

University of Maine

From the Selected Works of James Wilson

2013

Costly information and the evolution of self-organization in a small, complex economy.

James A Wilson, *University of Maine - Main*



Available at: https://works.bepress.com/james_wilson1/4/



Costly information and the evolution of self-organization in a small, complex economy

James Wilson^{a,*}, J. Hill^f, M. Kersula^a, C.L. Wilson^b, L. Whitsel^b, L. Yan^a, J. Acheson^g, Y. Chen^a, C. Cleaver^a, C. Congdon^c, A. Hayden^a, P. Hayes^a, T. Johnson^a, G. Morehead^f, R. Steneck^a, R. Turner^b, R. Vadas^d, C.J. Wilson^{a,e}

^a School of Marine Sciences, University of Maine, Orono, ME 04469, USA

^b School of Computing and Information Sciences, University of Maine, Orono, ME 04469, USA

^c Department of Computer Sciences, University of Southern Maine, Portland, ME 04104, USA

^d School of Biology and Ecology, University of Maine, Orono, ME 04469, USA

^e Maine Department of Marine Resources, West Boothbay Harbor, ME 04575, USA

^f Department of Mathematics, University of Maine, Orono, ME 04469, USA

^g Department of Anthropology, School of Marine Sciences, University of Maine, Orono, ME 04469, USA

ARTICLE INFO

Available online 29 December 2012

JEL classification:

B52
C63
D23
D64
D83
D80

Keywords:

Costly information
Learning
Adaptation
Self-organization
Learning classifier system
Evolutionary computation
Evolutionary economics

ABSTRACT

The core idea of evolution is that order in living systems emerges from a simple process of variation and selection. In biological systems we usually understand the source of variation as best described by the mechanisms of genetics. If human social systems are evolutionary systems, however, it would seem the variation that most explains the sources of change in these systems, occurs not from a genetic mechanism, but from individual learning. We use an evolutionary computational methodology to explore the way individual learning and adaptation lead to the evolution of persistent, self-organized social and economic activity. The basic idea behind these explorations is that the character and extent of self-organizing social and economic activity depends upon the way the environment frames the costs of individual learning and adaptation. We consider three different kinds of costs affecting learning and adaptation: the costs of autonomous searching, of communicating, and of deciding. Individuals respond to these costs by carefully, i.e., economically, choosing to learn about and interact with familiar agents in familiar arenas in which they have relatively secure expectations about the outcome of their actions. Emerging from these choices are persistent relationships among agents that lead to social and economic structure and to the imperfect coordination of aggregate production. The character and the extent of each are a function of the way the costs of information change with changing natural and human system conditions.

We use a learning classifier system (LCS) to model learning. The logic of an LCS closely mimics the mechanisms of Darwinian evolution, but is applied to the evolution of an agent's decision rules. We describe the environmental context necessary for an LCS to produce economizing behavior and apply the method to a multi-agent simulation of the Maine lobster fishery, which we treat as a metaphor for competition based on the search for useful knowledge. The structure and method of the model is similar to a conventional agent-based model except we use LCS to evolve decision rules for each agent rather than supplying those rules ourselves. This allows agents to change their behavior, i.e., to learn

* Corresponding author. Tel.: +1 207 581 4381; fax: +1 207 581 4990.
E-mail address: jwilson@maine.edu (J. Wilson).

and adapt, as their environment changes. Modeling learning and adaptation as the source of behavioral variation makes it possible to use evolutionary theory to address important questions of social and economic emergence not possible with current methods.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The core idea of evolution is that order in living systems emerges from a simple process of variation and selection. In biological systems we usually understand the source of variation as best described by the mechanisms of genetics. If human social systems are evolutionary systems, however, it would seem the variation that most explains the sources of change in these systems, occurs not from a genetic mechanism, but from individual and group learning. The work that we describe here is intended as an exploration of the way individual learning and adaptation lead to the emergence of persistent, self-organized social structure and dynamics.

Our perspective can be explained with a relatively simple proposition about a basic problem confronting any individual in a complex world: she cannot know everything about her environment; she has limited time and resources and must focus her learning to concentrate on what is most valuable to her interests. The kinds of things an individual might find valuable to learn and how she might adapt are tightly constrained by her cognitive abilities, by her experience, by the resources she has for searching out useful knowledge, by the diversity and variation of her environment, and especially, by the competitive and cooperative actions of other individuals. Faced with this complexity and its accompanying uncertainty, a self-interested individual might be expected to choose to interact with other individuals and with her environment in arenas where her experience helps her to recognize the regularities in both the social and biophysical realms. In these arenas she can arrange her affairs so that she has relatively secure expectations about the outcome of her actions. Consequently, her choice of activities is likely to favor interactions with familiar agents at familiar times and places. This preference can be expected to lead to further interactions that reinforce the advantages of familiarity, resulting in persistent relationships with other individuals and groups. We presume these social relationships are a necessary foundation for the articulation of collective action (Ostrom, 1990) and for the decentralized coordination of production leading to efficiencies that are greater than might be expected from the solitary actions of unconnected individuals.

Analysis of the sources of variation that lead to the emergence of structure and dynamics in social systems usually lies outside the realm of formal economics. Arthur (1992, pp. 1–2), for example, notes:

In dealing with problems that are complicated and not well-defined, economic agents are forced to reason somewhat differently from the way our standard notion of rationality pictures things. They are forced to use inductive rather than deductive means of reasoning. They must rely heavily on internal models of and hypotheses about the problems they are dealing with; and they constantly monitor, update, and revise these by importing feedback—new data—from their environment. Thus in dealing with problems of any realistic complication, economic agents unavoidably live in a world of learning and adaptation.

Even though formal methods in economics avoid problems of learning and adaptation, there is a strong verbal tradition in economics that attributes the structure and dynamics of social and economic organization to learning and adaptation. Hayek (1945) nicely captures this perspective. He describes a complex, particularistic and continuously changing world in which imperfectly informed individuals acquire useful knowledge. To conserve the resources needed to acquire that knowledge, they restrict their search to a relatively small sphere. Within that sphere, individuals are relatively well informed and able to adapt to change; outside, the relevant knowledge of events unknown to them are conveniently summarized by the prices of the goods and services that enter and leave their sphere. When similar behavior is followed by large numbers of individuals in large numbers of nearly independent spheres, it leads to the self-organized, information-efficient coordination of the activities of the individuals in those spheres. In a similar vein, Simon (1962, 2002) argues convincingly that this kind of nearly decomposed organizational structure is close to ubiquitous because it creates strong local feedback that reduces the costs of learning and accelerates the adaptation of both individuals and groups. The strong parallels with group selection theory in biology are hard to ignore (Wilson and Gowdy, 2013; Levin, 1998; Wilson and Wilson, 2007; Darwin, 1859). Nevertheless, as intuitively appealing as it might be, this canonical perspective conceals the mechanisms of learning and adaptation that lead to the emergence of self-organized social and economic activity.

The modeling we describe here attempts to address this problem. We borrow an evolutionary learning approach from computer science, modify it so it treats learning and adaptation as an economic problem, and apply the result to a relatively small, complex economy in which individuals compete and cooperate according to differences in their knowledge of a continually changing, valuable resource. In this instance, the economy is the activity that revolves around harvesting the lobster resource in the Gulf of Maine. This economy is a relatively closed system that is a concrete example of competition through the search for useful knowledge. The methodology we use to model learning is a version of John Holland's learning classifier system (LCS; Holland, 1975, 1991, 1996; Holland et al., 1989). Holland's conception was to create a computational analog of Darwinian evolution that could be used to address questions of adaptation in complex biological and human systems. The fundamental idea is that the evolutionary process is a method of distributed computation that can be used to better understand problems not amenable to conventional mathematical, statistical, or even verbal analysis. The idea of

distributed problem solving, although it has not been known by that name, has been part of the mental model of economists since at least 1776, but, so far as we know, has not been formally modeled.

An LCS can be thought of as a set of rules, or procedures, that govern the way an agent¹ searches its environment and learns by experience. The great advantage of LCS is that it allows us to model learning as an economic activity concerned with the competitive search for useful knowledge. Its framework allows us to develop a working definition of the cost of information and to describe heterogeneous knowledge held by heterogeneous agents. We use LCS to model the activities of multiple, adaptive agents operating in a small, complex environment. The term “adaptive agent” is used to describe a self-interested agent whose behavior evolves as it gains experience and learns about its environment. These adaptive agents are located in a large (for them) and complex world. Each agent has limited resources for learning, so must choose what is likely to be worth learning. The complexity of the environment means the agent's observations and its communications with other agents are always noisy and incomplete. Thus, the meaning and value of the information it acquires at any moment is ambiguous and can only be evaluated in light of its existing imperfect knowledge. Consequently, everything the agent learns is tentative and can only be refined or affirmed by testing, almost as if it were a scientific hypothesis (Holland, 1996).

Agents respond to the imperfection of their knowledge by focusing their actions, their acquisition of information about their environment and their association with other agents in familiar arenas. Repeated interactions between pairs of agents (Gowdy et al., 2013) allow them to better understand the context of each others' actions; as a result, the accuracy and value of communication between them grows, generating more secure expectations about the other agent's behavior and the outcome of their interactions with that agent. These relatively secure expectations lead to persistent individual and group relationships. We argue that the expectation of the continuing value of these arrangements motivates individual restraint and is a necessary precursor (Wilson et al., 2007) of successful collective action (Ostrom, 1990) and efficient aggregate production.

Pursuing these ideas has taken us down a path that departs in significant ways from standard economic approaches. In their introductory article, Gowdy et al. (2013) note three central cosmologies manifest in economic theory: “(1) rational economic man, (2) the invisible hand of the market, and (3) the existence of a general competitive equilibrium.” They argue “that current breakthroughs in evolutionary biology and neuroscience can help economics go beyond these simple cosmologies.” We would add computational advances to their list. If we think of formal, not simply mathematical, modeling as a way to rigorously extend our thinking about the world, then our available modeling methodologies are a significant constraint upon our mental explorations. Advances in the methods of artificial intelligence can significantly relax those constraints because they allow us to model learning and adaptation as an evolutionary, economic activity (Epstein, 2006). Our intuition is that this capability, even if imperfect, can provide new insights into (1) the behavior we consider rational, (2) our understanding of the way individual action leads to emergence of regular social structure and dynamics, and (3) to tractable ideas about non-equilibrating evolutionary systems.

In the next section of the paper, we describe the LCS procedures that govern learning. Following that, we lay out the assumptions needed to create an environment in which the constraints on agents' behavior cause them to act in an economizing way. These assumptions are close to those used normally in economics except we add important assumptions about the cost of information. Section 4 of the paper describes a fairly simple dynamic model of the biophysical environment of the lobster fishery; it is followed by an explanation of the way we couple LCS and the lobster environment. In Section 6, we set our agents loose to learn about their environment and show the results that emerge from their learning including the ways individual and group behavior changes with changes in the costs of search and communication. Tests of the coupled model's qualitative predictions are compared with similar outcomes in the real fishery. The final section summarizes the paper.

2. Agent learning and adaptation

In the paragraphs that follow, we provide a broad-brush description of the computational procedures governing learning in an LCS. Our intent is to give the reader an accurate view of the mechanisms without getting bogged down in unnecessary detail. Online Appendix III elaborates on this description. Someone interested in a more technical description should go to Holland's publications cited earlier. Urbanowicz and Moore (2009) and Lanzi and Riolo (2000) provide excellent reviews of the broad method and the various approaches to LCS, including their strengths and weaknesses. The LCS version we use is called an XCS (an eXtended Classifier System) and was developed by Wilson (1995).

An LCS closely mimics the mechanisms of Darwinian evolution; however, it applies the mechanisms of variation and selection to a population of decision rules governing an agent's behavior instead of a set of traits in a population of organisms. In their most broadly conceived form these mechanisms are (1) the genetic algorithm—a set of computational methods for creating variations in decision rules and (2) reinforcement learning—a means for selecting those rules that better adapt the agent to its environment. A decision rule, or classifier, associates a unique combination of environmental conditions and a particular action. Each classifier is also associated with a predicted reward, an expected error in that prediction, and a fitness value for the rule. A rule is a record of the agent's experience and a prediction of future results. Decision rules take the form:

¹ We use the sometimes irritating word ‘agent’ to make it clear that we are talking about the entities in the computer model. When we are not talking about computer entities we use ‘he,’ ‘she,’ ‘fishers’ and other words that indicate real people.

Environmental conditions = 1,1,2,3,2					Environmental conditions = 1,1,1,3,2				
Matching rules					Matching rules				
Conditions	Actions	p	E	F	Conditions	Actions	p	E	F
#,1,2,3,2	3	3,	.1,	.8	1,1,1,3,2	3	11,	.3,	.6
1,#,2,3,2	0	1,	.7,	.2	1,##,3,2	0	2,	.3,	.6
##,##,3,2	2	11,	.3,	.6	###,##,3,2	2	11,	.3,	.6
1,###,2	1	2,	.3,	.6	1,###,2	1	2,	.3,	.6

Fig. 2.1. General rules can match more than one specific environmental instance. The figure shows five conditions each with a discrete state. The “#” sign represents “wild card” or “do not care,” i.e., that specific input is not considered for that rule. Rules with more # signs are more general rules and apply to multiple states of the environment. A totally specific rule containing no # signs would apply to only one state. Expected payoff, prediction error and fitness are represented by p, E, and F, respectively. See [Appendix III](#).

“If these conditions and this action, then expect this outcome,” where “conditions” refers to a unique set of environmental conditions important to the agent’s decision; “actions” refers to one of a list of possible actions the agent might take, and “outcome” is the average outcome predictive accuracy the agent experienced when the rule was used in the past ([Fig. 2.1](#)).

Completely specific rules specify a narrow set of conditions that apply to a unique instance in the search domain. General rules specify a set of conditions that apply to a range of similar circumstances. For example, in the lobster environment a specific rule might specify a particular depth, a particular time of year, a particular bottom sediment type and an action. A general version of that rule might specify a particular depth and time but expresses indifference to the kind of bottom type. Thus, general rules apply to a broader range of environmental circumstances. Consequently, for any particular circumstance an agent might have several rules that match, e.g., possibly one specific rule and several general rules ([Fig. 2.1](#)). This means the agent is more likely to encounter the circumstances associated with general rules and, because of that, likely to learn their predictive accuracy faster. As we explain later, agents’ tendency to learn different kinds of rules in different circumstances has important implications for the kinds of behaviors agents exhibit in different environments.

An agent can have a large number of specific and general rules each applying to different circumstances in its environment. All its decision rules are called its rule set, which is simply a set of conditional expectations about the likely outcome of its action in a variety of circumstances, i.e., its mental model of its environment. The agent’s mental model is always imperfect because in a changing environment past outcomes are imperfect predictors of current outcomes, and it is always incomplete because the agent does not have the time and resources required to search out all the possible rules that might exist in its environment, the near equivalent of perfect knowledge.

The testing and selection procedures of the LCS allow the agent to refine its rules. In our approach when the simulation is initialized, an environment appropriate to the problem being modeled is created, e.g., a dynamic biophysical model of the lobster fishery, and the agent is given a set of random rules with random predicted outcomes and placed at a random location in the simulated environment. In the first iteration (and all subsequent iterations) of the simulation, the agent observes the conditions of the environment at its current location and checks its (initially random) rule set to see if there are any rules that match the circumstances at its location.

From that matching set, it chooses a rule that is expected to perform well and takes the action the rule advocates, effectively testing the expectation contained in the rule. The agent gets feedback about the performance of the rule and revises its sense of the likely outcome if it were to be used again. In the early stages of the simulation, when the agent has only random rules, the outcome of its actions is often bad. However, as the agent repeatedly tests the rules in its rule set, it begins to identify those rules with better outcomes. Rules that best predict the reward received from the action they advocate are identified as the fittest rules.

A significant source of improvement in agent performance comes through the creation of variations of existing rules. The importance of rule variation is that it allows the agent to explore novel circumstances that are outside the realm of its specific experience. Nevertheless, these variations are constrained by the agent’s experience, as explained below, and can be thought of as actions that are suggested by that experience. There are three ways that new rules are generated.

- (1) The first is through a genetic algorithm that preferentially uses fit rules as parents for crossover and as the source for mutation. Crossover mates the parent rules; for each parent rule, the list of conditions it stipulates is split (at the same random point in the list) with one part going to one new child rule and the remaining part to a second child rule. (See [Appendix III](#) for more detail.) The importance of using fit parents is that the new variations, the child rules, that result are slight variations of their parents and are also likely to be fit rules. In a mutation, a child rule is created with a simple random change in one (or more) of the conditions or the action of a single fit parent rule. If one considers the evolution of rules as a search process, the effect of creating new rules from fit parents tends to localize the agent’s search. The frequency with which crossover and mutation occur is set by the modeler.

- (2) When the agent encounters circumstances for which it has no existing appropriate rule, a new rule that consists of the conditions describing its current circumstances and a randomly selected action is invented. This is called “covering.” It tends to occur early in the learning process when the agent has limited experience.
- (3) In our model, rule variations are also created by the preferential imitation of other agents that appear to show good performance. This method also tends to bring rules that perform well into the agent’s rule set. The patterns of agent interaction that develop from this procedure are an important part of the coupled model and are described in more detail in Section 3.2.

The relatively large number of new rules created in these ways are added to the agent’s rule set. Each new rule is best thought of as an experience-based suggestion about what the agent might expect if it was to take the action prescribed by the rule. Except for the occasional mutation, these newly created rules tend to be conservative because their content, and the breadth of the search they suggest, is based on already fit rules or on the imitation of agents who are well known to the agent and likely to use similar fit rules.

By repeatedly following this procedure the agent begins to refine its expectations about the outcomes it might expect from particular actions in particular circumstances. Rules that accurately predict the outcome of the agent’s actions allow the agent to make relatively well informed choices about the kinds of actions that are likely to be most beneficial to its self interest. These fit rules are selectively retained, even if they are rules that predict bad outcomes. (The agent has to learn what not to do, otherwise it might repeat a costly error time after time). Rules that predict poorly tend not to be used and, depending on the computer memory given to agents, are eventually eliminated (Smith, 2012).

In short, rule creation is the counterpoint of rule selection. Rule creation expands the agent’s rule set by adding new, unexplored variations of existing fit rules. Selection, on the other hand, tends to retain variations that are beneficial while filtering out those that contribute least to the agent’s overall fitness. The combined effect reduces the size and refines the fitness of the set of rules used by the agent. Furthermore, the ongoing creation and testing of rules gives the agent the ability to “tinker” continuously with actions that lie outside its current adaptive solution. This is an essential attribute in an evolving environment.

3. Using LCS to model a social environment

What agents learn and how they adapt depends upon their objective and the environment in which they are placed. Because our intention is to model the resource- and time-constrained learning and adaptation of an agent in a complex environment, we create an environment that is similar to that usually imagined in neoclassical economics; except we treat knowledge as a scarce resource and its acquisition and use as a costly process. If knowledge is not costly to acquire, of course, learning is trivial and not important to economic behavior.

Williamson (1981) argues that transaction costs are important determinants of economic behavior and organization. He distinguishes three dimensions of transactions costs (uncertainty, frequency, and asset specificity) and contends the latter is the most important of the three. Our idea is close to Williamson’s except we emphasize the way path-dependent learning, i.e., experience-specific knowledge, leads agents to choose different actions and strategies. At any particular time and place, the incompleteness and specificity of an agent’s knowledge gives it certain expectations about the likely value of different actions. These expectations largely determine its choice of actions and the direction of its learning. Another agent in exactly the same situation but with a different history might value its alternatives in an entirely different way and take an entirely different course of action. In the paragraphs that follow we outline the modeling assumptions, i.e., the environment, that is consistent with intended, but path-dependent, knowledge-constrained, economizing behavior. We consider three ways in which the acquisition of valuable knowledge is costly.

3.1. Costs of learning through autonomous search

We assume:

- There are multiple, self-interested agents whose behavioral patterns (i.e., their rule sets) are the path-dependent and continuing result of the learning rules incorporated in the LCS described in Section 2.
- The biophysical domain is characterized by a patchy, valuable, and subtractable (but annually renewed) resource with multiscale patterns of various regularity (as described in Section 4).
- Each agent has limited time and resources to devote to the search for and use of valuable knowledge.
- The resource domain is larger than any single agent is able to search (i.e., the distribution and abundance of the resource changes at a rate that is faster than the agent is able to monitor), but at the same time, is small enough that it can be searched nearly completely by all agents.

Combined, these assumptions have the effect of placing all agents in a competitive race for the acquisition of scarce knowledge about a complex, valuable and scarce resource. Because the time the agent can spend searching (or experiencing) this environment is limited, every time the agent acts it incurs the opportunity cost of not knowing the current state of that part of the environment it chooses not to search.

3.2. Costs of observation of other agents

We assume observations of other agents are characterized by the following conditions:

- Agents are able to identify one another as distinct individuals.
- Observations of other agents are noisy and incomplete.
- The noisiness and incompleteness of observations among familiar agents, however, is mitigated by the agents' prior knowledge of the other agents' patterns of behavior.
- Agents tend to ascribe greater accuracy and value to the observation of familiar agents.²
- Agents can gain knowledge (acquire new rules) by observing and copying the action of another agent. (We do not assume articulate communications, simply observation.)

These assumptions allow agents to identify each other and to act differently toward other agents depending on the history of their interactions. We assume agent history is important because the more agents encounter one another the more likely they are to understand the patterns of each other's behavior. This understanding is an antidote to noisy communication; it gives agents a sense of context and the ability to attach greater meaning to such communications. In other words, if an observation of another agent is interrupted or is otherwise imperfect, the knowledge acquired in earlier similar interactions helps the observing agent to fill in the missing information. The result is a more accurate flow of information. For another agent with a different history, the meaning and the value of the same observation might be entirely different and might lead to an entirely different valuation. Nevertheless, the broad assumption is that the greater the accuracy of the information transmitted in an observation, the greater the value of that information. More accurate information gives the agent more secure expectations about the outcome of its possible range of actions and for that reason is valuable. Even accurate information about a likely poor outcome is valuable information. In short, we define the opportunity cost of information as the value lost due to noise and incompleteness in its transmission. (The computational method we use to implement this method is described in [Appendix I](#).)

The important implication is that because each agent has a different history of interactions with each other agent, its costs of communicating—the accuracy of its communications—with each other agent are different; this creates a preference for interactions with familiar agents. That preference leads to still more imitations of and better understanding—positive feedback that becomes the foundation of preferential individual relationships. The cost of this improved understanding, of course, is less understanding of the broader environment and other agents.

As the psychiatrist Marco [Iacoboni \(2012\)](#) writes:

Imitation in humans is widespread and often automatic. It is important for learning and transmission of culture. We tend to align our movements (and even words!) during social interactions without even realizing it. However, we don't imitate other people in an equal way. Perhaps not surprisingly, we tend to imitate more people that are like us. Soon after birth, infants prefer faces of their own race and respond more receptively to strangers of their own race. Adults make education and even career choices that are influenced by models of their own race. This is a phenomenon called self similarity bias. Since imitation increases liking, the self similarity bias most likely influences our social preferences too. We tend to imitate others that are like us, and by doing that, we tend to like those people even more.

The advantages of relationships built on more accurate communications do not increase indefinitely because these advantages are only present when the two agents have different, valuable knowledge ([Fosco and Mengel, 2011](#); [Boyd and Richerson, 2005](#)). But when agents associate closely and imitate one another, they become more similar; the act of imitation reduces the value of continued cooperation, increasing the likelihood and value of autonomous action (or of cooperation with another, third, agent). With the passage of time and changes in the environment, autonomous action (or cooperation with another agent) leads to the acquisition of new and different knowledge that makes cooperative action again beneficial. This dynamic implies serial complementarity between cooperative and autonomous behavior, rather than the usually imagined mutual exclusion.³

3.3. Costs and method of deciding

Agent learning in an LCS requires repeated experience with similar environmental circumstances. This implies an LCS is likely to be most applicable when modeling small, relatively stable environments in which the agent is able to frequently sample its environment with relatively low cost. However, a complex social environment implies a large, dynamic, and

² The ecology literature distinguishes between social exploitation, i.e., unilateral observation and imitation, and the active sharing of information through communication. Our assumption here corresponds to social exploitation, but our presumption is that the familiarity that results from our other assumptions creates the circumstances in which more articulate communications are facilitated, i.e., that social exploitation can lead to active sharing, but the results we report do not depend on that presumption.

³ We distinguish the terms "imitation" and "cooperation" by using cooperation only when the interaction among agents is the result of repeated, path-dependent mutual imitations. In other words, if agents imitate one another persistently with symmetrical frequency because of a historical preference, we describe their actions as "cooperative."

costly search space and a potentially intractable problem for an LCS. However, a decision hierarchy that partitions the agent's search space into a series of relatively small search spaces, using a separate LCS for each space, makes the LCS methodology applicable in a large environment. In this smaller search space, the agent is able to acquire a proportionately larger sample of the environment. Consequently, the agent is able to learn and adapt to its regularities faster than it might adapt to a broader space, just as the people in Hayek's and Simon's views of the world do.

We assume each decision an agent makes can be partitioned into a simple hierarchy that corresponds closely to the spatial and temporal hierarchy of the agent's environment. We design each step in the decision hierarchy so that the conditional variables used by the LCS are congruent with the relevant temporal and spatial scale of the environment, e.g., variables that capture fast-paced, relatively local events for which the value of knowledge decays rapidly characterize the first LCS. Variables characterized by longer-term, more durable knowledge are the focus of the next scale and nearly permanent aspects of the environment are the focus of the third. This spatial/temporal congruency creates tight, coherent feedback that facilitates learning and adaptation (Frenken et al., 1999; Hayek, 1945; Riquelme et al., 2000; Simon, 2002). As Frenken et al. (1999, p. 18) note:

In an evolutionary environment, the decomposition of the system into sub-systems, even if it is only an approximate decomposition, allows agents to improve their fitness in a very short time, even though they usually are not able to find the global optimum. Thus, satisficing strategies aiming at a high rate of improvement outcompete optimizing strategies which aim at the maximum end result.

Finally, the hierarchy is designed so that each step is linked to the scale above and below. This allows the agent to make comparisons between adjacent temporal and spatial scales and to evolve decision rules that reflect what it has learned about, say, the appropriate time to switch away from its current short-term strategy. For example, in the coupled model described below, the decision about whether to continue exploiting the resource at a particular spot or to explore elsewhere involves a comparison between (a) what the agent knows about its immediate spatial and temporal circumstances and (b) its broader-scale, longer-term expectations that are drawn from more durable knowledge about the likely results of exploration. The former is an expectation based on relatively certain, current information; the latter is an expectation that is more uncertain because the system may have changed in ways that the agent's earlier experience might not be able to account for. Thus, in a system with a fair degree of regularity, a hierarchical approach that permits scale-linking rules gives the agent the opportunity to learn the probable trade-offs between current exploitation and a costly search for new knowledge.

A decision hierarchy designed in this way makes the agent's problem tractable in a way that is consistent with the perspectives of both Hayek (1945) and Simon (1962, 2002). The price the agent must pay for tractability is ignorance of the opportunities and threats that exist in a large part of its environment. In a dynamic and volatile environment this could be a high price. However, (in the model below) there are two factors that mitigate the problem: (1) all other agents are similarly impaired, which levels the competitive playing field and (2) observations of other agents give the individual access to broader knowledge of its environment without significantly reducing its ability to localize its own actions.

4. The lobster fisher's environment

4.1. The natural and social domain

Commercial fishing is a good metaphor for models of learning and adaptation. Fishing is all about searching for knowledge about the location of a valuable resource. In the lobster fishery, learning how and where to catch lobsters is not easy. Lobsters are located in a complex oceanographic environment. Tides, currents, and weather create a difficult physical environment. The topography is rugged; bottom types—mud, rock, sand, kelp and other seaweeds—are highly variable at a very fine scale leading to patchy lobster habitat. Dramatic seasonal changes in water temperatures lead to differences in the metabolism and catchability of lobster at different places and depths. Lobster do not form schools like fish, but their local abundance varies, often dramatically, from place to place and from season to season due to storms, changes in water temperature, and especially, the harvesting activities of fishers (Steneck and Wilson, 2001; Incze et al., 2010).

Learning, or as fishers would say, “staying on top of” these changing patterns is essential for competitive success. Fishers use traps to sample their environment and to harvest lobsters. Up to 800 traps (the legal limit) are fished in “strings” of 10–60 traps. A string might be spread out over a quarter square mile of roughly similar depth and habitat. Traps are moved frequently in a non-random pattern. The information driving these movements comes from fishers' own observations of the natural environment and from the action of other fishers and is always incomplete and uncertain. The catch from any single trap is an ambiguous indicator of local abundance; multiple traps in the same locality provide a better picture. Fishers often place a few traps where catch rates are not high but might be expected to increase soon; but the information obtained from these traps is never, by itself, conclusive. Fishers also observe the buoys that mark other fisher's traps, but they can only guess at how well those traps are doing. They can observe the changing spatial patterns of other fishers' traps and guess what that means about where lobsters are most catchable. They occasionally are close enough to watch another fisher as he hauls his traps, but find it difficult to know exactly how much was caught. At the buying station and through local gossip, they can learn whether other fishers are doing well or poorly, and they can keep track of the longer-term success of their competitors by the age and size of their pickups and the size of their houses, although these indicators sometimes say more about personal debt than skill. Fishers continually try to find associations within this scattered and incomplete data; they

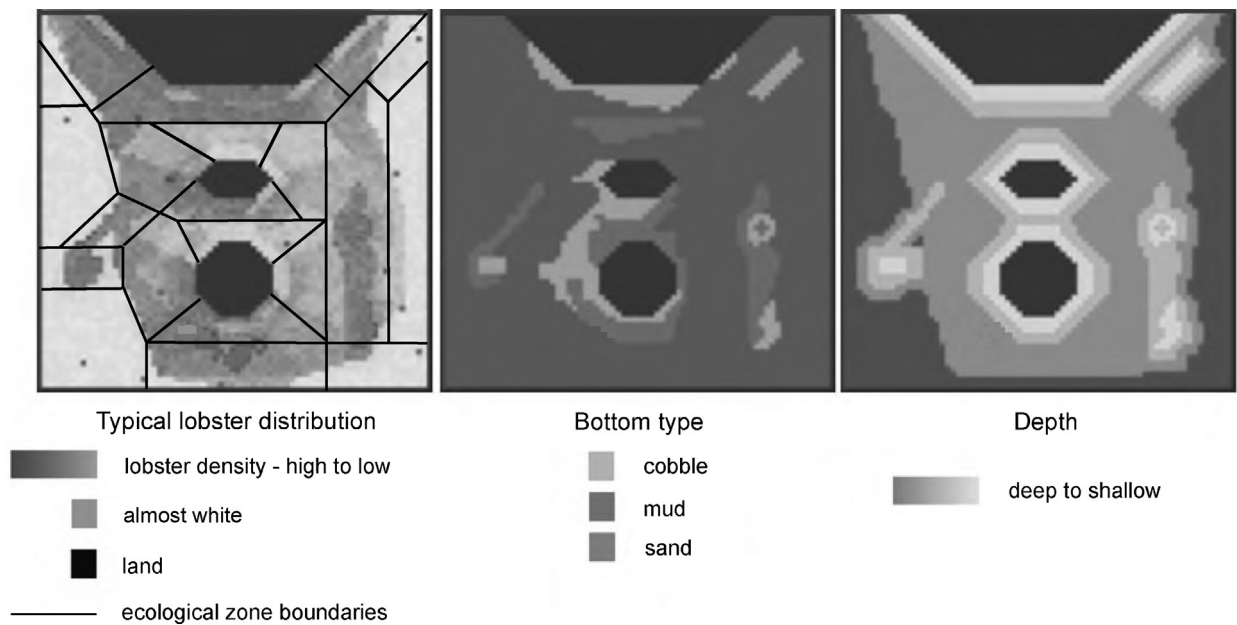


Fig. 4.1. The lobster environment.

build elaborate mental models of lobster biology and of other fishers' behavior. Each trap placement is a test of a hypothesis derived from these mental models.

Familiarity with the behavioral patterns of other fishers is important because it gives a lobsterman a sense of context that helps greatly in his interpretation of his own observations and the detection of patterns in his environment. The value of this information motivates an intensely social and costly activity. Fishers closely observe one another; they imitate one another; and they form long-term individual and group relationships (Acheson, 1988, 2003; Wilson, 1990). The spatial extent and patterns of their behavior mirror the relatively sedentary behavior of lobsters. This increases the frequency of fishers' encounters, leads to fairly durable lobster-fishing communities and to a strong tendency to maintain exclusionary group territories. All this provides a strong foundation for successful collective action (Acheson, 2003; Ostrom, 1990). Interestingly, the extent of a territory, the nature of its boundaries (permeable or rigid), and the size of the groups in each territory depend upon the topography of the fishing area. In the mid-coast area narrow bays ("drowned river valleys") restrict the spatial patterns of fishing, putting fishers in frequent contact with one another and providing convenient landmarks to demarcate boundaries. As a result the groups, or "harbor gangs" (Acheson, 1988), are relatively small and tightly knit. In the eastern and western parts of the coast, in contrast, the topography is more open with fewer natural boundaries. Repetitive contact is less pronounced, and fishing strategies emphasize more frequent individual searching. As a result, territories at either end of the coast tend to be larger; boundaries are more permeable, and groups tend to be larger and less tightly knit.

4.2. Modeling the biophysical domain

We create a relatively simple, spatially explicit, dynamic model of the biophysical domain. The model operates on a 70×70 grid with four depth strata and three bottom types (sand, mud, and cobble) (Fig. 4.1). Each cell is the equivalent of about 200 m on a side; this level of resolution obscures very fine-scale details such as the occurrence of a boulder field or habitat edge or other small features that many fishers argue are hard to learn and important determinants of their fishing success. There are 24 ecological zones (Fig. 4.1) that correspond with different combinations of bottom types, depths, orientation to prevailing winds, and proximity to land.

At the beginning of the (real) fishing year, a large population of undersize lobsters molt, grow, and enter the legal size range, supplementing those left from the previous year. In contrast, at the beginning of each year the model creates a new population (of constant size) with stochastic variation in each cell; we do not model the year-to-year and long-term dynamic of the lobster population because our immediate interest is in the search problem faced by fishers, not the long-term population dynamics of the exploited resource.

Water temperatures and, consequently, the metabolic rate and the catchability of lobster, change in a regular pattern at different depths as the seasons progress (Fig. 4.2). In the summer, catchability tends to be high in warmer, shallow water; in the winter, catchability is higher in deep and, in that season, relatively warmer water. In the real fishery, younger lobster, which are the bulk of the catch, are relatively stationary (Cowan et al., 2007). In the model, we treat all lobsters as stationary within a cell. In spite of this simplification, the effect of changes in water temperature on metabolic rate and catchability




Fig. 4.2. Water temperature by depth and month in lobster season.

produces roughly the same spatial/temporal distribution of lobster catchability as exists in the real fishery; that distribution is what is important to the agent's fishing location decision at the scales modeled here.

There are three spatial and temporal scales in the model. The predictability of the location and catchability of the resource and, consequently, the durability of agents' knowledge are different at each scale. At the cell level, the distribution of lobster is somewhat unpredictable. The numbers in each cell at the beginning of each year vary by no more than 20%; however, later in the year much more significant changes occur due to the harvesting activities of agents. In any neighborhood, a few cells on a side, fishing can change the abundance of lobster rapidly, almost overnight, leaving an unpredictable local patchiness whose pattern depends on the particular spots fished (by any agent) recently. As a result, for an individual agent, the value of knowledge about the abundance of lobster in any particular cell level or small neighborhood has a rather short life measured in days.

At the moderately broad scale of the ecological areas, there is a fairly regular spatial/seasonal pattern of lobster catchability driven by changes in water temperature. Compared with the cell level, this scale shows a slower rate of change; it is not affected as much by agent activity and is more predictable. The dominance of seasonal biological processes at the scale of ecological areas imparts much more regularity to these areas than occurs at the cell level, but fishing can disturb the patterns of abundance within a year and from year to year. At the broadest scale, the topography of the environment—water depth, bottom type, and orientation to prevailing winds—is held constant. The constancy of the broad-scale physical environment means that agents easily acquire durable knowledge about phenomena at this scale.⁴ Consequently, this knowledge is of less competitive value than up-to-date knowledge about the abundance of lobster in particular localities—knowledge that other agents might not be able to acquire easily.

In summary, the biophysical model is a relatively simple representation of the spatial and temporal dynamics of the lobster population; it generates patterns whose regularity, or predictability, varies according to scale. At the local (cell) level, especially, these patterns tend to be much less predictable than they might be in an unfished state because they are altered, often rapidly, by the harvesting activities of the agents themselves. At the next broader scale, ecological areas and seasons, the patterns are driven by seasonal changes and are more regular, slower, and less disturbed by the activity of agents. At the broadest scale, the physical circumstances of the ocean bottom are treated as if they were permanent and not changed by the agents. The differing pace and predictability of biophysical change at these different scales (O'Neill et al., 1986) creates different learning problems for agents and, therefore, different reasons to cooperate and compete at each scale. Further, even though the fine-scale distribution of lobster is close to unpredictable at the cell level, the regularity of human behavior at that scale generates a fairly regular dynamic in which persistent individual relationships, groups, and networks can be built (Cohen et al., 2001) as we show below.

5. A coupled biophysical and multi-agent LCS

In this section of the paper, we describe how the LCS procedures were coupled with the dynamics of the biophysical model. The objective of the coupled model is to create the circumstance in which an agent has to learn how to use its limited resources to search for and harvest lobsters. This means the agent must learn the natural patterns of lobsters, the way its own activities and those of other agents disturb those patterns, and how it might adapt its own behavior to the patterns of both lobsters and other agents. This learning takes place in an environment that is always disturbed by the co-evolving, competitive behavior of the agents who exploit it—in the language of ecology, “a dynamic fitness landscape” (Wright, 1932).

⁴ In the real fishery, seasonal changes in water temperature can begin a month earlier or later than usual. We don't introduce this variability; the resulting constancy means agent knowledge about the ecological scale is durable.

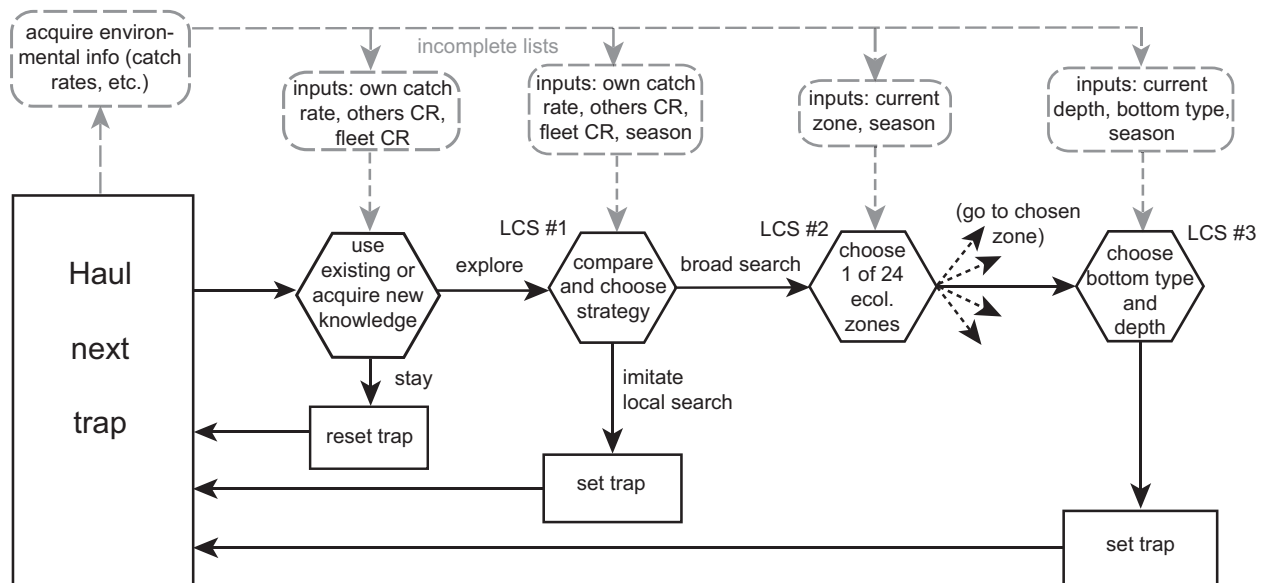


Fig. 5.1. The agent's decision hierarchy.

5.1. The agent's decision

The agent's decision is the equivalent of the decisions all economic agents must make to "position" their product or service so that it competes well against their competitors', e.g., the decision a retail chain makes about the placement of a new store, an airline makes about where and when to schedule flights, a manufacturer makes about where to place its products in "product space," or an individual makes to acquire a specialized skill or other knowledge. Each of these decisions affects the agent's immediate economic environment and provokes a competitive or cooperative response.

In the lobster model, there are usually 30 agents. The behavior of each one is independently governed by the learning procedures of an LCS, that is, each one has a unique set of experiences and rule set. The agents' problem is learning when and where to place their traps. Agents make three trap⁵ placement decisions a day, 240 days a year (July through February); we usually run the model for 20 years. At the beginning of the simulation, all agents are placed on a level playing field. They have no tested knowledge of their environment; instead, each is given a unique initial random rule set, and traps are placed at random locations on the map. All agents are assumed to have comparable boats, the same number of traps, the same fixed daily operating costs, and they all receive the same price for their product. A simple accounting program keeps track of agents' daily costs of operation and revenues; longer-term costs are not included. Both prices and daily operating costs are held constant; consequently, profits are determined solely by the agent's catch. During the year as the resource becomes depleted and daily profits fall, an agent may decide to leave the fishery for the remainder of the year. All agents resume fishing at the beginning of the next year, retaining the rules they learned previously.

Fig. 5.1 and Table 5.1 illustrate the decision hierarchy agents use to address the question about where to fish next. The process starts when the agent hauls a previously placed trap. If lobsters were caught, they are removed, creating a small change in the state of the biophysical environment. The agent notes the catch and revises its expectation of the likely performance of the rule used to place the trap. At the same time, the agent observes the current state of its environment and obtains information about local and broader conditions—bottom type, depth, its own catch rate, average fleet catch rate, season. This information is used at various points in the decision process as shown in Fig. 5.1 and in more detail in Table 5.1.

In the first step of the hierarchy, the agent compares the current performance of the trap it just hauled with the current average performance of the traps of all other agents. Using that comparison, it makes a simple decision about whether to exploit its current knowledge (a) by continuing to fish in the same place(s) or (b) whether to explore or imitate so that it can acquire new knowledge about the state of the resource at other locations. At this point in the hierarchy, all the information used by the agent is current and its value has a very short life. There is no learning associated with this decision; we supply the rule as in a regular agent-based model. The rule is simply: if the agent's catch is above average, it stays; otherwise, it moves. If the agent decides to let the trap stay in its current location, it resets the trap and goes to haul its next trap. If it decides to move to the location of one of its better-performing traps, it moves and sets its current trap and goes to haul

⁵ We use the term "trap" here, but fishers set strings that may consist of 10–60 traps. The spatial resolution of the model is appropriate to the location decision about where to place a string of traps. Consequently, at the scale we are modeling the fishery, three decisions a day is a reasonable match with the frequency of location decisions in the real fishery.

Table 5.1

The agent's decision hierarchy.

Steps in the decision hierarchy	Conditions—information inputs	Action possibilities	Feedback
Step 1—Exploit or acquire new knowledge	Compare own current catch rate (CR) relative to current broad scale average (simple fixed rule—no learning)	(a) Stay in same cell, reset trap , or (b) Move to ownbest and set trap , or (c) Acquire new knowledge, i.e., go to step 2	None
Step 2—LCS #1 Compare performance and choose strategy for acquisition of new knowledge	<ul style="list-style-type: none"> • Current CR vs. own best • Current CR vs. other's best CR • Own best CR vs. other's best CR • Own current avg. CR vs. own previous avg. CR • Current avg. public CR vs. own avg. CR • Current avg. public CR vs. previous public CR 	(a) Local search, set trap (b) Imitate, set trap (c) Broad area exploration, and ... go to step 3	The feedback used to revise the performance measures of all LCS rules is simply the agent's catch after it sets or resets the trap.
Step 3—LCS #2 Explore, choose broad area	<ul style="list-style-type: none"> • Current subarea (1 of 24) • Current month (1 of 8) 	Choose one of 24 subareas, and ... go to step 4	
Step 4—LCS #3 Choose bottom type and depth	<ul style="list-style-type: none"> • Current month (1 of 8) • Current depth (1 of 4) • Current bottom type (1 of 3) 	Choose one of • ... 12 actions giving the depth and bottom type combinations, plus • Deeper, other bottom • Shallower, other bottom • Same, other bottom and ... set trap	

its next trap. If the agent decides the catch of its current trap and its other traps is too low it proceeds to LCS #1 where it considers how or where it might move its trap to acquire new knowledge about the resource.

In the second step, LCS #1, the agent has to learn whether to acquire new knowledge (a) by imitating another agent, (b) by searching in the immediate neighborhood, or (c) by engaging in autonomous exploration on a broader scale. In all three cases, the lifetime of the knowledge the agent uses is longer and the certainty is less than the instantaneous and short-lived knowledge it relies upon in the first step (because the older knowledge may no longer reflect the current conditions in the fishery). Imitation depends on observation of another agent's recent catch (but that observation is tempered by the longer-term knowledge it has of the other agent). Imitation leads directly to the placement of the trap near the imitated agent's trap. Local search is an action based on the agent's very recent experience in the cell it currently occupies and its longer-term knowledge of how local search has paid off in the past. A decision to search locally means the agent moves its trap to a random location in the immediate neighborhood. If the agent decides to explore, the principal data the agent uses are its current performance relative to other agents and its fairly durable, longer-term knowledge about the regularity of seasonal conditions, i.e., the changing catchability of lobsters at different depths in different seasons. When a decision to explore is made, the agent moves on to the second LCS.

In the second LCS (step three in the hierarchy), the agent's decision about where to explore depends upon what it has learned about the particular attributes of the 24 ecological zones. This knowledge is based upon the seasonal regularity of the biophysical environment; agents can disturb lobster abundance at this scale, but not to the extent or with the rapidity of events at the cell level. As a result, knowledge of the regularities at this scale is more certain and relatively long-lived.

Once the agent has moved to the chosen zone, it uses LCS #3 (the fourth step in the hierarchy) to decide the combination of depth and bottom type to place the trap. This decision depends upon durable knowledge about the catchability of lobsters in different seasons at different water depths and bottom types. After the agent places the trap, it moves to and hauls the next trap.

As the agent moves through this hierarchy it does not consider all the possibilities suggested by the permutations and combinations of conditions and actions. Instead each decision is constrained to a choice among the actions that appear on the list of its decision rules that match the conditions at its current location. In other words, the decision hierarchy and the methods of the LCS lead to a computationally efficient and parsimonious use of information. For each agent, this procedure is repeated 720 times each year, 14,400 times for a 20-year run. For all 30 agents, the total is 432,000. Over time the procedure has the effect of increasing the agent's focus on a particular part of its environment; it also increases the agent's ignorance of other large parts of its environment. As the agent repeatedly goes through this hierarchy, its experience accumulates in three imperfect rule sets, each focused on a different spatial and temporal scale. All three LCSs together give the agent an imperfect mental model that guides its daily and longer term activities.

6. Emergent dynamics of the coupled system

The model generates a rich set of outcomes that vary in important ways depending on the parameter values chosen and the inputs provided. To understand the model, we went through a set of exercises that allowed us to explore its internal consistency and its qualitative correspondence with the real world of the lobster fishery. Our initial goal for the current model⁶ was to locate a set of parameter values and model specifications that produced a qualitatively reasonable picture of the fishery. We conducted extensive explorations of the model, changing parameter values and specifications until the behavior was internally consistent. If the model produced output that was unlike our expectations, we explored the result, comparing it with qualitatively known behavior until we understood whether our expectations were not well formed and wrong (this happened often) or whether the specification needed to be changed (this also happened often). This process helped us to understand the dynamics produced by the LCS. Eventually we settled on a default model whose parameter values were internally consistent and yielded a good qualitative fit with the data available for the actual fishery (as described in Section 6.4). We used the default model as a base from which we conducted extensive multi-parameter sweeps, emphasizing changes in the costs of search and of communication. The results of those sweeps and our understanding of the model dynamics are reported in the paragraphs that follow. We explain the behavioral patterns of individual agents and two important aggregate phenomena that emerge from that individual behavior: (1) efficient aggregate production without centralized coordination and (2) persistent individual and group organization. Where possible we compare the patterns generated by the model with patterns observed in the real fishery. The model lends itself to extensive “what if” experimentation. Because of space limitations, however, we do not report any experimentation here (see Wilson et al., 2007, for an example of how this can be done).

6.1. Dynamic patterns of the model

The first day of the simulation occurs in the summer when water temperature and lobster catchability are rising rapidly in shallow water. The map (Fig. 4.1) is populated with lobsters at different depths and locations in a fairly regular pattern according to what is known about their distribution (Steneck and Wilson, 2001; Incze et al., 2010). We create 30 autonomous agents, each employing an independent implementation of the LCS decision hierarchy; each agent is given an initial set of random rules (one set for each LCS in the hierarchy) and three traps. The traps are placed at random locations on the map. On the second day, the agents haul their traps, learn what they have caught, observe the local environment (water depth, bottom type, etc.) and estimate the catch of nearby agents (Section 3.2 above). Using this information, they decide where to place the trap next. This simple procedure is repeated three times daily and simultaneously for all agents through the 8 months of the fishing year. At the beginning of the next year, new lobsters are recruited into the harvestable population, and agents resume fishing using the experience (rules) they acquired in the previous years. Usually the model is run for 20 years.

6.2. Patterns of individual behavior

As agents gain experience, they begin to associate non-random differences in their catch with the different environmental circumstances where their traps are placed. They judge the adequacy of their own performance (measured as catch per trap haul) at any moment according to their imperfect understanding of the performance of other agents. Depending on what they observe, they may or may not decide to move the trap to a new location. For example, the initial random (summer time) trap placement may have put one of the agent's traps in 600 ft of water on a mud bottom where it might produce one or no lobster; on the same day, one of the agent's other traps located in 30 ft of water on a rocky bottom might produce four lobsters. This limited experience might lead the agent to move the trap in 600 ft of water to the same location as the trap in 30 ft of water and to not move the trap in 30 ft of water.

Over the course of the simulation, each agent repeatedly samples its environment in this way. It compares its performance with the catch in its other traps and with its estimate of the catch of other agents and then decides what to do next based on its experience to that point. Typically, when an agent moves to a new location, it exploits the resource at that location until its catch rate falls below that of other agents (Fig. 6.1). Then it moves to a new location and continues moving until it finds a spot where its catch is equal to or above what it believes other agents are getting. This pattern continues indefinitely, creating a catch record that resembles a “ragged saw tooth.” When viewed in aggregate in both the model and in the real fishery, this continuous repositioning generates a coherent seasonal pattern that reflects the changing abundance of lobsters (Fig. 6.2a and b).

At different times of the year and when the model is run with different levels of lobster abundance, the saw tooth shows different patterns. At the beginning of each year, when the resource is abundant and has a fairly regular spatial gradient, agents tend to move less often because it takes longer to fish out a local area, and when they do move, they find a good spot rather easily. As the year progresses, however, harvesting by all agents reduces average resource abundance at a broad

⁶ The computational model was built on code for another version of LCS (ZCS, Wilson, 1994) that had been applied in the same (lobster) environment (Wilson et al., 2007).




Fig. 6.1. Agents move traps when their catch rate falls below average.

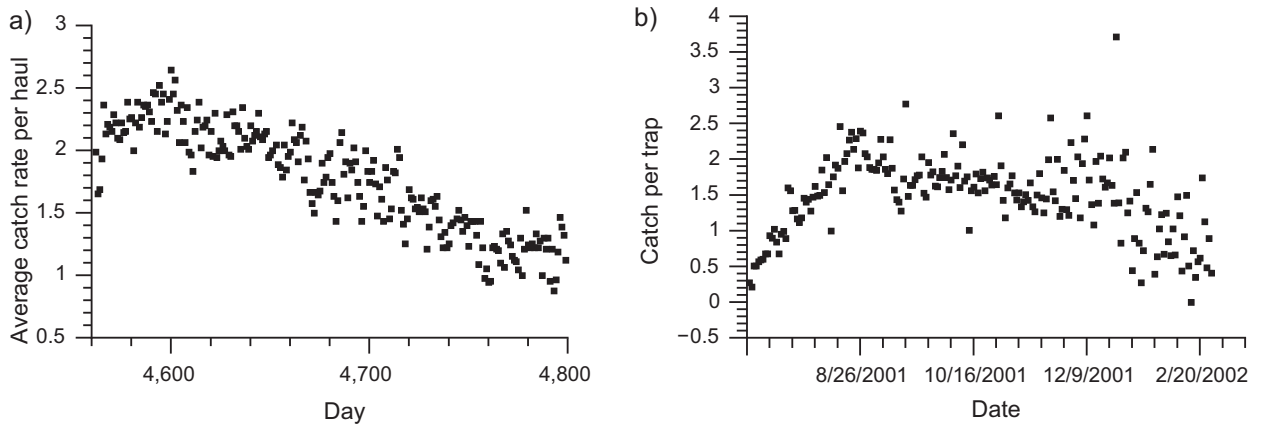


Fig. 6.2. Average catch per haul per day for thirty agents during the final year of the simulation (year 20) as shown in panel a. The maximum catch for each string is set to five. Catch per haul from actual fishery is shown on right, panel b (Wilson, unpublished data). The differences that occur in the early part of the year between the model and actual fishery data are due to differences in beginning of the year recruitment. In the model it is abrupt; in the real fishery recruitment is gradual. We have not incorporated a population recruitment function in the model.

scale and, importantly, at the same time transforms the original fairly regular local lobster distribution into a nearly random patchwork of depleted areas that get larger as the year progresses. In response, agents choose to explore and imitate more often and when they do are less likely to find a productive spot. Similar shifts in the kinds of actions agents choose are observed when different runs with different levels of lobster abundance are compared. When lobsters are scarce, agents emphasize actions meant to acquire new knowledge about the resource; when lobsters are abundant, agents favor the exploitation of their current location (Fig. 6.3).




Fig. 6.3. Exploit and explore choices. The proportion of strategy actions classified as exploit actions and explore actions, for different levels of lobster abundance. The error bars indicate the range of variation over ten runs of the default model at each level of abundance.

Spatio-temporal scale	Resource regularity	Average rule fitness (standard deviation)	Average rule generality (standard deviation)
Fine scale - LCS #1	Unpredictable	0.2332 (0.0030)	0.5946 (0.0047)
Ecological scale - LCS #2	Regular	0.3184 (0.0010)	0.2571 (0.0041)
Global scale - LCS #3	Stochastic stable	0.8286 (0.0009)	0.0226 (0.0008)

Fig. 6.4. Predictability, fitness and generality of all rules by LCS (year 20 of the default model).

6.3. Predictability, general and specific rules

Initially, as agents are learning how to fish, they rely upon relatively general rules because they encounter the broad conditions that match general rules with greater frequency than they encounter the less likely circumstances that characterize specific rules. Consequently, they more easily learn the predictive accuracy of more general rules. As time passes, however, agents may be able to gain more experience with specific environmental circumstances, and their understanding of the predictive accuracy of specific rules improves. For example, initially they may learn that fishing inshore in the summer is a good general rule; however, as time passes a more specific rule that advocates fishing inshore in the summer on rocky bottom on an exposed shore may develop.

This pattern of usage of general and specific rules is also observable at the different spatial scales of the model. At a fine spatial scale (cell level), the abundance of lobster can vary dramatically from one day to the next, depending on the activity of agents. Even when agents are thoroughly familiar with the environment, they have little ability to predict at the scale of the cell; only those agents who have fished the cell very recently have relatively accurate knowledge about its current state. The resulting idiosyncratic distribution of lobsters means agents have a hard time developing fit rules. Nevertheless, they do find fairly general rules that perform (predict) better than others and, if the rules predict a good payoff, they use those rules more frequently (Fig. 6.4).

At a broader scale, ecological zones, the resource is less variable because it is dependent on longer-term, regular seasonal changes in water temperature. Thus, even though the ability to predict lobster abundance in any cell within a zone on a particular day may be low, on a broader scale, the season of the year is a good indication of the ecological zones where lobster catchability is likely to be highest. These differences in predictability at different scales are reflected in differences in the fitness of agent's rules at the relevant scale (Table 6.1). They also lead to differences in the way agents compete at different scales. At the very fine scale, an agent who possesses certain knowledge of a particular spot has a competitive advantage, but that advantage may quickly disappear because the agent or other agents can rapidly deplete the spot. These rapid changes in the local environment mean agents must continually monitor and respond to their perceptions of other agents' catch. However, at the broader spatial and longer temporal scales associated with ecological zones and the broadest scale environment, the seasonal regularity of changes in lobster catchability means that all agents receive fairly consistent feedback and quickly learn the relevant patterns and develop fit rules. Knowledge of the broadest scale confers no competitive advantage since all agents have the same knowledge. On the other hand, if an agent does not possess this knowledge, it is severely disadvantaged.

6.4. Aggregate efficiency

Because agents are able to observe one another and compare catches, even if imperfectly, they are individually and collectively able to find better places to fish. When agents compare their performance with others, they move traps so that few are placed in areas or depths where the catch rate is significantly below or above the current average. If an agent finds an area where the catch rate is high, other agents quickly imitate, bringing down the catch rate in that area. If an agent places a trap in a cell with a low catch rate, it immediately moves the trap to an area of higher catch. In short, even though an agent has only a limited ability to learn about the environment using its own resources, its imperfect ability to observe other agents, to compare its performance with theirs and to change its behavior in response, allows the agent to focus its activities in a way that is beneficial to its own interests. Consequently, although each agent, when viewed in isolation, is working with limited and imperfect information, the aggregate effect is to refine all agents' information, thereby leading to an efficient spatial allocation of traps. This aggregate effect shows up strongly in both the modeled and the real fisheries (Figs. 6.5 and 6.6).

In a stationary or regular environment, accurate communication among agents might be expected to lead to an "ideal free distribution" in the language of biology or a state of "equal marginal profit" in the language of economics. In this environment, however, the dynamic is more accurately described as the active movement away from equilibrium because the continuous, competition-driven alteration of the environment by the agents themselves means each agent has to continuously change its current "solution" to the problem of catching more lobsters.

A disinterested observer of this process might reach a different conclusion about the basic behavior of the system depending on the temporal scale of the statistics she used. When the results are averaged over the long term as in Fig. 6.5a and b, there appears to be a remarkable level of efficiency associated with trap placements. However, if the dynamic is observed for

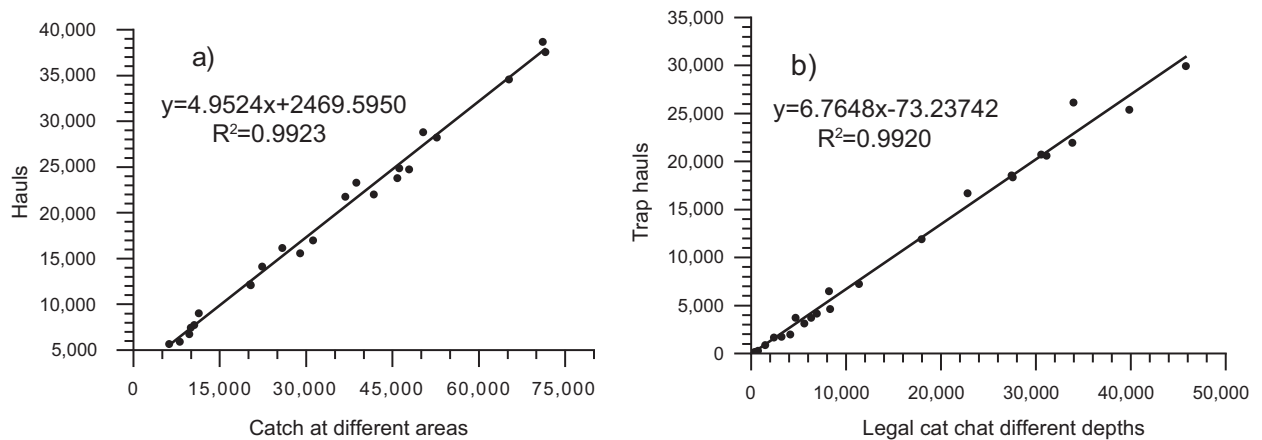


Fig. 6.5. When the spatial placement of traps relative to the availability of lobster is viewed as an average over a long period (20 years in the model, four in the real fishery) (Wilson, unpublished data) it is apparent that agents (panel a) and fishermen (panel b) allocate their traps in space very effectively. As the slope of the line depends on the cost of search and density, it is the regression value and not the slope that serves as a metric for efficiency. Default model.

a shorter time, the observer would likely conclude the system is continuously struggling to correct dynamic inefficiencies and that the extent of the inefficiency was a function of the cost of communication. Fig. 6.6 shows how the pace of adjustment, the dynamic inefficiency, is different depending on the costs of information. The high frequency response on the left (Fig. 6.6a) occurs when the costs of information are low; the less rapid response in the right-hand panel of the figure reflects higher information costs.

An aspect of this dynamic that is not apparent at first is that the incompleteness and diversity of agent knowledge plays an important role in aggregate efficiency (Page, 2007). Because each agent has a different history, each agent's vision of the environment is slightly different, and as a result, each agent has a different perception of the opportunities, especially the very local opportunities, that might exist in the system. In a patchy and complex environment, the result is a diversity of individual histories and knowledge that distributes the activities of all agents in a way that fairly comprehensively reflects the fine-grained distribution of the opportunities in the system.

6.5. The emergence of persistent individual relationships and groups

As described in Section 4.1, one of the more interesting aspects of the lobster fishery is the spontaneous formation of harbor gangs that maintain exclusive territories (Acheson, 1988, 2003). These gangs are important because they generate persistent individual relationships and group social structure that is the foundation of a formal fisheries management system. The size and persistence of territories and their associated harbor gangs is not uniform, but we hypothesize, varies according to the geographical circumstances and changing fishing practices along the coast. Geography is important because it has a strong effect upon the spatial patterns of fishing. It affects the frequency with which fishers contact with one another and, consequently, the relative value of autonomous search and communication with other fishers. We cannot document these




Fig. 6.6. Fifty day running average of the proportion of exploit actions (stay and own best) for year 10 through year 12 for all agents. The left graph (panel a) is derived from a run with moderate communication costs; panel (b) is from a run with high communication costs. The horizontal line represents average proportion of time agents decide to exploit. Communication cost is an arbitrary parameter value; 0.75 is the equivalent of very low costs; 1.0 very high.




Fig. 6.7. Imitations and individual relationships as a function of communication costs. Average bilateral imitations show the increasing tendency to rely upon ‘tighter’ relationships as the costs of communication rise. Average maximum imitations show the tendency to rely heavily on strong bilateral relationships as communication costs rise. (Average values for ten runs at each level of cost. Error bars are not shown because they are very small.)

differences along the coast except by personal observation; however, in the model, we can alter the relative costs of search and communication and show how differences in the cost of search and communication change group size and attributes.

In the model, groups evolve from the associations developed by pairs of agents. In the early days of a model run, two agents may begin to interact simply because their initial random locations put them in proximity. If this happens with some regularity, they become more familiar. After a while if an agent observes two other agents who have the same actual catch, the observing agent will give greater weight to its perception of the agent with whom it is more familiar and, if imitation is called for, will imitate that agent. Each imitation creates another “encounter” that reinforces the agents’ familiarity and increases their tendency to imitate one another. Thus, the formation of individual relationships is a gradual process that builds toward circumstances in which both agents imitate one another with about the same frequency. The longer the model runs, the more symmetrical the relationship and the more the agent is biased toward imitations of its partner. (The evolution of group formation is shown for high and low communication costs in [Online Appendix II](#).) In spite of the strength of these relationships, it is always possible that the agent perceives that another, less familiar agent’s catch might be high (because of an erroneous optimistic perception of a moderate catch or an accurate perception of a good catch) and, as a result, the agent might imitate a less familiar agent.

Nevertheless, especially when the costs of communication are high, the positive feedback of familiarity is strong and individual agent relationships are characterized by intense bilateral interactions. Agents often establish these relationships with more than one other agent, so for the system as a whole this leads to a pattern of persistent, multiple, small groups

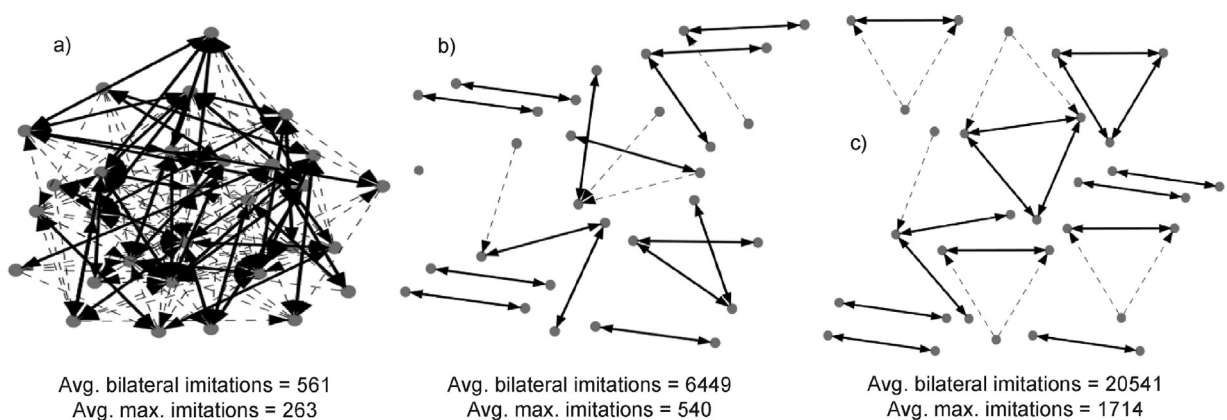


Fig. 6.8. Individuals and groups with different costs of communication. The three panels show individual and group relationships with three different costs of communication. Panel (a) on the left is equivalent to almost free communications: panels (b) and (c) show increasing costs. Heavy lines show symmetrical bilateral relationships; light dashed lines show asymmetrical relationships. The average bilateral relationships noted at the bottom of each panel indicates the number of imitations taking place within symmetrical bilateral relationships; average maximum imitations indicates the strength of individual ties. In panel (a), the strong presence of unilateral imitations (light lines) demonstrates that when the cost of communication is low, agents have a weak or no preference for the agents they choose to imitate, even symmetrical bilateral ties are weak. Thus no apparent social structure emerges. In panel (b), moderate costs, the larger number of bilateral imitations among agents indicates that as the costs of communication rise, there is a stronger tendency to imitate familiar agents and groups start to form. In panel (c) when the costs of communication are high the bilateral relationships among agents are very strong leading to stronger group ties. Each diagram is for a single run of the model. Variations among runs with the same parameter values are small.

(Figs. 6.7 and 6.8 panel c and Appendix II). Within each group, agents tend to communicate intensely with one another, but have infrequent communications with agents outside the group.

When the costs of communication are low, bilateral relationships are much less intense, i.e., less positive feedback and less frequent interactions. Each agent tends to imitate a larger number of agents, and the size of groups is larger when compared with the results when communication costs are high. If the costs of communication fall toward zero, all agents tend to imitate all other agents, but very infrequently; personal relationships among agents weaken, and there is either one or no group, depending on how one defines a group (Figs. 6.7 and 6.8, and in Online Appendix II). Put differently, when agents are able to communicate without cost, the group circumstances that are consistent with Ostrom's preconditions for successful collective action (1990; Gowdy et al., 2013) are not likely to occur because there is no reason for individuals to cooperate.

7. Summary

Evolutionary theory is undoubtedly one of humankind's great ideas. Darwin proposed that an (unspecified) mechanism for producing variation in behavior and a means for selecting those variations that better adapted the organism to its environment was sufficient to explain the evolution of order in the natural world. In the past century biologists have generated a deep understanding of the genetic mechanisms of variation and ecologists have clarified the order that arises from persistent selection. Ultimately the intention of economics is to understand the source of dynamic order in human social systems. The basic ideas of economic competition capture half of Darwin's idea but the mechanisms of variation have not had much attention in economics because of the difficulty of understanding the aggregate patterns that arise from individual learning. The broad idea motivating this work is that the character and extent of self-organized social and economic structure depends upon the way the environment frames the costs of individual learning and adaptation.

In the work described here, we cast the modeling problem in the context of boundedly (Simon, 1955), or ecologically (Smith, 2008), rational self-interested agents competing through the costly acquisition of valuable information. We use John Holland's learning classifier system, an evolutionary computation technique developed in computer science, to model agent learning. The structure and method of our model is similar to a conventional agent-based model except we use the learning classifier system to evolve agent rules, rather than supplying them ourselves. Unlike a conventional agent-based model the LCS methodology can generate for each agent separately a very large set of decision rules (we allow over a thousand in the three LCS together) as is appropriate for an agent adapting to a complex environment. An LCS can be thought of as a set of computational procedures that defines the way an agent can learn by experience. Alternatively, it can be thought of as a prescription for dynamic rational behavior in a complex and uncertain environment. The conceptual approach is an almost strict analogy with the mechanisms of Darwinian evolution but the evolution here concerns an agent's experience-based ideas about the outcome of its actions, i.e., the agent's mental model of its environment. That mental model consists of a list of condition:action:outcome rules and several measures of their contribution to the agent's well being. The agent treats each rule as tentative knowledge—the equivalent of a scientific hypothesis. At each iteration of the model, the agent selects those rules that are appropriate to its immediate environment, chooses an action that has performed well in the past, implements it and records the result. Over time it selects those rules that predict more accurately than the others. Because the agent's sources of information are always ambiguous, the accuracy of its interpretation of any observation or communication with another agent and the resulting value of the message, that is, its ability to understand the context and discern accurately the meaning in the message depends upon its experience. Consequently, the agent's history strongly biases its understanding of the information it obtains, the opportunities it perceives, its choice of actions, and from that, the path of its experience and learning.

We define three different kinds of opportunity costs affecting learning and adaptation: the costs of autonomous searching, of communicating, and of deciding. Individuals respond to those costs by carefully, i.e., economically, choosing to learn and interact with familiar agents in familiar arenas where they have relatively secure expectations about the outcome of their actions. Emerging from these choices are persistent individual relationships that lead to equally persistent social structure and the implicit coordination of aggregate production. The extent of the social structure and of coordination is a function of the various costs of information. In environments in which the costs of communications are high, there is a strong tendency for agents to form small groups with frequent individual interactions. These strong relationships imply close coordination within groups, but weak connections and coordination with individuals outside the groups. In other environments in which the costs of communication are low, agents imitate widely, developing relatively weak connections with a large number of agents. As the costs of communication fall toward zero, everyone communicates with everyone but at a low rate that is indistinguishable from random connections. Agents have little or no reason to maintain individual or group relationships and there is no basis for self organizing order in the system.

In conclusion, evolutionary theory seeks to explain order in living systems through an understanding of behavioral variation and selection. The broad sweep of the biological sciences, from genetics to ecology, is meant to refine that understanding. We suggest that an evolutionary theory of economic behavior and organization has to incorporate an understanding of the broad patterns of variation that arise from selective individual learning. Our intuition is that that understanding will better illuminate the circumstances in which self organization can be expected and broaden our sense of the dynamics of cooperation and competition. Advances in the field of artificial intelligence open the door to these kinds of investigations, as we try to illustrate with our model of a small, complex economy.

Acknowledgments

We would like to acknowledge the support of the National Science Foundation Program in Coupled Natural and Human Systems (Project 0909449) and Maine Sea Grant, which funded our initial work in this area. We received thoughtful and helpful comments from David Sloan Wilson, Carl Simon, Rick Riolo, Charles Lucas, Brian McGill, Tim Waring, Ted Ames, and many skeptical fishermen. James Wilson is the principal author and lead P.I. on the project. Yan wrote the original code for the simulations. J. Hill, C.L. Wilson, and L. Whitsel modified the code, ran the simulations and analyzed the results. The remaining listed authors are project members who contributed in a variety of important ways. And finally we would like to thank Barbara Harriy for her wonderful editorial help.

Appendix A. Supplementary appendices

Supplementary appendices associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jebo.2012.12.019>.

References

- Acheson, J.M., 1988. The Lobster Gangs of Maine. University Press of New England, Lebanon, NH.
- Acheson, J.M., 2003. Capturing the Commons: Devising Institutions to Manage the Maine Lobster Fishery. University Press of New England, Hanover, NH.
- Arthur, W.B., 1992. On Learning and Adaptation in the Economy, Working Paper. Santa Fe Institute, Santa Fe, NM.
- Boyd, R., Richerson, P.J., 2005. The Origin and Evolution of Cultures. Oxford University Press, New York.
- Cohen, M.D., Riolo, R.L., Axelrod, R., 2001. The role of social structure in the maintenance of cooperative regimes. *Rationality and Society* 13, 5–32.
- Cowan, D., Watson, W., Solow, A., Mountcastle, A., 2007. Thermal histories of brooding lobsters, *Homarus americanus*, in the Gulf of Maine. *Marine Biology* 150, 463–470.
- Darwin, C., 1859. On the Origin of Species by Means of Natural Selection, or The Preservation of Favoured Races in the Struggle for Life. John Murray, London.
- Epstein, J.M., 2006. Generative Social Science: Studies in Agent-Based Computational Modeling (Princeton Studies in Complexity). Princeton University Press, Princeton, NJ.
- Fosco, C., Mengel, F., 2011. Cooperation through imitation and exclusion in networks. *Journal of Economic Dynamics and Control* 35, 641–658.
- Frenken, K., Marengo, L., Valente, M., 1999. Interdependencies near-decomposability and adaptation. In: Brenner, T. (Ed.), *Computational Techniques for Modelling Learning in Economics*. Kluwer, Dordrecht, pp. 145–165.
- Gowdy, J.M., Dollimore, D., Witt, U., Wilson, D.S., 2013. Economic cosmology and the evolutionary challenge. *Journal of Economic Behavior and Organization* 90, s11–20.
- Hayek, F., 1945. The use of knowledge in society. *American Economic Review* 35, 519–530.
- Holland, J.H., 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press, Ann Arbor.
- Holland, J.H., 1996. Hidden Order: How Adaptation Builds Complexity. Helix Books, Reading, MA.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E., Thagard, P.R., 1989. Induction: Processes of Inference, Learning and Discovery. The MIT Press, Cambridge, MA.
- Iacoboni, M., 2012. Like attracts like. From edge.org, <http://edge.org/response-detail/2925/what-is-your-favorite-deep-elegant-or-beautiful-explanation> (accessed 13.07.12).
- Incze, L., Xue, H., Wolff, N., Xu, D., Wilson, C., Steneck, R., Wahle, R., Lawton, P., Pettigrew, N., Chen, Y., 2010. Connectivity of lobster (*Homarus americanus*) populations in the coastal gulf of Maine. Part ii. Coupled biophysical dynamics. *Fisheries Oceanography* 19, 1–20.
- Lanzi, P.L., Riolo, R.L., 2000. A roadmap to the last decade of learning classifier system research (from 1989 to 1999). *Lecture Notes in Computer Science* 1813, 33–61.
- Levin, S.A., 1998. Ecosystems and the biosphere as complex adaptive systems. *Ecosystems* 1, 431–436.
- O'Neill, R.V., DeAngelis, D.L., Waide, J.B., Allen, T.F.H., 1986. A Hierarchical Concept of Ecosystems. Princeton University Press, Princeton, NJ.
- Ostrom, E., 1990. Governing the Commons. Cambridge University Press, New York.
- Page, Scott., 2007. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools and Societies (New Edition). Princeton University Press, Princeton, NJ.
- Riquelme, J., Aguilar, J., Toro, M., 2000. Discovering hierarchical decision rules with evolutive algorithms in supervised learning. *IJCSS* 1, 73–84.
- Simon, H.A., 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics* 69, 99–118.
- Simon, H.A., 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* 106, 467–482.
- Simon, H.A., 2002. Near decomposability and the speed of evolution. *Industrial and Corporate Change* 11 (3), 587–599.
- Smith, C.H., 2012. Natural selection: a concept in need of some evolution? *Complexity* 17, 8–17.
- Smith, V.L., 2008. Rationality in Economics: Constructivist and Ecological Forms. Cambridge University Press, New York.
- Steneck, R.S., Wilson, C., 2001. Long-term and large-scale spatial and temporal patterns in demography and landings of the American lobster *Homarus americanus*, in Maine. *Marine and Freshwater Research* 52, 1303–1319.
- Urbanowicz, R.J., Moore, J.H., 2009. Learning classifier systems: a complete introduction, review, and roadmap. *Journal of Artificial Evolution and Applications* 2009, 1–25.
- Williamson, O.E., 1981. The economics of organization: the transaction cost approach. *American Journal of Sociology* 87, 548–577.
- Wilson, C. Thistle Marine Electronic Logbook Database. Maine Department of Marine Resources, Boothbay Harbor, ME, unpublished data.
- Wilson, D., Gowdy, J., 2013. Evolution as a general theoretical framework for economics and public policy. *Journal of Economic Behavior and Organization* 90, s3–10.
- Wilson, D.S., Wilson, E.O., 2007. Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology* 82, 327–348.
- Wilson, J., 1990. Fishing for knowledge. *Land Economics* 66, 12–29.
- Wilson, J., Yan, L., Wilson, C., 2007. The precursors of governance in the Maine lobster fishery. *Proceedings of the National Academy of Sciences of the United States of America* 104, 15212–15217.
- Wilson, S.W., 1994. ZCS: a zeroth-level classifier system. *Evolutionary Computation* 2, 1–18.
- Wilson, S.W., 1995. Classifier fitness based on accuracy. *Evolutionary Computation* 3, 149–175.
- Wright, S., 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the Sixth International Congress on Genetics* 1, 356–366.