrast to the rich set of organizational structures allowed by Miller (2001), here just two organizational forms are considered, hierarchy and team. A hierarchy comprises three levels where the lowest level has nine agents (referred to as analysts) who receive the data. The data consisting of N bits are partitioned into nine sub-vectors with each analyst receiving one of them. In response to observing an element from $\{0, 1\}^{N/9}$, an analyst puts forth a recommendation, either *profitable* or *unprofitable*, to an agent (manager) at the next level. There are three managers and each receives recommendations from three analysts. At the top is a single agent (CEO) who receives recommendations from the three managers and makes a final evaluation regarding profitability of the project. A team is also comprised of nine analysts but has just one level. Each analyst makes a recommendation in response to their input, and the organizational decision is based on majority rule. Though the number of decision makers varies between the two organizational forms, the number of agents receiving information about the problem is the same.

Agents engage in experiential learning about the latent mapping between $\{0, 1\}^N$ and $\{\textit{profitable}, \textit{unprofitable}\}$. After the organization makes its decision, all agents observe the true state of the project. Each agent keeps track of how information relates to the true state. For example, an analyst keeps a running tab of how many times a project was *profitable* for each observed input from $\{0, 1\}^{N/9}$. Similarly, managers and the CEO keep track of how many times a project was *profitable* for

each observed element from $\{profitable, unprofitable\}^3$. The specified behavioral rule is that an agent reports *profitable* (*unprofitable*) in response to his information when the fraction of times that the true state was *profitable* (*unprofitable*) for that given information exceeds 50%. When it is exactly 50%, the agent randomizes.

The task varies in terms of complexity and decomposability. Complexity is measured by the length of the data vector. More data means more problems, with less opportunity to see a particular problem repetitively, and also a bigger set of possible mappings to sort among. A problem is referred to as *decomposable consensual* when all analysts are given the same sub-problem.²³ For example, the task 110110110 is decomposable consensual to three analysts. Since the more frequent bit value for each analyst is also the more frequent bit value for all N bits, in principle an individual analyst can come to correctly identify a project's true state based only on his own $N/9$ bits of data. By contrast, a non-decomposable task is when the accuracy of an individual's prediction is dependent upon information possessed by others. For example, the task 111010000 is nondecomposable to three analysts as one analyst receives 111, a second receives 010, and the third receives 000. This information is insufficient to determine whether 1 is in the majority and thus that the project is *profitable*.

One of the unique and interesting features of this model is personnel turnover. According to a Poisson process, an agent may be replaced with a new agent.²⁴ Analysts can be replaced with someone who has no experience ("novice"), someone who has experience with 500 sub-problems generated by the same stochastic process ("good fit"), and someone who has experience with 500 sub-problems in an organization with a slightly different problem-generating process ("poor fit"). Managers can also be replaced, although their replacements are restricted to be novices. Given that agents are learning, replacing experienced agents with possibly less experienced ones obviously deteriorates performance. Less clear is what type of organizational structure better handles such disruptions.

In contrast to Miller (2001), the organization is not necessarily given enough data to correctly solve the problem. Thus, performance is measured by the accuracy of solutions. The average percentage of correct assessments in the final 200 of 2500 periods measures long-run performance, while the average number of periods it takes to reach 60% accuracy serves as a measure of the speed of learning. As there are only two true states and the organization is endowed with no experience, it is initially guessing and so starts with 50% accuracy.

For either organizational type, performance is greater with a less complex task and when the task is decomposable. Teams learn significantly faster than hierarchies (though an important exception is noted below). A key force at work here is information loss. Analysts convert information defined on a space with $2^{N/9}$ elements into a

²³The modifier "consensual" is added because this task is more restrictive than the standard definition of decomposability (see Section 2). A problem can be decomposable but not involve identical sub-problems.

²⁴Here, turnover is exogenous though in other models it is endogenous. An agent may decide to leave, as in Axtell (1999), and managers may decide whether to hire someone, as in Glance, Hogg, and Huberman (1997).

signal from a two-element space. In the hierarchy, managers take information defined on an eight-element space (the three possible recommendations from those at the next lower level) to a two-element space. On these grounds, one expects teams to perform better because there is less information loss; it occurs twice for a hierarchy but only once for a team. However, when turnover is sufficiently high, hierarchies perform better for both decomposable and nondecomposable tasks. It is unclear whether this is due to hierarchies being less sensitive to the recommendation of a single rogue analyst or to their managers having more experience.

Barr and Saraceno (2002) A similar exercise to that of Carley (1992) is performed in Barr and Saraceno (2002) though a distinctive feature of their approach is to model the organization as an artificial neural network (ANN). The organization’s task is to identify the latent relationship between information that lies in $\{0, 1\}^{10}$ and the true state that lies in $\{0, 1, \dots, 1023\}$ (as the latent function converts 10 binary digits to its equivalent number in base 10). The organization is an ANN with three layers (see Figure 5). The input layer is comprised of ten input nodes, each of which receives one of the ten bits of data. The next (hidden) layer is made up of n nodes - which can be interpreted as the lowest level in the organization with each node being an agent. Each of these agents takes a weighted sum of the data from the input layer and transforms it into an output. These n outputs then go to the top level where they are weighted and summed to produce the organization’s output. This output is a prediction of the true state.

On a broad level, learning is equivalent to that in Carley (1992) though the specifics differ both because of the type of function being learned and the use of an ANN. The state of the organization is represented by the weights that each node in the low level uses to produce output for the high level and the weights that the high level uses to produce organizational output. Initially, these weights are randomly selected. After receiving data, the organization produces an output, denoted \hat{y} , and then agents observe the realization of the latent function, y . Each agent calculates the gradient of the mean squared error, $(1/2)(y - \hat{y})^2$, with respect to their weights and incrementally adjusts them in the direction that reduces mean squared error, taking other agents’ weights as fixed.

While Carley (1992) fixes organizational size and varies structure, here structure is fixed at the two-level hierarchy and the role of size, as measured by the number of low-level agents, is explored. Interestingly, a bigger organization is not necessarily better. The reason lies in two types of prediction error. Approximation error is associated with the limited capacity of an ANN to represent a latent function. By expanding the space of approximating functions, more agents reduce approximation error. Of course, better fit also depends on the efficiency with which the coefficients (weights) of the ANN are estimated. The authors refer to this as estimation error and it measures how badly the ANN performs relative to maximal performance for a given size. The trade-off is that a larger organization reduces approximation error but, with more agents and thus more weights to be estimated, estimation error can rise. Clearly, with enough data, a bigger organization means better predictions; but,

as in the real world, the simulations have only a limited number of problems from which to learn. Small firms are interpreted as having a simpler class of functions - they don't need many problems to get low estimation error - while large firms have a richer class of functions - they are slow to learn but may ultimately have a more sophisticated solution.

An organization is faced with a set of feasible problems, each of which is a random draw from $\{0, 1\}^{10}$. The complexity of the environment is measured by the size of that feasible set, which numbers at most ten. Stability is measured by the probability that an element of the feasible set is replaced with a fresh draw from $\{0, 1\}^{10}$. This random event occurs each period.²⁵ Performance depends on the accuracy of an organization's solution and, more specifically, equals the inverse of the squared error less the cost of the network. Network cost is composed of a cost per agent plus the cost of delay, which is linear in the number of operations performed on data. Larger firms experience greater network costs but may have less error.

Optimal firm size is typically found to be an interior solution, reflecting the trade-off from a bigger organization: less approximation error, more estimation error, and a higher network cost. The most interesting results concern the interaction between stability and complexity. When complexity is high, the optimal number of agents is lower when the environment is less stable. With the set of problems to be learned changing at a faster rate, agents have to adapt their weights more frequently, and this is done less effectively when there are more weights to adjust. When instead complexity is low, optimal firm size is higher in unstable environments than in a near-stable environment. With low complexity, there are only two problems to be learned and this doesn't require many agents. As stability falls, the set of examples is changing at a faster rate and having more agents allows the organization to adapt faster. More broadly, these results seem to suggest a rising marginal cost to the number of agents. With only a few problems to be learned, the organization is initially small so that reduced stability is best handled by adding agents. However, if there are a lot of problems, then the organization is already large and adding agents in response to less stability means having to adjust far too many weights. It is preferable to reduce the number of agents, thereby trading off lower estimation error for higher approximation error.

Barr and Saraceno (2005) In an ensuing paper, the authors make a modelling advance that is innovative from both a computational and economic perspective. They allow two ANNs - each representing a firm - to coevolve in a competitive market situation. The situation is the classic symmetric Cournot game in which two firms make simultaneous quantity choices. The demand function is linear and its two parameters follow an *iid* stochastic process. The task before a firm is to learn its optimal quantity where the data it receives pertain to the unknown demand parameters. Learning is modelled as in Barr and Saraceno (2002). A firm chooses a quantity then learns *ex post* what would have been the profit-maximizing quantity.

²⁵Unfortunately, the model is designed so that a less complex environment implies a more stable one, which means any comparative statics with respect to complexity confounds these two effects.

Learning occurs in the face of an exogenously stochastic demand function and an endogenously stochastic quantity for the other firm.

In comparison with Barr and Saraceno (2002), the environment is stable and there is no network cost so performance equals profit. Given the absence of network costs, the only reason not to have more agents is greater estimation error. The coevolving system always converges to Nash equilibrium; that is, each firm's quantity converges to that which maximizes its profit. As this occurs for each realization of the demand parameters, firms are learning how the signals map into the true state of demand. Further analysis shows that average profit is initially increasing in a firm's own network size but, due to estimation error, is eventually decreasing. More interesting is that a firm's performance is initially increasing in the other firm's network size. We conjecture the reason is that a smaller rival learns slower, which means it takes longer for its quantity to settle down. This would translate into a more volatile environment for a firm and serve to lower its profit. Interestingly, it may be in the best interests of a firm that its competitor be sophisticated.

4.3 Adaptation and Evolution of Organizational Structure

In performing comparative statics to explore the impact of organizational size and structure on performance, a critical question is begged: To what extent can an organization find and adopt better structures? When dealing with complex entities such as an organization's architecture, it isn't sufficient to characterize optimal structure and presume an organization somehow finds it. Actual organizations are endowed with a structure and find large-scale change difficult. It is then worthwhile to know whether incremental changes can lead to superior designs. In addition, models of the previous section consider a very limited set of structures. By instead specifying a large class of organizations and a flexible dynamic for moving among them, new structures can emerge that are truly novel. To address these issues, we review Carley and Svoboda (1996), where simulated annealing searches for better organizations. We also return to discussing Miller (2001), who utilizes the forces of selection and adaptation through a genetic algorithm (GA). The driving question is, how effectively can an organization evolve to efficacious structures and what do those structures look like?

Carley and Svoboda (1996) With some minor modifications, Carley and Svoboda (1996) adapt the organizational model of adaptive agents of Carley (1992) by appending an organizational design dynamic to it. Thus, structure is adapting at the same time that agents are learning. A key feature of this type of model is the class of organizations over which search occurs. An organizational structure is defined by the number of agents, which agents receive data, and how agents are connected. The set of feasible organizations is limited to those with at most three levels (where each level can only report to the next higher level), at most fifteen agents on each level, and at most nine pieces of information on a task. In the event that the highest level has more than one agent, those agents use majority rule to determine the organization's choice with an equality of votes being broken through randomization.

Upon this space of organizations, a dynamic is applied which constructs a new feasible organization through four operations: i) firing (the elimination of agents); ii) hiring (the addition of agents); iii) re-tasking (a link to the data queue is redirected from one agent to another); and iv) reassigning (a link between two agents is changed so that an agent reports to a new agent). Faced with a new design, the process by which it is adopted is modelled using simulated annealing. First, an off-line experiment is performed whereby the organization’s performance (as measured by the accuracy of the organization’s decisions) is projected out for 100 tasks under this new design. If this performance exceeds the performance of the existing design then the new design is adopted. If performance is lower - and here lies a singular feature of simulated annealing - it is adopted with positive probability where this probability decreases with the existing design’s performance during the preceding 500 periods (where there is one task each period) and also exogenously declines every 200 tasks.²⁶ The minimum time between new design adoptions is 100 periods. The initial organizational structure is randomly selected and there is a training period of 500 periods before the design dynamic is turned on.

As a theoretical benchmark, the optimal design is to have a one level organization with nine agents, each receiving one of the nine bits of info, and making their decision by majority rule. Simulated annealing never finds it. Compared to random organizations, the organizations that emerge after 20,000 periods have noticeably more agents on average, a lower span of control (the average number of links to a higher level agent), and fewer links to the data queue though none of these differences are statistically significant.²⁷ Though the results of the analysis are ambiguous, the approach represents a pioneering step in modelling the evolution of organizational structure.

Miller (2001) Finally, let us return to Miller (2001) whose work on randomly generated organizations was reviewed earlier. Recall that the task is associative and thereby decomposable. As all solutions are accurate, the performance criterion is speed. Using a genetic algorithm (GA), a population of fifty randomly created organizations coevolve.²⁸ In each generation, there is a sequence of problems that each of the fifty organizations solves. Two organizations are then randomly selected and replaced with two copies of the one with greater speed. This operation is performed fifty times with replacement. These organizations are then randomly paired to engage in two genetic operations - crossover and mutation. For crossover, a node (other than one that is attached to the data queue) is randomly selected from each organization and the subtree beginning with each node (that is, the node and all of its children) are exchanged. Each organization also has a chance of mutating, which means a

²⁶The purpose of this feature is to try to keep the organization from getting stuck on bad local optima. By accepting performance-deteriorating designs, the organization might get kicked into the basin of attraction for a better local optimum.

²⁷They actually run two experiments and the results referred to here are for the case of “dual learning.”

²⁸Also see Bruderer and Singh (1996) for an early use of a GA in organization theory. For more detailed discussions of GA learning, see Brenner (2005) and Duffy (2005).

change in links. A single run has fifty generations and the output for analysis is the best organization after fifty generations. Results are based on an average over fifty runs. Miller (2001) considers the four possible cases associated with random versus ordered firing and single versus multiple problems.

To begin, is a GA outperforming random search? For comparison purposes, random search means starting with a set of randomly generated organizations (comparable in number to what the GA handles over its fifty generations) and choosing the best performer. GA is also identifying a best performer but uses crossover and mutation as well. For the case of ordered firing and a single problem, the GA impressively reduces speed by 25% compared to random search. For the other three cases, the reduction is considerably more modest at 2%. Still, the GA is creating better structures.

Whether the organization is trying to solve a single problem or a sequence of problems, results show that synchronizing the activation of agents sustains larger organizations with more levels. For a single problem, a GA produces, on average, an organization with 34 agents and eight levels under ordered firing while organizations are quite small under random firing with only three agents and less than two levels. Adding agents allows more processing to be done but at the cost that information has to travel through more levels. This can create delay, which makes ordered firing critical in keeping it under control. The superior performance of larger organizations is even stronger with multiple problems (and ordered firing) as the average size of 48 is pushing the upper bound of 50 agents. The range of size is 43 to 50 for the 50 runs (with a standard deviation of 1.9) which further suggests that to be a top performer requires being big. In contrast, for the case of a single problem, the range is vastly greater; it runs from 7 to 50 with a standard deviation of 14.2. When an organization has a light workload, a wide range of structures can perform well; when pushed harder, it becomes crucial to be larger so problems can be effectively handled without much delay.

In conclusion, a challenge for analysis is developing informative summary statistics for emergent structures. Miller (2001) goes to considerable lengths by also reporting mean path length, highest level attached to the queue, and maximum number of nodes at a level. Still, it's hard to see from these measures what the architecture looks like. One suspects it wouldn't "look like" a typical corporation. Having meaningful summary statistics for designs is essential for drawing insight and comparing results across studies. Indeed, two studies could produce organizations with a comparable number of nodes and levels but result in quite different structures. This is a challenge for future work.

4.4 Summary

Contrary to the models of search and learning in Section 3, the models explored in this section focus on organizational size as a critical factor in connection with an organization as an information processing network. Generally, more agents available to process information acts to improve predictions and produce better decisions. The analysis of Carley and Svoboda (1996) and Miller (2001) both find that their adaptive

design dynamics produce organizations with more agents. This advantage to size is more acute when the task is more complex, as the organization needs the additional processing power that comes from more agents. But bigger is not universally better. This is obvious when one assumes there is a cost to more nodes in a network, but as shown in Barr and Saraceno (2002, 2005), more agents to “train” may slow down an organization’s rate of learning. While the long-run efficiency of a network is increasing with the number of agents, smaller organizations can outperform in the intermediate run. This advantage from fewer agents is particularly relevant for a less stable environment where perpetual training occurs.

A second but more tentative piece of insight is that while bigger is typically better, organizational structure and coordination among agents may be more critical for bigger organizations. Miller (2001) finds that, when lower-level agents are activated prior to higher-level agents, the best performing organizations are vastly larger than when activation is random. Synchronization is then critical for taking advantage of larger size. This relationship between size and structure requires further examination.

In conclusion, research on information processing is trying to develop a “production function” for organizational decision-making, a difficult and challenging problem. Though significant progress has not yet occurred, the modelling approaches have been rich, novel, and provocative.

5 Effort, Norms, and Endogenous Hierarchies

While the vast majority of computational agent-based models of organizations focuses on search and information processing, there are many other organizational issues tackled. Here, we provide some of the best of this other work and in doing so touch on issues of effort and shirking, norms, and endogenizing organizational structure.

5.1 Effort and the *Commons* Problem in Organizations²⁹

[H]ardly a competent workman can be found in a large establishment ... who does not devote a considerable part of his time to studying just how slow he can work and still convince his employer that he is going at a good pace. [Frederick W. Taylor, *The Principles of Scientific Management* (1919), p. 21]

The models of organization considered thus far have assumed that the efforts required of agents – be it associated with production, innovation, or information processing – are achieved costlessly. Of course, effort is, in practice, costly and, more importantly, poses the organizational challenge of inducing agents to work hard. Organizations suffer from the “tragedy of commons” (Hardin, 1968) whereby agents shirk from a collective perspective. The essential problem here arises from the possibility that an agent may have to *share* the returns to his costly effort with other agents in the organization. While all agents would be better off if all were to exert effort, shirking with the intention to free-ride may turn out to be the dominant

²⁹The issues addressed in this section are closely related to the concerns of Janssen and Ostrom (2005).

strategy for each individual agent. As an individual's share of the returns to his/her effort is likely to depend on the number of other agents in the firm, the incentive to shirk tends to be affected by firm size. This intuition plays an important role in the ensuing analysis.

Axtell (1999) Consider a population of (non-competing) firms with workers being able to partially control their exposure to the intra-firm commons problem by switching firms or even starting their own firm. As the mobility of the workers implies that the size of the existing firm can change, it has implications for the extent to which workers will free-ride. A central focus of Axtell (1999) is on the dynamics of a population of firms whose number and size are endogenous.

A firm having $M \geq 2$ workers engages in production through the joint efforts of its members. Let $e_i \in [0, 1]$ denote worker i 's level of effort and $E \equiv \sum_{i=1}^M e_i$ be the total effort of the firm. The firm's value, $V(E)$, takes the following form: $V(E) = aE + bE^\beta$ with $a, b > 0$ and $\beta > 1$. Assume an egalitarian sharing rule so that each worker receives exactly $V(E)/M$. Denote by $U_i(e_i, E_{-i}; M)$ the utility of worker i in a firm of M workers, where he supplies e_i and everyone else supplies $E_{-i} (\equiv E - e_i)$. Workers are assumed to have Cobb-Douglas preferences for income and leisure such that

$$U_i(e_i, E_{-i}; M) = \left(\frac{V(e_i + E_{-i})}{M} \right)^{\theta_i} (1 - e_i)^{1 - \theta_i}, \quad (5)$$

where θ_i is worker i 's relative weight for income over leisure (which equals $1 - e_i$). Preferences are heterogeneous in the population as θ_i is an independent draw from a uniform distribution on $[0, 1]$.

To characterize the population of firms, let $J(t)$ be the number of firms operating at t and $M^j(t)$ denote the size of firm $j \in \{1, \dots, J(t)\}$. $e_i^j(t)$ and $E^j(t)$ represent, respectively, the effort exerted by worker i in firm j and the total effort level of firm j . The initial configuration for the computational experiment assumes a population of N workers and N single-worker firms.

In any given period, a fixed number of workers are randomly selected to alter their behavior. Workers are myopic optimizers in that, in period t , each chooses effort to maximize period t utility under the assumption that the period t total effort of the other members equals what it was in the previous period, which is denoted $E_{-i}^j(t-1)$ for firm j , and the number of its workers is the same as previously, which is $M^j(t-1)$. In this case, i was a member of firm j in $t-1$. If he remains at firm j , then worker i 's optimal effort level, $\hat{e}_i^j(t)$, is

$$\hat{e}_i^j(t) = \arg \max_e U_i(e, E_{-i}^j(t-1); M^j(t-1)). \quad (6)$$

This gives expected utility from remaining at firm j .

Alternatively, worker i can join another firm or start up a new firm (which will, at least initially, consist only of himself). As regards the former option, worker i is (randomly) endowed with a network of ν_i other workers and can consider joining the

firms to which they were members at $t - 1$. The baseline simulation assumes $\nu_i = 2 \forall i$. For each of these alternatives, the worker computes the maximal utility level using the procedure described in (6). Out of the (at most) $\nu_i + 2$ firm-options, a worker chooses the one yielding the highest expected utility.

Simulations show that the stochastic process by which firms are created, expand, and contract never settles down. Furthermore, there is considerable intertemporal fluctuations in the number of firms, average firm size (as measured by the number of workers), and average effort. Though average firm size is only four, firms can grow to be much larger. The basic forces are that, as firm size grows, increasing returns to total effort enhances marginal productivity - thereby making it more attractive for a worker to join the firm and thus leads to growth - but the free-riding problem is exacerbated with more employees - which serves to contract firm size. Firms expand when they offer a high value per worker as it induces workers to join. Now recall that a worker's optimal effort is based on the *previous* period's firm size and effort. Thus, a firm that currently has a high value per worker will experience a high inflow of new workers and, furthermore, this will continue to result in a high value per worker because each of those workers base their effort on a smaller sized firm so there is less free-riding than is appropriate for a firm of that size. This serves to attract yet more workers to join and, as long as the flow of workers into the firm remains high, increasing returns in total effort stays ahead of the intensifying free-riding problem. In this manner, a firm can experience sharp growth but it is also why it cannot maintain large size because once the flow of new workers subsides (which is sure to occur since there is a finite population of workers) then free-riding becomes the dominant force; value per worker declines and this leads to a rapid exodus of workers. Firms grow but then, like the bursting of a Ponzi scheme, eventually collapse. The model is parsimonious as a rich set of dynamics is generated by three factors: increasing returns, free-riding, and worker mobility.

Though focusing on a different set of issues, the work of Axtell (1999) has a predecessor in Glance, Hogg, and Huberman (1997). The latter authors model two organizational dilemmas: the lack of accountability in large organizations with the free-riding that ensues, as is in Axtell (1999), and the risk associated with training workers who are mobile. An organization realizes the benefits from training employees only if they remain with the organization but, once trained, a worker may leave to join another organization. Towards encompassing this latter issue, Glance et al (1997) enrich the flat organizations of Axtell (1999) by assuming each firm has a manager whose role is to decide whether to train workers and whether to add workers. A worker can join a firm only upon invitation by its manager. Both of these distinctions result in modelled firms closer to real firms than those in Axtell (1999).

In an early model of the commons problem in a team production setting, Alchian and Demsetz (1972) proposed a top-down organizational solution to free-riding. The firm is hierarchical with salaried workers and a capitalist who is motivated to monitor worker effort by virtue of being the residual claimant of firm profit. In contrast, Glance et al. (1997) and Axtell (1999) take a bottom-up approach to the issue by assuming that the workers themselves can independently control their exposure to the

commons problem by moving from one firm to another and that they also share in the firm's profit. Augmenting these models with the mechanism of Alchian and Demsetz (1972) would move these models in a useful and realistic direction. In particular, firm size is greatly limited in these models because of the intensity of the free-riding problem. Allowing multiple layers with each layer monitoring the one below them could allow for larger firms and perhaps even persistently large firms, which is a feature of the data (see, for example, Mueller, 1986) but not a property of the model.

5.2 Organizational Norms

At one point during his investigations, [consultant] Sym-Smith asked [Sears managers] how controversy was handled at the upper level of Sears. He was told that there was no controversy. Senior Searsmen were trained from their corporate infancy to participate in a veritable cult of contrived harmony and consensus. [Donald R. Katz, *The Big Store: Inside the Crisis and Revolution at Sears* (1987), p. 28.]

As discussed earlier in the context of organizational search and learning, the long-run performance of an organization depends crucially on the way it balances exploration with exploitation. There are two issues central to this trade-off. First, exploitation at the organizational level relies upon diversity at the agent-level; there must be someone who knows something special in order for the rest of the organization to learn something new and possibly useful. When agents engage in independent innovation, diversity is naturally generated, thereby providing the raw material for exploitation by the organization. However, the very process of global exploitation reduces the degree of diversity - replacing ideas with what are considered to be superior ones - so that eventually improvements in organizational performance disappear.

The second issue is how the global exploitation of local knowledge gets carried out in the organization. We've considered exploitation being done under centralization (for example, the top-down mandate of a superior practice) as well as decentralization (for example, agents share information and individually decide on whether to adopt an idea). March (1991) considers a particular form of decentralized learning in which the agents learn from organizational norms - "accepted wisdom" as to the proper way in which to do things - but where the norms themselves evolve as they are shaped by the behavior of the more successful agents within the organization. The coevolutionary dynamics between organizational norms and agent behavior drive performance by influencing the extent of diversity in the population and, therefore, the delicate balance between exploration and exploitation.

March (1991) Consider an organization facing an external reality that takes values from $\{-1, 1\}$ on m dimensions. The external reality is known only to the modeler and is assumed to be fixed for the initial set of analyses. The organization has n agents who in each period hold beliefs about the external reality. Agents' beliefs on each dimension lie in $\{-1, 0, 1\}$ as does the organizational norm (or code). These beliefs coevolve and only indirectly connect to external reality. In any period, if the code

is 0 for a particular dimension, then agents do not modify their beliefs about that dimension. It is as if the code has nothing to prescribe for that dimension. If instead the code is -1 or 1 and differs from an agent's belief, then the belief of that agent switches with probability p_1 to what the code dictates. It is natural to interpret p_1 as a measure of socialization since it controls the degree to which an agent is influenced by organizational norms. As agents learn from the code, the code itself evolves to conform to the beliefs of those agents whose beliefs are closer to external reality than that of the code. To be specific, if the code differs from the majority view of those agents whose beliefs (over all dimensions) are closer to reality, then the code remains unchanged with probability $(1 - p_2)^k$ where k is the difference between the number of agents whose beliefs differ from the code and the number with the same belief. p_2 then controls how effectively the code responds to the beliefs of the "best" agents.

The performance of the organization is measured by two levels of knowledge: the accuracy of the organizational code (which is the proportion of the organizational code that matches reality) and the average accuracy of the organization's members (which is the average proportion of individual beliefs that match reality). As agents and the code influence each other, they converge over time. An equilibrium is reached when the organizational code and the individuals share common beliefs over all m dimensions. At that point, no further learning is possible though these beliefs need not match up with external reality.

Given the mutual learning dynamics between an organization's members and its norms, slower socialization (that is, a lower value for p_1) enhances the equilibrium level of knowledge. Furthermore, there is an interesting interaction between socialization and the adaptivity of the organizational code (as measured by p_2). When socialization is slow, an increase in code adaptivity raises the average level of knowledge; when socialization is fast, a more adaptive code reduces knowledge. The equilibrium knowledge level is maximized when norms respond quickly and the population is comprised of slow-adjusting agents. The key to understanding these results is to recognize from where the raw material for learning is coming. In that agents and the organizational code learn from each other to the extent that their beliefs differ, what drives mutual learning is *sustained* diversity in beliefs. Rapid socialization causes agents' beliefs to converge to the organizational code before the code has been able to match the beliefs of the agents whose beliefs are most accurate. In contrast, slow socialization coupled with a rapidly learning code maintains a sufficient amount of diversity in the population during the code's adaptation. This augments the spreading of correct beliefs throughout the organization, with these correct beliefs ultimately becoming embedded in the code.

Just as diversity of beliefs is conducive to knowledge accumulation, so is heterogeneity in learning rates among agents. For the same average rate of learning, a mix of fast and slow learners leads to more aggregate knowledge than a homogeneous group. The slow learners provide the raw material that the organization needs to adapt in the long run, while the fast learners take advantage of the code capitalizing on this diversity; they perform the exploitation function. Providing a dilemma for organizations, the individual performance of slow learners is worse than that of fast

learners, as reflected in the inaccuracy of their beliefs, which means that fast learning can be good for the agent but bad for the organization.

A similar set of forces comes into play when personnel turnover and environmental turbulence are introduced. Suppose that, in each period, an agent is replaced with probability p_3 by a new agent with a fresh set of beliefs. When socialization is slow, an increase in p_3 decreases the average level of knowledge as these new agents replace accurate beliefs with randomly selected ones. However, when socialization is rapid, long-run knowledge is maximized with a moderate rate of personnel turnover as it serves to introduce diverse beliefs and thus to prevent premature fixation on homogeneous (but incorrect) beliefs. The impact of environmental turbulence is examined by stochastically shifting external reality. If the rate of environmental change is such that the population reaches an equilibrium before effectively responding to the turbulence, organizational performance tends to degrade as the homogeneous population lacks the raw material to respond to a changing reality. Once again, personnel turnover can enhance knowledge by injecting new beliefs into the organization.

In evaluating this model, it clearly lacks the richness of structure of the previous models reviewed. Learning is occurring in an unstructured space, thus the model does not deliver the type of insight obtained when there is the additional structure of a landscape. Also, the focus on beliefs without an explicit specification of how they map into performance omits an essential step in the norm-formation process. All these weaknesses aside, the paper makes a singular contribution in providing a plausible and parsimonious feedback mechanism for the determination of organizational norms.

5.3 Growing an Organization

If you don't zero in on bureaucracy every so often, you will naturally build in layers. You never set out to add bureaucracy. You just get it. [David Glass, CEO of Wal-Mart, quoted in Sam Walton with John Huey, *Sam Walton: Made in America* (1992), p. 232.]

Thus far, the primary approach to studying organizational structure has been to exogenously specify various structures – in terms of the communication network and the allocation of information and decision-making – and to compare their performance. While these models are bottom-up to the extent that organizational behavior is the product of the interactions of individual agents' acting according to their decision rules, they are top-down in terms of organizational structure, as it is pre-specified by the modeler. Though organizational structure is endogenized in such work as Carley and Svoboda (1996), DeCanio, Dibble, and Amir-Atefi (2000), Miller (2001), and Ethiraj and Levinthal (2002), this is done by specifying a super-agent process as reflected in, for example, applying a genetic algorithm on a population of organizations. It fails to produce organizational structure from the bottom-up by having it be the product of the decisions of individual agents within the organization. This all-important task - using the bottom-up approach of agent-based models to generate the structure of an organization - is initially attempted in Epstein (2003). Though, as we'll later argue, the model has features running counter to real organizations, it is

a novel and thought-provoking initial salvo on this challenging fundamental problem.

Epstein (2003) In this model, individual agents in the organization endogenously generate internal hierarchy in response to their environment. The environment for the organization is represented by a flow of “opportunities” that are met by the available pool of labor (agents). The central organizational problem is how to allocate the fixed pool of labor within the organization so as to most effectively respond to these opportunities.

The type of task faced by the organization is visually summarized in Figure 6. There is a fixed number N of sites (where $N = 8$ in Figure 6), each of which may receive a profit opportunity. One might imagine a site corresponds to a geographic or product market and an opportunity is demand to be met. The baseline organization consists of a fixed number of workers and level-1 managers. Each worker is assigned to a market site and the organization earns profit when a worker is at a site when it receives an opportunity. Using Epstein’s colorful terminology, a worker “intercepts” an opportunity if present when one arrives and a “penetration” occurs when an opportunity arrives without a worker there to intercept it. In Figure 6, the firm has five workers who are positioned at sites 1, 3, 4, 7, and 8 and there are four level-1 managers, each being in charge of two adjacent sites. Opportunities are coming into sites 1, 2, 4, 5, 6, and 7 with the workers at sites 1, 4, and 7 positioned to intercept. The opportunities coming into sites 2, 5, and 6, on the other hand, are wasted and represent penetrations. Finally, the workers at sites 3 and 8 are idle for lack of opportunities. Penetrations and idle workers are monitored by level-1 managers. For instance, the level-1 manager in charge of sites 1 and 2 recognizes the need for a worker to meet the opportunity at site 2. Concurrently, the level-1 manager in charge of sites 3 and 4 recognizes that the worker at site 3 is underutilized. Clearly, an appropriate move for the organization is to shift the worker from site 3 to site 2.

The organizational problem in this model is to efficiently allocate its workforce. However, workers cannot, by themselves, move among sites but may be reallocated by upper management. Epstein considers two approaches to solve the allocation problem, though only the hierarchy approach will concern us here.³⁰ This approach has managers creating higher level managers to solve the allocation problem. In the example above, a level-1 manager would “activate” a level-2 manager who would be in control of the four sites (two sites each from the two subordinate managers) and thus have the capacity to move workers among those sites. For instance, the level-1 manager controlling sites 1 and 2 can activate a level-2 manager who has control over sites 1 through 4 and who can thus observe and respond to the excess demand at site 2 and the excess supply at site 3. Being in charge of sites 1 through 4, he has the authority to shift the worker in site 3 to site 2 and balance the demand and supply of the workers for the sites that are under his control.

A manager’s decision rule for activating an upper level manager is defined by three

³⁰Indeed, the primary objective of the paper is to characterize the optimal solution - hierarchy or a trade mechanism - and how it depends on the organization’s objective. Our interest is more in terms of it as a modelling approach to endogenizing hierarchies.

parameters: two penetration threshold values, denoted T_{\min} and T_{\max} , and a finite memory of length m . Given the number of penetrations recorded in their memory, a level- k manager computes the average number of penetrations per period, P , over the 2^k market sites he controls. If $P \geq T_{\max}$ then, with some upward inertia, a manager of level $k + 1$ is created. If $P < T_{\min}$ then, with some downward inertia, the manager cedes authority to managers at level $k - 1$. This inertia captures the reluctance of a manager to relinquish control. There is no change in the current hierarchical structure when $T_{\min} \leq P < T_{\max}$. The threshold values can vary across levels, though they are specified to be the same within a managerial level. Given a pattern of opportunities arriving at these sites over time, the baseline organization can endogenously grow its hierarchy to as many as $\log_2 N$ levels.

Suppose the flow of opportunities is continual and concentrated on a set of sites for which there are, initially, no workers. The hierarchy mechanism creates additional managerial layers to handle this misallocation as long as the penetration thresholds and the upward inertia parameter are set sufficiently low. The emergent hierarchy, even after the workers have been properly allocated to effectively intercept all incoming opportunities, tends to remain in place when the downward inertia of the top managerial level is sufficiently high. When downward inertia is instead low at all management layers, then the generated hierarchy quickly dissolves after successfully reallocating labor. The flexibility with which the organization restructures itself internally - both to effectively reallocate labor and then to dismantle itself when no longer needed - depends on the inertia embedded in agents' decision rules as well as the thresholds for inducing a change in structure.

As a theory of organizational structure, Epstein (2003) offers a rich and novel approach to organizations but it has a critical feature which runs counter to our understanding of real organizations.³¹ In this model, managerial layers emerge from below, as managers create levels above them to coordinate the behavior of what were originally independent divisions. To begin, in most organizations (such as corporations and governments), there is always a manager at the top who, at least in principle, can reallocate resources as desired. More importantly, managers only have authority over reorganizing what lies beneath them in the hierarchy so that, as a result, new managerial layers are created from above. A commonly purported motivation for adding middle level managers is that upper managers perceive themselves as overburdened and thus distribute tasks and authority to newly created managerial levels. In Epstein's model, organizational structure is created in a direction running counter to reality.

In spite of this weakness, Epstein (2003) is a provocative study. Epstein lays down an important issue for future research - to model the internal organizational pressures that create a need for a new organizational structure and the process by which change is realized. This would represent the acme of agent-based models of organizations; it closes the circle in that an organization can re-invent itself through the decisions of the organization's members.

³¹In fairness, Epstein (2003) states that the model is not intended to represent any existing organization.

6 Critique of the Past and Directions for the Future

You don't want to learn a science in its early stages. ... You have to think about ... your mind as a resource to conserve, and if you fill it up with infantile garbage it might cost you something later. There might be right theories that you will be unable to understand five years later because you have so many misconceptions. You have to form the habit of not wanting to have been right for very long. If I still believe something after five years, I doubt it. [Marvin Minsky quoted in Stewart Brand, *The Media Lab: Inventing the Future at M.I.T.* (1987), pp. 103-104]

Recent research in the computational agent-based literature has provided a new and fresh perspective to exploring organizations. There is real promise that theory can produce precise results while encompassing the rich institutional features of corporations, governments, political parties, and other organizations. But if we are to effectively traverse the learning curve associated with this new modelling approach, we must maintain a healthy level of skepticism. Research builds its own momentum as assumptions initially considered problematic are routinized, arbitrary modelling conveniences become entrenched and leave unexplored the sensitivity of results to them, and standards for acceptable work form when methods are rudimentary. As March (1991) discovered, stability during an intense learning phase can be quite deleterious. We are at such a point and it is wise that we be on guard against acquiring bad habits. Towards this end, we'll make three methodological points in this section and conclude with a few suggested directions for research.

The first point is that, while there is always a disconnect between our models and what they are intended to represent, this can be a more serious issue with computational agent-based models. This concern does not come from modelling simplicity - indeed, the models are quite rich compared to their predecessors - but rather that insufficient attention may be given to relating a model to reality. Many of the modelling components - artificial neural networks, simulated annealing, genetic algorithms, and the like - were originally developed for very different purposes and some work has used them without adequately explaining how these theoretical constructs map into real-world entities and processes. For example, what is the correspondence between the components of an artificial neural network and the components of a firm? Is it appropriate to interpret a node as a person? If not, what additional structure would make a node a reasonable representation of a person? What is the correspondence between a genetic algorithm and the process of imitation and innovation conducted within and between organizations? Is crossover a descriptively accurate model of some organizational process? Before "off-the-shelf" modules are deployed in modelling organizations, the researcher should map it to what is being modelled. Doing so will not only lead to more confidence in the model but is likely to suggest useful modifications.

One of the reasons that neoclassical economists resist bounded rationality is that there are so many ways to model it, and often which is selected is arbitrary. This is a legitimate concern, although it should not deter one from engaging in such work.

Indeed, the equilibrium assumption - an agent “understands” the world around him in the sense of, for example, knowing the behavioral rules used by other agents - is as *ad hoc* as most assumptions of boundedly rational agents in that, in most instances, it is not based on empirical evidence and often no credible story can be told to make the assumption convincing. (The appeal of the equilibrium assumption is not its empirical validity but rather its power in generating precise results and its accordance with the faith of many neoclassical economists in equilibrating processes.) This leads to our second point. Given there are many ways in which to model bounded rationality, a feature to the broad research program should be assessing the robustness of insight to the particular way bounded rationality is instantiated in agents and how the tasks facing agents are represented. In finding a solution, does it make a difference whether an organization is modelled as an artificial neural network or as a collection of myopic hill-climbing agents? Do results depend on the organization solving a decomposable problem or a binary classification problem or minimizing a distance function? Does it make a difference whether organizational structure evolves as represented by simulated annealing or a genetic algorithm? Rather than consider one particular task, it may be more useful to allow for a variety of tasks, exploring how the optimal organizational structure depends on the task and identifying those structures that perform reasonably well for an assortment of tasks.

The third point to make about this literature is that results can be inadequately explained. This is partly due to models being too complex and researchers forgetting that parsimony is a virtue, not a weakness emanating from a lack of computing power. Indeed, the poignancy of Einstein’s well-known apothegm - “Everything should be as simple as possible, but not simpler.” - is nowhere greater than with agent-based computational modelling.³² As the power of computing allows us to solve models of increasing complexity, there is a natural tendency to complicate. This is a mistake. Even with Moore’s Law sailing at full mast, computing constraints continue to make our models gross simplifications of what we are trying to understand. The deliverable of formal models remains what it has always been - insight. A model that is so complex that its implications elude explanation is a model that has not altered our understanding.³³

Complexity aside, a more disturbing feature of this work is the sometimes perceived lack of necessity to carefully explain results. An attractive feature of a mathematical proof is that it provides a paper trail that can be used to explain results. Though computational results are also the product of logical operations, there is a tendency to think that if the model cannot be solved analytically then there is little point to trying to carefully sort out how output is produced. Anyone who has worked with computational models knows that results can be the product of arbitrary assumptions of convenience or coding errors, which makes it all the more critical

³²Indeed, some work in the computational agent-based literature seems to be guided by the axiom: “Make it simple enough to be computable and complex enough to be incomprehensible.”

³³It is in this light that we decry “emergent phenomena” when it is meant to refer to results unanticipated by virtue of a model’s complexity. If one could not, upon proper reflection, anticipate the possibility of some results then it is hard to see how one can *ex post* explain them and, if one cannot explain them, in what sense do we understand more.

that they be adequately explained. Though computational work may not leave an analogous paper trail, it can offer a way in which to “test” an explanation. If one conjectures that a result is due to a particular force, then it may be possible to “turn off” that force. If the result persists then one’s conjecture is wrong; if the result goes away then the “evidence” is consistent with that explanation. Furthermore, explanation is not only essential to gaining insight but also to assessing robustness. Convincingly arguing that the forces driving the results are not peculiar to those examples is the way in which to develop confidence that the insight uncovered is broadly applicable. The bottom line is that researchers must apply the same standards for explaining results that are used in the assumption-proof-theorem tradition.³⁴

Given research on organizations using a computational agent-based approach is in its incipency, there are multitudes of research directions. Rather than propose specific lines, which would only serve to scratch the surface and deplete what minuscule scholarly wealth the authors possess, let us instead provide three general directions for research.

One direction is to take bounded rationality another step further. While agents are modelled as being limited in their decision-making capacity, they are often assumed to have an unrealistic amount of information, either before or after acting. A common assumption in rugged landscape models is that an agent observes *ex ante* the true performance associated with an idea and, based on that information, decides whether or not to adopt it. In some cases, this can be plausibly motivated by imagining that the idea is temporarily implemented with (noiseless) performance being observed. Learning is then occurring offline. If, however, there is noise, then learning will have to occur in real time - an organization may need to continue the experiment for a non-trivial length of time in order to get a sufficiently informative signal. Before even experimenting with an idea, it will want to make an assessment of its potential but then the agent must have a “model” so as to make such a judgment. That is a feature lacking in most agent-based models (Gavetti and Levinthal, 2000, being an exception). Depending on how one models the evolution of an agent’s model of the landscape, biased and not just noisy evaluations could emerge.

An analogously strong assumption is made in many information-processing models, which is that agents learn *ex post* what was the true state; an organization receives data, makes a choice, observes the outcome, and is able to infer from the outcome what would have been the right choice. In practice, the true state is rarely observed and, while performance may be observed, it provides noisy information regarding what would have been the best decision. In addition, when there are more than a few members, organizational performance is a highly uninformative signal of what an agent outside of the upper-most levels should have done. Agents need to know about their “local” performance rather than the global performance of the organization. Models have to come to grips with how an organization measures an individual agent’s contribution to total performance.

A second direction is to bring in more structure. Thus far, models have been

³⁴This comment was distinctly improved by a stimulating dinner conversation between one of the authors and Patrick Rey in Siena, Italy.

too generic. The results generated by models of search and learning are extensions or applications of insight regarding search on rugged landscapes. If we're to move beyond that, we need to impose more structure so that a variable is not some faceless dimension but concretely corresponds to an actual practice. This would allow one to explore not only how many dimensions should be centralized but also *which* dimensions should be centralized. What determines whether, say, marketing is controlled by the corporate office or a product manager? What determines which dimensions a store manager controls rather than assistant managers? An important step is to further pursue the approach of building a landscape from economic primitives by modelling specific functions - pricing, product selection, training practices, marketing, inventory policy, etc. Such an approach could open up an entirely new set of questions and make these models more powerful both in explaining organizational behavior and serving a normative role for organizations.

More structure is also needed in information processing models where, thus far, agents are excessively simple-minded and too heavily "programmed," even by the standards of the computational agent-based literature. Endowing them with preferences and giving them choices - such as how much effort to exert and what information to pass onto the next node in the network - is vital for the distance between models and reality to lessen.

At present, organization theory is partitioned into the neoclassical economics approach and the agent-based computational approach and "ne'er the twain shall meet." It is obvious that these two research lines should not be moving independently. Each has its virtues - computational work provides a rich modelling of organizational structure and how agents interact while neoclassical work is sophisticated in its modelling of incentives - and a superior theory of organizations is to be had if the two can be integrated. This challenge is the third direction.

References

- [1] Aghion, P. and J. Tirole (1997), "Formal and real authority in organizations", *Journal of Political Economy* 105: 1–29.
- [2] Alchian, A., and H. Demsetz (1972), "Production, information costs, and economic organization", *American Economic Review* 62: 777–795.
- [3] Aldrich, H. (1999), *Organizations Evolving* (Sage Publications, London).
- [4] Axtell, R. (1999), "The emergence of firms in a population of agents: local increasing returns, unstable Nash equilibria, and power law size distributions", CSED Working Paper No. 3, Brookings Institution.
- [5] Barr, J., and F. Saraceno (2002), "A computational theory of the firm", *Journal of Economic Behavior and Organization* 49: 345–361.
- [6] Barr, J., and F. Saraceno (2005), "Cournot competition, organization and learning", *Journal of Economic Dynamics and Control* 29: 277–295.

- [7] Baum, J. A. C., and J. V. Singh, eds. (1994), *Evolutionary Dynamics of Organizations* (Oxford University Press, New York).
- [8] Baum, J. A. C., ed. (2002), *The Blackwell Companion to Organizations* (Blackwell Publishers, Oxford).
- [9] Brand, S. (1987), *The Media Lab: Inventing the Future at M.I.T.* (Viking, New York).
- [10] Brenner, T. (2005), “Agent learning representation: advice on modelling economic learning”, this handbook.
- [11] Bruderer, E., and J. V. Singh (1996), “Organizational evolution, learning, and selection: a genetic-algorithm-based model”, *Academy Management Journal* 39: 1322–1349.
- [12] Burton, R. M., and B. Obel (1980), “A computer simulation test of the M-form hypothesis”, *Administrative Science Quarterly* 25: 457–466.
- [13] Carley, K. M. (1992), “Organizational learning and personnel turnover”, *Organization Science* 3: 20–46.
- [14] Carley, K. M., and L. Gasser (1999), “Computational organization theory” in: G. Weiss, ed., *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence* (The MIT Press, Cambridge, MA).
- [15] Carley, K. M., and M. J. Prietula, eds. (1994), *Computational Organization Theory* (Lawrence Erlbaum Associates, Hillsdale, NJ).
- [16] Carley, K. M., and D. M. Svoboda (1996), “Modeling organizational adaptation as a simulated annealing process”, *Sociological Methods & Research* 25: 138–168.
- [17] Carroll, G. R., and J. R. Harrison (1998), “Organizational demography and culture: insights from a formal model and simulation”, *Administrative Science Quarterly* 43: 637–667.
- [18] Carroll, T., and R. M. Burton (2000), “Organizations and complexity: searching for the edge of chaos”, *Computational and Mathematical Organization Theory* 6: 319–337.
- [19] Chandler, A. D. Jr. (1977), *The Visible Hand: The Managerial Revolution in American Business* (Belknap Press of Harvard University Press, Cambridge, MA).
- [20] Chang, M.-H., and J. E. Harrington, Jr. (1997), “Organizational structure and firm innovation in a retail chain”, *Computational and Mathematical Organization Theory* 3: 267–288.

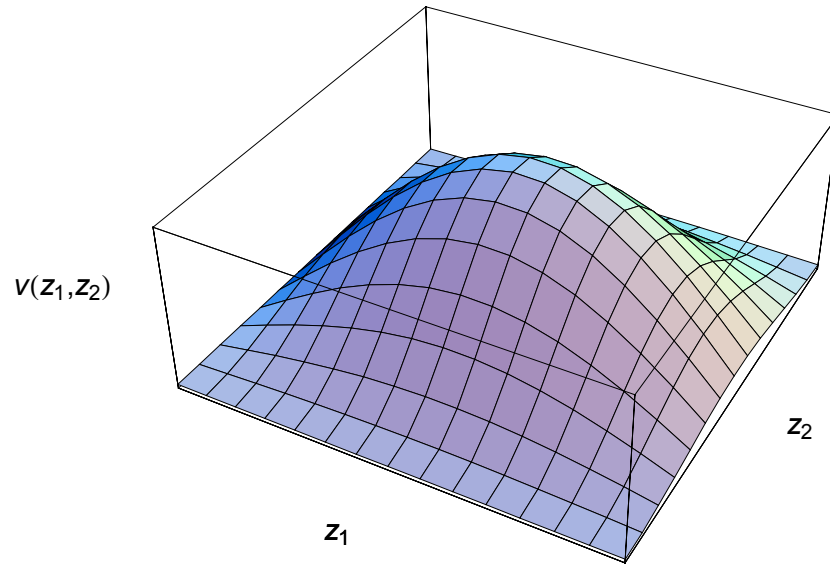
- [21] Chang, M.-H., and J. E. Harrington, Jr. (2000), “Centralization vs. decentralization in a multi-unit organization: a computational model of a retail chain as a multi-agent adaptive system”, *Management Science* 46: 1427–1440.
- [22] Chang, M.-H., and J. E. Harrington, Jr. (2003), “Multi-market competition, consumer search, and the organizational structure of multi-unit firms”, *Management Science* 49: 541–552.
- [23] Chang, M.-H., and J. E. Harrington, Jr. (2004), “Organization of innovation in a multi-unit firm: coordinating adaptive search on multiple rugged landscapes” in: W. Barnett, C. Deissenberg, and G. Feichtinger, eds., *Economic Complexity: Non-linear Dynamics, Multi-agents Economies, and Learning*, ISETE, Volume 14 (Elsevier, Amsterdam), pp. 189–214.
- [24] Cohen, M. D., and L. S. Sproull, eds. (1996), *Organizational Learning* (Sage Publications, Thousand Oaks, CA).
- [25] DeCanio, S. J., and W. E. Watkins (1998), “Information processing and organizational structure”, *Journal of Economic Behavior and Organization* 36: 275–294.
- [26] DeCanio, S. J., C. Dibble, and K. Amir-Atefi (2000), “The importance of organizational structure for the adoption of innovations”, *Management Science* 46: 1285–1299.
- [27] Dessen, W. (2002), “Authority and communication in organizations”, *Review of Economic Studies* 69: 811–838.
- [28] Dow, G. K. (1990), “The organization as an adaptive network”, *Journal of Economic Behavior and Organization* 14: 159–185.
- [29] Duffy, J. (2005), “Agent-based models and human subject experiments”, this handbook.
- [30] Durfee, E. H. (1999), “Distributed problem solving and planning” in: G. Weiss, ed., *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence* (The MIT Press, Cambridge, MA).
- [31] Epstein, J. M. (2003), “Growing adaptive organizations: an agent-based computational approach”, The Santa Fe Institute, Working Paper 03-05-029 [forthcoming in: Epstein, J. M. (2005), *Generative Social Science: Studies in Agent-Based Computational Modeling* (Princeton University Press, Princeton, NJ)].
- [32] Ethiraj, S. K., and D. Levinthal (2002), “Search for architecture in complex worlds: an evolutionary perspective on modularity and the emergence of dominant designs”, Wharton School, University of Pennsylvania, pdf copy.
- [33] Gavetti, G., and D. Levinthal (2000), “Looking forward and looking backward: cognitive and experiential search”, *Administrative Science Quarterly* 45: 113–137.

- [34] Ghemawat, P., and D. Levinthal (2000), “Choice structures and business strategy”, The Wharton School, University of Pennsylvania, pdf copy.
- [35] Glance, N. S., T. Hogg, and B. A. Huberman (1997), “Training and turnover in the evolution of organizations”, *Organization Science* 8: 84–96.
- [36] Hardin, G. (1968), “The tragedy of the commons”, *Science* 162: 1243–1248.
- [37] Harrington, J. E. Jr. (1998), “The social selection of flexible and rigid agents”, *American Economic Review* 88: 63–82.
- [38] Harrington, J. E. Jr. (1999), “Rigidity of social systems”, *Journal of Political Economy* 107: 40–64.
- [39] Harrington, J. E., Jr., and M.-H. Chang (2005), “Coevolution of firms and consumers and the implications for market dominance”, *Journal of Economic Dynamics and Control* 29: 245–276.
- [40] Hayek, F. A. (1945), “The use of knowledge in society”, *American Economic Review* 35: 1–18.
- [41] Holmstrom, B. R., and J. Tirole (1989), “The theory of the firm” in: R. Schmalensee and R. D. Willig, eds., *Handbook of Industrial Organization*, Volume I (Elsevier, Amsterdam).
- [42] Janssen, M. A., and E. Ostrom, “Governing social-ecological systems”, this handbook.
- [43] Jensen, M. C., and W. H. Meckling (1995), “Specific and general knowledge, and organizational structure”, *Journal of Applied Corporate Finance* 8: 4–18.
- [44] Katz, D. R. (1987), *The Big Store: Inside the Crisis and Revolution at Sears* (Penguin Books, New York).
- [45] Kauffman, S. (1993), *The Origins of Order* (Oxford University Press, New York).
- [46] Kollman, K., J. H. Miller, and S. E. Page (2000), “Decentralization and the search for policy solutions”, *Journal of Law, Economics, and Organization* 16: 102–128.
- [47] Levinthal, D. A. (1997), “Adaptation on rugged landscapes”, *Management Science* 43: 934–950.
- [48] Levinthal, D., and J. G. March (1981), “A model of adaptive organizational search”, *Journal of Economic Behavior and Organization* 2: 307–333.
- [49] Lomi, A., and E. R. Larsen, eds. (2001), *Dynamics of Organizations: Computational Modeling and Organization Theories* (AAAI Press/The MIT Press, Menlo Park, CA).

- [50] Mackie-Mason, J., and M. Wellman (2005), “Automated markets and trading agents”, this handbook.
- [51] March, J. G. (1991), “Exploration and exploitation in organizational learning”, *Organization Science* 2: 71–87.
- [52] March, J. G., and H. A. Simon (1958), *Organizations* (Wiley, New York).
- [53] Maskin, E., Y. Qian, and C. Xu (2000), “Incentives, information, and organizational form”, *Review of Economic Studies* 67: 359–378.
- [54] Milgrom, P., and J. Roberts (1992), *Economics, Organization and Management* (Prentice Hall, Englewood Cliffs, NJ).
- [55] Miller, J. H. (2001), “Evolving information processing organizations” in: A. Lomi and E. R. Larsen, eds., *Dynamics of Organizations: Computational Modeling and Organization Theories* (AAAI Press/The MIT Press, Menlo Park, CA).
- [56] Minsky, M. (1986), *The Society of Mind* (Simon and Schuster, New York).
- [57] Mueller, D. (1986), *Profits in the Long Run* (Cambridge University Press, Cambridge).
- [58] Nelson, R. R., and S. G. Winter (1982), *An Evolutionary Theory of Economic Change* (The Belknap Press of Harvard University Press, Cambridge, MA).
- [59] Page, S. E. (1996), “Two measures of difficulty”, *Economic Theory* 8: 321–346.
- [60] Prendergast, C. (1999), “The provision of incentives in firms”, *Journal of Economic Literature* 37: 7–63.
- [61] Prietula, M. J., K. M. Carley, and L. Gasser, eds. (1998), *Simulating Organizations: Computational Models of Institutions and Groups* (The MIT Press, Cambridge, MA).
- [62] Qian, Y. (1994), “Incentives and loss of control in an optimal hierarchy”, *Review of Economic Studies* 61: 527–544.
- [63] Radner, R. (1993), “The organization of decentralized information processing”, *Econometrica* 61: 1109–1146.
- [64] Rivkin, J. (2000), “Imitation of complex strategies”, *Management Science* 46: 824–844.
- [65] Rivkin, J., and N. Siggelkow (2003), “Balancing search and stability: interdependencies among elements of organizational design”, *Management Science* 49: 290–311.
- [66] Sathe, V. (1985), *Culture and Related Corporate Realities: Text, Cases, and Readings on Organizational Entry, Establishment, and Change* (R.D. Irwin, Homewood, IL).

- [67] Schumpeter, J. A. (1939), *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process*, Volume I (McGraw-Hill, New York).
- [68] Siggelkow, N., and D. A. Levinthal (2003), “Temporarily divide to conquer: centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation”, *Organization Science* 14: 650–669.
- [69] Simon, H. A. (1996), *The Sciences of the Artificial* (MIT Press, Cambridge, MA).
- [70] Sorensen, O. (2002), “Interorganizational complexity and computation” in: J. Baum, ed., *The Blackwell Companion to Organizations* (Blackwell, Oxford).
- [71] Taylor, F. W. (1919), *The Principles of Scientific Management* (Harper and Brothers Publishers, New York) [original edition, 1911].
- [72] Van Zandt, T. (1999), “Real-time decentralized information processing as a model of organizations with boundedly rational agents”, *Review of Economic Studies* 66: 633–658.
- [73] Vriend, N. (2005), “ACE models of endogenous interactions”, this handbook.
- [74] Walton, S. with J. Huey (1992), *Sam Walton: Made in America* (Doubleday, New York).
- [75] Weiss, G., ed. (1999), *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence* (The MIT Press, Cambridge, MA).
- [76] Wilhite, A. (2005), “Economic activity on fixed networks”, this handbook.

(a) Smooth Landscape



(b) Rugged Landscape

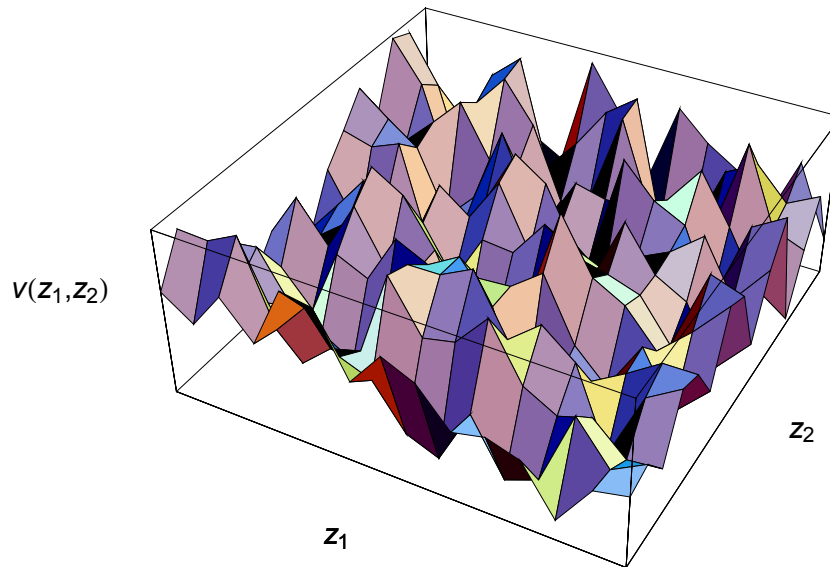


Figure 1: Search Landscapes

(a) Independence ($K = 0$)

	1	2	3	4	5	6
1	x					
2		x				
3			x			
4				x		
5					x	
6						x

(b) Full Interaction ($K = 5$)

	1	2	3	4	5	6
1	x	x	x	x	x	x
2	x	x	x	x	x	x
3	x	x	x	x	x	x
4	x	x	x	x	x	x
5	x	x	x	x	x	x
6	x	x	x	x	x	x

(c) Block-Diagonal ($K = 2$)

	1	2	3	4	5	6
1	x	x	x			
2	x	x	x			
3	x	x	x			
4				x	x	x
5				x	x	x
6				x	x	x

(d) Random Interaction ($K = 2$)

	1	2	3	4	5	6
1	x	x			x	
2	x	x		x		
3		x	x			x
4	x			x		x
5		x			x	x
6	x		x			x

Figure 2: Adjacency Matrix ($N = 6$)

(a)
Hierarchical and
Nearly Decomposable

	a1	a2	a3	b1	b2	b3	c1	c2	c3
a1	x	x	x						
a2	x	x	x						
a3	x	x	x						
b1			x	x	x	x			
b2				x	x	x			
b3				x	x	x			
c1						x	x	x	x
c2							x	x	x
c3							x	x	x

(b)
Non-Hierarchical and
Nearly Decomposable

	a1	a2	a3	b1	b2	b3	c1	c2	c3
a1	x	x	x						
a2	x	x	x						
a3	x	x	x	x					
b1			x	x	x	x			
b2				x	x	x			
b3				x	x	x	x		
c1						x	x	x	x
c2							x	x	x
c3							x	x	x

(c)
Hierarchical and
Non-Decomposable

	a1	a2	a3	b1	b2	b3	c1	c2	c3
a1	x	x	x						
a2	x	x	x						
a3	x	x	x						
b1	x	x	x	x	x	x			
b2	x	x	x	x	x	x			
b3	x	x	x	x	x	x			
c1				x	x	x	x	x	x
c2				x	x	x	x	x	x
c3				x	x	x	x	x	x

(d)
Non-Hierarchical and
Non-Decomposable

	a1	a2	a3	b1	b2	b3	c1	c2	c3
a1	x	x	x	x					
a2	x	x	x	x	x				
a3	x	x	x	x	x	x			
b1	x	x	x	x	x	x	x		
b2		x	x	x	x	x	x	x	
b3			x	x	x	x	x	x	x
c1				x	x	x	x	x	x
c2					x	x	x	x	x
c3						x	x	x	x

Figure 3: Adjacency Matrix ($N = 9$)

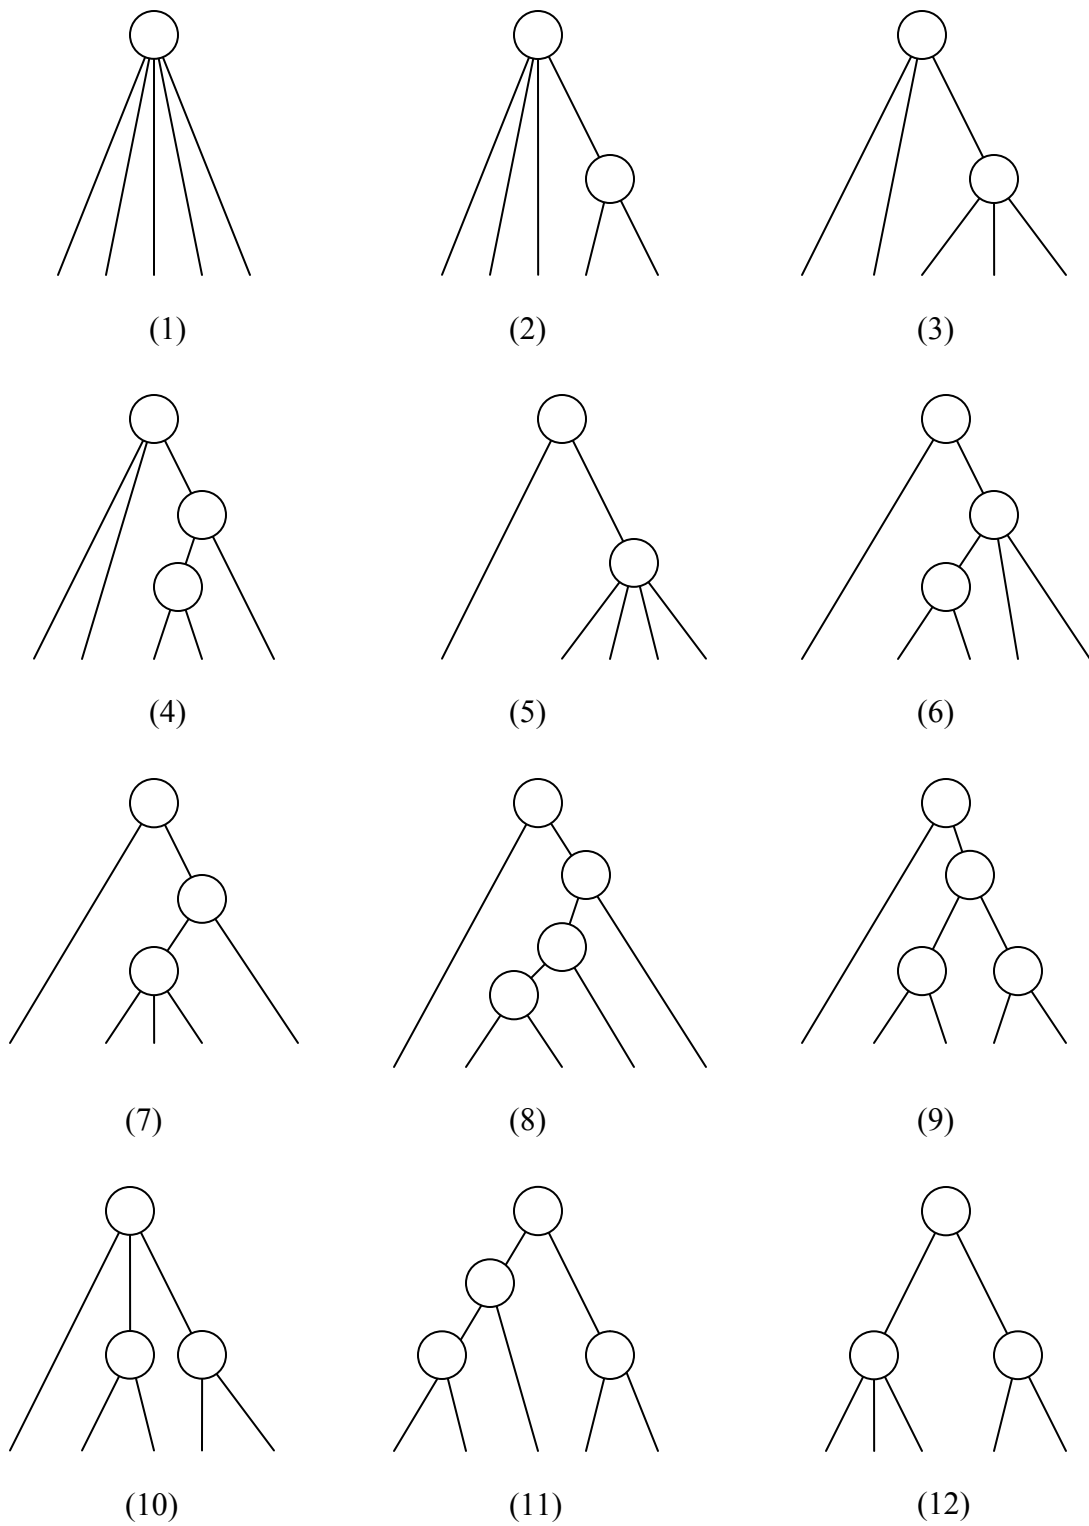


Figure 4: Networks with 5 Bits of Information

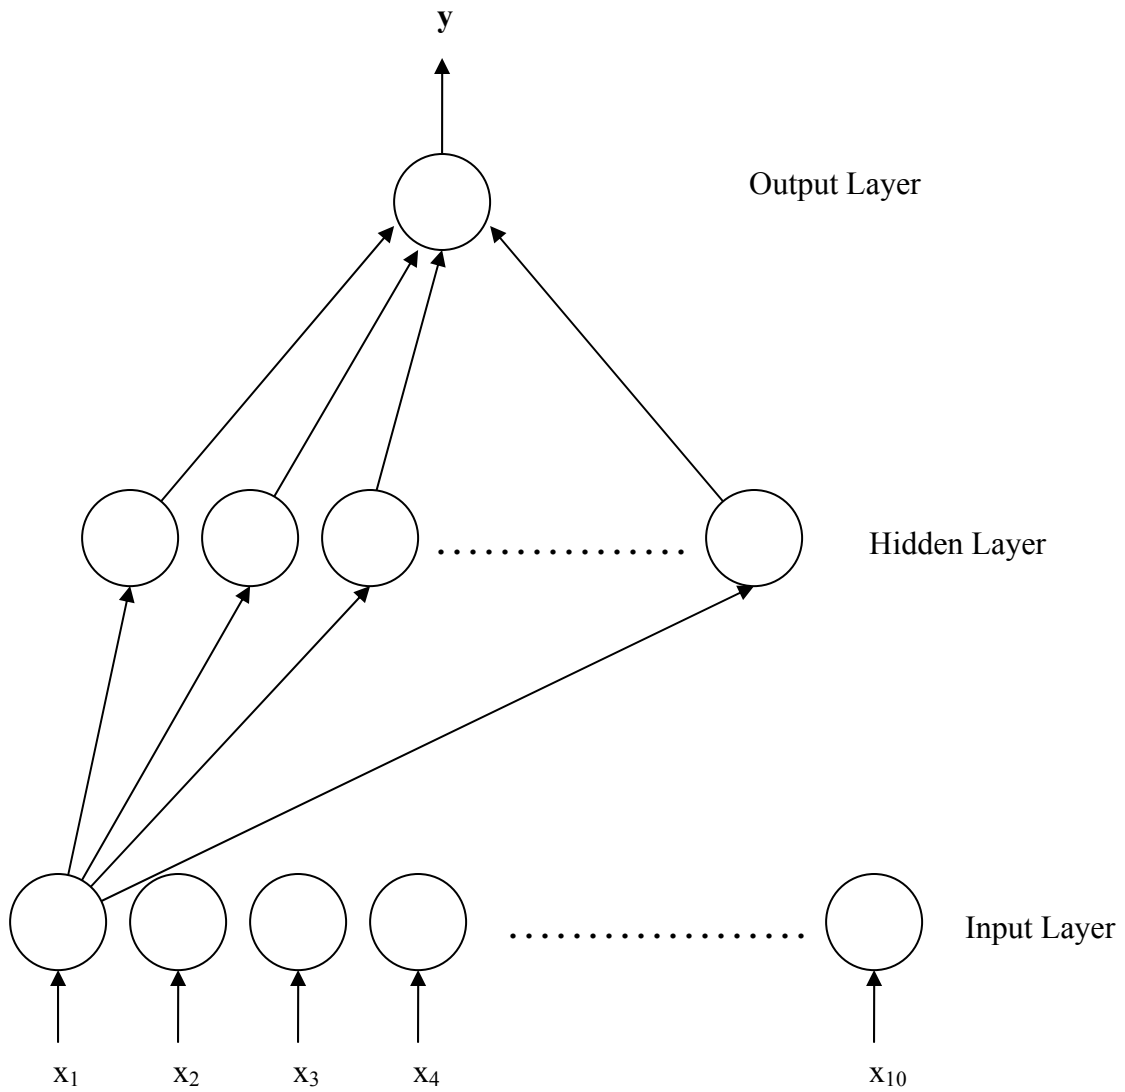


Figure 5: Artificial Neural Network

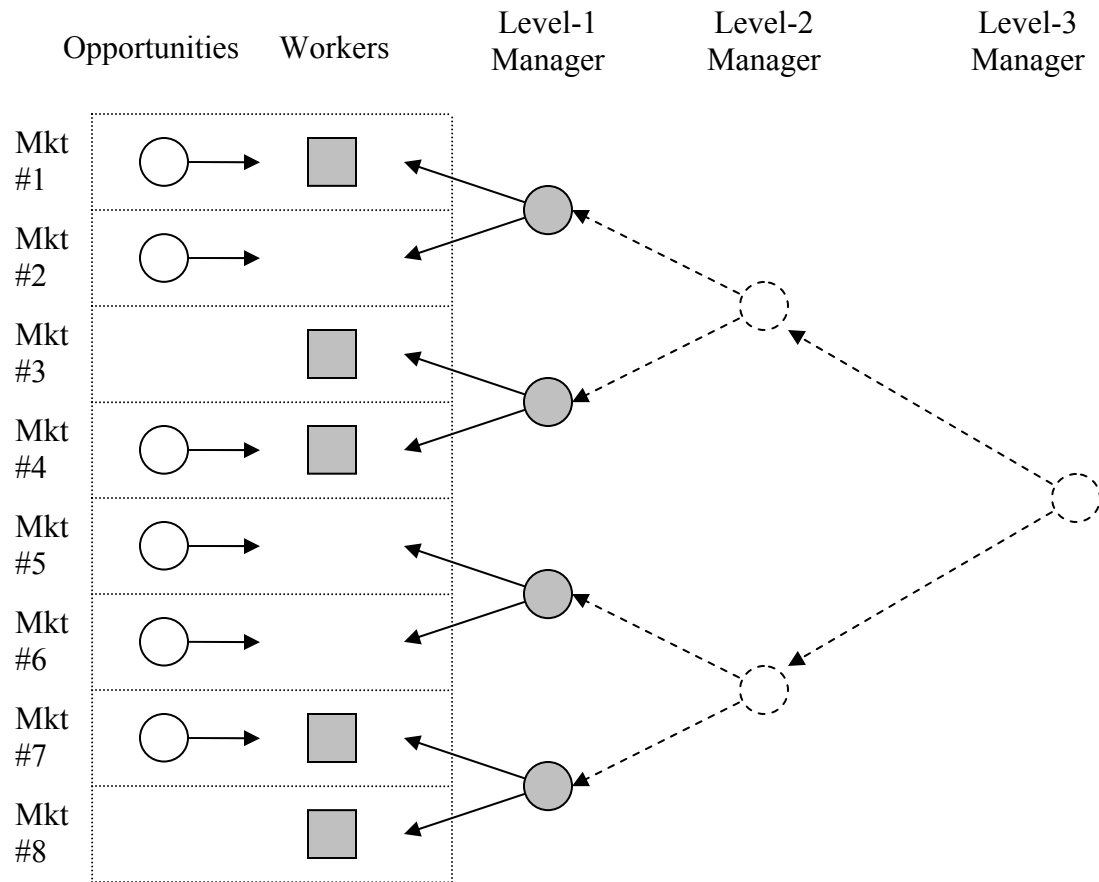


Figure 6: Epstein's Essential Problem