4

Evolution and Market Complexity

John E. Mayfield

Abstract

Many complexities in our world come about through the use of preexisting purposeful information. This information may be structured in various ways (e.g., instructions, recipes, algorithms, rules, rules of thumb, business plans, and expert knowledge) and, if followed, directs the formation of something which otherwise would not have existed. This chapter argues that information organized in this way must ultimately arise as the output of an evolutionary computation. Because of this, an evolutionary process underlies most everything that characterizes human existence. This principle includes economics and markets. This chapter addresses whether or not understanding the fundamental role of evolutionary computation for enabling human and biological complexity provides useful insight into market behaviors and introduces the basic concepts necessary to have this discussion.

Introduction

Market complexity can mean several things. It can refer to market behaviors, to the structure of the market, or to strategies employed by market participants. The structure of any market is determined by rules that define it, and these rules can generally be seen to evolve over time. A market is also characterized by the various strategies employed by its traders; strategies also evolve. A market exists in the context of an external environment, and aspects of this environment may also evolve. Market behavior, however, is the outcome of interactions between the market rules, strategies employed, and the market environment. As such, market behaviors do not, strictly speaking, evolve. They may change, but they do not evolve. To explore these ideas further, we need to work with a clear definition of evolution.

The word "evolution" is used in various ways in scholarly discourse, and the value of evolutionary concepts for understanding markets depends on which meaning one chooses to use. I distinguish three different meanings. The

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

first is change that builds in a nearly continuous way upon the past. With this meaning, the current state of a system is derived directly from its previous state in some incremental way. This definition is very general and includes much of physics; concepts of force and potential are often quite useful for explaining how and why change occurs. When a cosmologist speaks of the evolution of the universe, this is the usual meaning intended. The second meaning is a subset of the first. It describes situations in which change comes about not because of a force, but by means of selection. Selection is a very common feature of our world, and selection and force can combine as agents of change. Systems which change as a result of selection are quite common. Examples of things formed through the action of selection include snowflakes, marble statues, and the branching of neurons in the developing brain. The third meaning is what Darwin introduced to the world. It also depends on selection, but here selection acts on stored information that is used to make something happen. It is the body of information that is updated during the selection process. This is the meaning used when we talk about the "evolution of life" or the "evolution of a technology." We are directly aware of changes in life over time because of fossils and of technological progress as well as preserved artifacts and written accounts. In both cases, however, it is important to recognize that objects do not evolve; it is the underlying information which creates the objects that evolves.

I use the third meaning when I say that market behaviors do not evolve. Within this meaning, market rules and market strategies evolve, aspects of the market environment may even evolve, but market behavior does not. Behavior is the result of interactions between market-defining rules, internal strategies, and external events. Market behavior characterizes a complex adaptive system, which changes all the time. What evolves is information that defines the structure and strategies.

Evolution in the Darwinian Sense

Darwin revolutionized scientific thinking by identifying a simple, but powerful mechanism for progressive change that accounts, in a general way, for many features of the observed biological world (Darwin 1859). Almost immediately, social scientists began adapting the core idea to explain various aspects of human society. For much of the intervening 150 years, most social scientists saw the relationship between biological evolution and social evolution as one of analogy. We now know that it is more than that. Focusing on the central role played by information makes the relationship clear (Mayfield 2013).

Darwin did not know about information in the modern sense of the word, just as he did not know about genes; however, he clearly saw that if evolution were to occur (changes in lineages were to occur over time), there must be variation of traits within a population of organisms and these traits must be passed down from generation to generation (i.e., they must be heritable).

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

Variation is essential so that selection is possible, and heritability is essential so that selections made are "remembered" and passed on. Lamarckism, as a concept in biology, foundered because no hereditary mechanism has ever been found for the remembering of traits acquired by use.

We now know, of course, that the principal way that living organisms pass on "remembered" information is in the structure of molecules in their DNA cells. DNA encodes information in the order of its chemical subunits, called nucleotides. The four nucleotides that make up DNA function much as the zeros and ones in binary computer languages, or the 40-odd letters and symbols in European languages. Their order encodes purposeful information—information to be used to make something specific come into being. In biology, that "something" is the building and maintenance of an organism. The details are quite complicated, but the big picture is pretty simple.

The essential role in life of encoded information has been clear since the structure of DNA and its implications were deduced by Watson and Crick (1953). Understanding the structure also revealed how DNA encoded information could readily change. Any manipulation of information can be viewed as a computation, and given that DNA encodes information, it follows that changes in DNA over time can be treated as a computation. Figure 4.1 illustrates evolution as a computation, without the use of biology-specific concepts such as organisms, cells, genes, and traits.

In the figure, inputs and outputs are bodies of encoded information; outputs once selected become inputs to the next cycle. In biology, the inputs and outputs are DNA sequences. Outcomes can be most anything that results from the information present in outputs. Outcomes play a key role in selection, but do not become future inputs. In biology, the outcomes are organisms and their

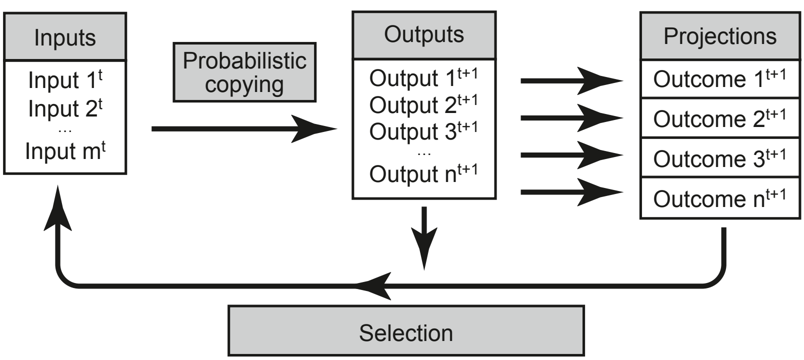

Figure 4.1 Diagram of the computational strategy that underlies all evolutionary processes. Superscripts t and t + 1 indicate cycle number; m^t represents the number of inputs for a particular cycle, which must on average be less than or equal to n^{t+1}, the number of outputs in the previous cycle indicating that selection must occur. Both inputs and outputs encode information (from Mayfield 2013; reprinted with permission from the Columbia University Press).

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

behaviors. Probabilistic copying may be viewed more generally as "copying with change." The strategy works with a wide variety of change mechanisms, so long as some of the outcomes are no worse than the inputs (according to the criteria used for selection). A virtue of the computational definition of evolution is that it applies equally well to systems outside of biology. In fact, biological evolution can be seen as a special case of a more general computational strategy. Other examples of this computational strategy in action include genetic algorithms (computer programs), the refinement of antibody gene sequences in the human immune system, technological change, advances in science, and a large number of other phenomenon observed and studied in the social sciences. A case can be made that the human brain also employs the same strategy (Mayfield 2013).

Most complicated things in human society are based on preorganized, purposeful information. Most often this information is in the form of "instructions" (or recipes or algorithms). Instructions are purposeful because their only function is the creation of something specific. A fundamental philosophical question is how instructions (or purposeful information generally) are formed.

In computational theory, every message, instruction, or recording of data is the output of a computation, and if we embrace the Kolmogorov (algorithmic) notion of information, then the information content of that message or instruction must have been present in the input to the computation that produced it. If we allow random input to be possible, the information content of any body of information then has two possible origins: prior information or random input to the computation that output it. For every existing body of organized information we can envision a long chain of prior computations, where the input of one is the output of a previous computation. For each computation in the chain, there may or may not have been some random input (i.e., random bits introduced). Following such a chain of computations back in time, all information we currently deal with ultimately originated either with input to the first computation or with randomness. But what was the first computation and when did this occur? The obvious situation is for the first input to have originated with the origin of the universe, and most current ideas of the "big bang" favor an origin with no initial information. If this is so, then we are left to conclude that all current information originated from randomness. This view of the origin of purposeful information, however, creates a conundrum; even though the set of all possible randomly generated bit strings contains every possible sequence (and therefore must include instructions for producing anything one might ever want or need), a systematic search for an instruction with a particular desired outcome is completely hopeless. This is because the number of strings that must be searched is 2^n, where n is the length of the desired string. Thus, to search all 1000 bit strings (looking for a useful instruction of that length) would take longer than the lifetime of the universe, and many instructions are much longer than 1000 bits.

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

Since we employ lots of very useful bit strings (or their equivalent in human or computer languages), there must be a way around this "probability problem," and indeed there is. It is the computational strategy shown in Figure 4.1. To emphasize the centrality of this strategy to our very existence, I have called it the "engine of complexity" (Mayfield 2013). Some might be more comfortable calling it the evolutionary algorithm. The power of this computational strategy for finding a bit string with a desired global property is illustrated by comparing the time required to solve a simple computer science problem called "one max." The problem is to convert an arbitrary binary string into a string of all ones without telling the computer what a "one" is. The goal of the computer programmer is to get the computer to generate a string of length n that consists of all ones when the only information available to the computer is the arithmetic sum of the bits that comprise the string (a global property). In Table 4.1, the number of strings (equivalent to time) required to produce a string of all ones, using a simple evolutionary algorithm, is compared with the number of strings required to solve the problem by generating random strings of length n. For each n, the algorithm maintained a population of 10 strings and mutated the strings by flipping bits in the string with a probability of $1/n$.

As one can see, the relative power of the evolutionary approach increases with the length of the string, and the computation easily generates a string of 1000 ones, even though this is totally impossible by random string generation. Impossible because the generation of 10^{300} strings, 1000 bits long, would take all the computers on Earth vastly longer than the projected future lifetime of the universe. The reason this strategy works is that whenever information is useful for something, the probability of that particular configuration of bits is $1/n^2$. The probability that changing one bit in that string will lead to a better (or worse) outcome is of the order of 1/2, and sequential steps with a probability of 1/2 each have reasonable probability whereas steps with a probability of $1/n^2$ do not. Selection of a sequence of steps each with reasonable probability is relatively easy in comparison to a single step that will never happen.

Table 4.1 Comparison of the time (number of strings that must be generated) to create a string of all ones by genetic algorithm or by random means (Mayfield 2013).

Length of the string (n)	Number of strings needed by genetic algorithm (on average)	Number of strings needed to generate the target randomly (on average)	Speed-up by the evolutionary algorithm
10	60	500	8
50	600	10^{15}	1,666,666,666,667 (1.7×10^{12})
100	1,330	10^{30}	10^{27}
1,000	34,000	10^{300}	3×10^{297}

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

From this argument we conclude that if a string of some sequence of ones and zeros has some special significance to you (such as being a useful instruction), then the origin of the sequence that encodes this meaning must have been repeated small random changes selected incrementally many times. Otherwise there is no way that a large body of purposeful information could ever have come into existence. Saying "someone thought it up" is not an acceptable scientific alternative, nor is assembling it from smaller bodies of existing information because the assembly of each of these smaller bodies of information also had to overcome the same probability problem.

Selection plays a key role in the evolutionary process, but the power of the strategy derives from more than just incremental selection. Two other features are very important for understanding how this computational strategy can achieve what seem to be hopelessly improbable results so quickly. First, the mutation rate is critical: if changes are either too great or too small, the power is lost (see Figure 4.2). Obviously, with a mutation rate of zero, nothing will happen. Less obvious is that a mutation rate that is too big can also result in lack of progress. Not explicit in Figure 4.2 is the observation that if an existing body of previously organized information is subjected to an evolutionary process with a mutation rate that is too large, the body of information can be quickly reduced to random gibberish. In addition, an evolutionary strategy works best when there are multiple inputs. The process shown in Figure 4.1 will continue on with only a single input during each cycle, but it proceeds toward a desired goal much faster when there are multiple inputs.

The mathematically informed will be quick to point out that in the general case there is no guarantee that an evolutionary strategy will always work. It is easy to create hypothetical environments in which long-term improvements

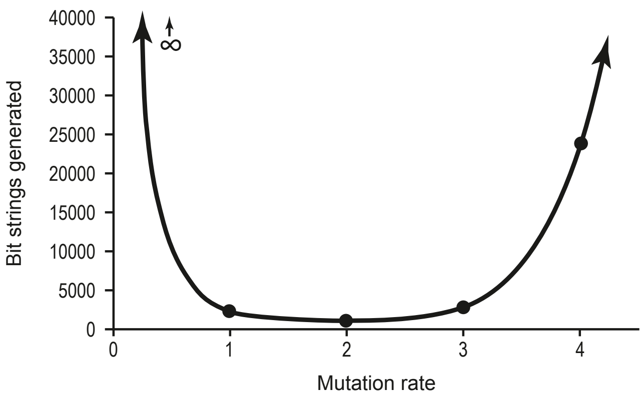

Figure 4.2 Number of strings that need to be generated by the one-max-solving genetic algorithm to find a string of 100 ones plotted against the average number of changes made to strings in each generation (mutation rate). At a mutation rate of zero, the algorithm takes forever; at a 100% mutation rate, roughly 10^{30} strings are required (from Mayfield 2013; reprinted with permission from the Columbia University Press).

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

based on accumulated small changes are not possible. Fortunately, such situations seem to be rare in the real world.

Social Evolution

Does the process diagrammed in Figure 4.1 properly describe social evolution? Two points relevant to deciding the answer to this question are: (a) everything uniquely human in human culture is based on language and (b) language allows the efficient transmission and recording of purposeful information. Everything that we describe as cultural is made possible by recorded (i.e., remembered) information. Thus, Mozart concertos, grocery stores, presidential elections, the Roman Catholic Church, and the theory of general relativity all depend on recorded language. Sophisticated language use allows us to share complex ideas and knowledge with others. From a computational perspective, communication consists of transmitted and received information, and whenever a chunk of information is communicated and remembered widely among members of a society, we recognize it as being part of the culture. According to this viewpoint, a culture consists of many remembered communications and the resultant behaviors this remembered information elicits.

Over the course of human existence, culture has been utterly transformed at least twice: by the biological evolution of language ability and by the invention of writing. Both transformations provided for expanded remembering and communication and were accompanied by vast increases in cultural complexity. Spoken language makes possible villages, traditions, crafts, and trade. The invention of writing greatly increased the amount of information that could remain current in a culture and also increased the number of people with access to it. Books and libraries allowed storage of information that no single person could remember in detail. We may be in the midst of a third transformation brought about by computer-encoded information and the Internet—time will tell.

Culture is not just about ephemeral things like fashions, pop music, democratic institutions, and religious practices, but also about physical artifacts, things like kitchenware, computers, and automobiles. All of these are made possible by the communication of purposeful information: instructions, recipes, laws, rules, and rules of thumb.

Accepting the informational basis of culture and cultural artifacts is a necessary starting place for any evolutionary theory of cultural change, but the process diagramed in Figure 4.1 requires more. Three additional things are needed: copying with modification, selection, and the maintenance of multiple versions of the information that underlies various aspects of a culture (this amplifies the power of selection).

The first question that comes to mind when considering social evolution is where is the "DNA" that encodes the accumulated cultural information? Clearly, a lot of it is stored away in books and other written materials, but

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

the working information that characterizes culture is stored in people's heads. Learned information lies at the heart of everything cultural. Written information plays a huge role in modern culture, but it must always be read and learned by someone. Cultures are characterized by shared information, and the key activities of individuals participating in cultures involvtransmission of previously learned information and selective remembering. But, culture is more than communicated information; it is information that has acquired a degree of stability through wide acceptance. For information to contribute to culture, there must be closed communication loops. When a concept or technique is communicated from one to another, and then passed on to yet other individuals, it eventually finds its way back to the original communicator in modified form. Only when this return communication is similar to the original message is a sufficient degree of stability achieved for the information to have more than a fleeting impact on a culture.

Some authors see a big difference between biological evolution and social evolution because social evolution seems to have different, or at least additional, change mechanisms. In particular, changes in DNA are not made with foresight, whereas changes in information structure that underlie social evolution frequently are made with foreknowledge of the impact that a change will make. Other apparent differences are that social evolution can make huge leaps (lots of new information can be introduced at one time), and that continuous viability is not required in a lineage (Page 2010). To give an example: in technological evolution, intermediate prototypes need not be commercially viable. For living organisms, however, every link in a lineage must be viable or the lineage terminates.

These differences, though important for understanding social change, do not invalidate the basic role of the engine of complexity in social evolution. Many of the apparent differences between social and biological evolution are because human mental activity is involved at every step in the process. Others are because social evolution occurs on multiple levels. I will argue later in this chapter that human thought and creativity also rely on the engine of complexity. If this is accepted, then the foresight and creativity that seems to make social change different is the outcome of evolutionary processes occurring in the brain. Thus, social evolution can be seen to occur on two levels: in people's heads and in society. These levels inform one another, but the underlying computational strategy is the same in both. As for the lack of a need for continuous viability in social evolutionary lineages, understanding of this flows directly from the differences in how information is encoded. In biology, information is encoded in DNA, which until recently could only be maintained in a living organism. Social information has no such requirement; it can be maintained in a library, or in a computer, or in a human brain and later retrieved without any need for interim viability.

So, by considering again the process diagrammed in Figure 4.1, we can see how this process describes the transmission of cultural information by

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

relabeling. Inputs are the concepts, ideas, plans, etc., that underlie a human behavior or explain how to do or make something. Probabilistic copying becomes the communication of information from one person to another with change. It may be probabilistic because human memory and communication is less than perfect, but changes can also include the products of human thought (which itself may be the product of a similar process operating in the human brain). The outputs in the diagram represent information learned by another, and projections are consequences of the learned information. The process is clearly a parallel one (many inputs) because many people may be communicating similar concepts, and clearly there is selection because people often do not remember what they have heard or choose not to pass it on. So, from this brief argument, it can be concluded that human cultures exhibit and employ exactly the same underlying computational strategy that characterizes biological evolution (the engine of complexity) with differences in the details of how information is encoded, how changes are introduced, and how selection operates.

Relevance of Evolutionary Thought to Economics

It is generally agreed that economies result from agents making and exchanging things with one another. In simple economies, the agents are people; in more complex economies, agents can also include organizations of people such as businesses and governments. Money and credit allow exchanges to occur on larger scales and in more complicated ways than would otherwise be possible. But always, from the simplest exchange to the most complex, both sides of an exchange must perceive that they gain something by making a trade.

Economies were once based on what people could individually raise or make, but this is no longer the case. Economies are now based almost entirely on what businesses or other human institutions create. Business firms play such a dominant role because they can produce large quantities of things that people want at a low cost. They are able to do this by taking advantage of the efficiencies inherent in various forms of organized cooperative activity (Chandler 1977). The opportunities are legion, but always a firm must produce things that people (or other firms) want and are willing to buy at a price that is greater than the cost of production. To accomplish this, every firm must have detailed organizational and production plans. Taken together we will call this their business plan (Beinhocker 2006). This plan includes everything from instructions for assembling whatever the firm makes, to the organization of its employees, to a plan for marketing and distributing. A business plan encodes organized purposeful information and, as I argue here and elsewhere, owes its assembly to an evolutionary process. The logic behind this claim is the necessity to overcome the probability problem. Large bodies of organized purposeful information can never be created *de novo*: they are always created through a process of selective tinkering, borrowing, and accumulation. In an

From "Complexity and Evolution: Toward a New Synthesis for Economics," David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19, series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

ever-changing business world, a healthy business is constantly tinkering with its business plan. To take full advantage of the engine of complexity, a firm needs to pay close attention to the rate of innovation and to the advantages of making tinkering with the business plan a parallel process.

Modern economies are only possible in a context of rules and the only effective creator and enforcer of such rules is government. Without effective rules, everyone would attempt to cheat everyone, and the only protection anyone could rely on would be physical force. In such an environment, modern firms would not be possible and neither would the economy as we know it. Rules and laws are also characterized by organized purposeful information no less than business plans, and their change is also an evolutionary process.

Markets play a central role in economies and have rules that allow them to function and give them their specific characteristics. Also, traders have strategies. Rules and strategies consist of organized purposeful information no less than laws and business plans. So, from this very brief discussion, it is pretty clear that modern economies are enabled, defined, and characterized by purposeful information in the form of laws, rules, plans, and strategies. I claim that all this information was assembled by evolutionary processes. To make this claim credible, we must explore human creativity and implications of the probability problem.

Evolution and Human Thought

The computational view of evolution can also be considered a theory of creativity, and no theory of creativity can ignore what goes on in the human brain. As discussed above, from the computational perspective, randomness can be seen as the ultimate source of everything new. But, it is not a very convenient source because of the probability problem (large bodies of information cannot be generated by random means because they are too improbable). The evolutionary process diagrammed in Figure 4.1 (the engine of complexity) provides a practical way for this unruly source to be corralled, or channeled for the production of useful structure. Reflection on the notion of creativity leads one to the conclusion that creativity must have its roots in randomness. Without the introduction of randomness, there is only determinism, and in a deterministic universe, creativity is an illusion.

When learning something new, the human brain must overcome the same probability problem faced by any activity that generates organized purposeful information. The brain encodes information in the form of neural connectivity patterns. When we learn something new, appropriate connectivity patterns must be created. But how the brain does this is not obvious. The new pattern must not only represent the new thing learned, it must also interface correctly with myriads of other already established connectivity patterns. This seems like a tall order. Given that the total number of potential brain connectivity

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

patterns is vastly large, how does the brain find one that satisfies all the requirements for the learning task of the moment, especially when that task has never before been encountered? This is a question that neuroscientists would dearly love to answer. For this procedure to be deterministic requires the brain to be hard wired to respond with just the proper circuitry whenever anything new is encountered—even things that the human race has never before encountered in its entire evolutionary history. Anyone proposing such a theory must also deal with the substantial problem that everyone's brain is wired differently at the level of individual neurons and their synapses.

A very promising approach that is easily adapted to admit random input has been popularly called the "Bayesian brain" (e.g., Lee and Mumford 2003; Knill and Pouget 2004; Friston 2010). In this model of brain function, what we perceive as our world is actually an internal world model. So when we look out the window and see a tree, we are not directly accessing the visual data stream from our eyes, but are drawing on an internal model or representation of the tree that has been created in our brain. The visual data stream does not paint the picture, but is used to continuously update the internal model. The brain modifies the neural connectivity patterns that portray the tree by minimizing the mismatch between the visual input and predictions of what that input should be based on the internal model. In the simplest situations, this pattern modification is simply an optimization of the synapse strengths in an unchanging architecture, but what does the brain do when something new enters the visual field? If it is a familiar object, such as a common bird, then it is easy to imagine the brain quickly accessing previously learned patterns which can then be updated and integrated into the tree model. But what if something enters your visual field that you have never seen before? What is the brain to do if it has no remembered neural patterns that even approximately match the input coming from your sense organs? The number of neural patterns it might try out is overwhelmingly large, and initial neural pattern attempts will likely be pretty bad. A rather simple solution to this seemingly hopeless probability problem is for the brain to start by producing a few wild guesses and then evolve connectivity patterns that provide increasingly better matches. By using the word "evolve" here I am hypothesizing that the brain employs the general information processing strategy diagrammed in Figure 4.1.

If the brain does indeed evolve appropriate and effective neural patterns when faced with a novel learning situation, to be efficient, it must produce a population of trial connectivity patterns, select the best, modify these to produce new populations, select from these, and repeat the cycle long enough to achieve a good match to the task at hand. Central to this is selection, but if it is to work efficiently, it needs to employ more than simple selection. It must maintain a "population" of mental models and rapidly update these models with random or preprogrammed modest changes, all while assessing the mismatch between the predicted and actual input. Internal models with the largest mismatch would be discarded while those that match best would

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

serve as the basis for a new population of models all differing somewhat from one another.

It is easy to imagine this same strategy being engaged when we "think up" something new. Creative people are constantly generating novel concepts, novel ways of looking at the world, or novel solutions to problems. All are based on complex neural connectivity patterns in the brains of the creators. There is no doubt that all new ideas incorporate bits and pieces of old ideas. But how are these bits and pieces put together into something meaningful? How were the bits and pieces originally created? Reflecting on childhood development, how does a baby, with no previous concepts to draw upon, create internal models that make sense and work? If we evolve neural connectivity patterns by maintaining populations of patterns all somewhat different, and we are continually making small changes to existing patterns, and patterns are selected by minimizing the mismatch between predictions of each pattern and current sensory data and/or expectations, then it is not hard to imagine how our brain could generate new patterns that satisfy complex criteria.

I am particularly intrigued by this theory of brain function and creativity, because if the brain does indeed work this way, and I must emphasize that this has not been firmly established, then everything in the realms of biology and human activity ultimately owes its genesis to the engine of complexity operating at three levels: DNA, human cultural information, and human brain encoding. One could say that accumulated information defines everything that characterizes our existence and that a single mechanism, a computational strategy, makes it all possible.

Evolution and the Generation of Complexity

A fundamental feature of most complexities we encounter in our world is that they are specified by, or are a consequence of, previously specified information. Thus, most of the complexities that characterize life owe their existence to information stored in DNA, and most technological wonders owe their existence to information encoded in detailed instructions and blueprints. A similar case can be made for everything within the realm of social science, which relies on recorded or remembered information (Cziko 1995; Wilson 2002; Mayfield 2013). Complex thoughts can also be characterized in informational terms, and all of our actions result from neuronal connections that encode purposeful information.

Purposeful information, by definition, has a purpose. It only exists because it leads to something else. Whether such information is encoded in DNA, the English language, or neural connectivity patterns, it is always created in advance of its use. Information and complexity are related concepts but are not the same thing. There is no single quantitative measure of complexity, but various aspects of it can be defined and measured. A complex

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

object is generally thought of as having lots of different kinds of parts that interact in non-simple ways. Complex behaviors exhibit unexpected but non-random actions.

Two commonly used measures of information are also applied to the measurement of complexity: entropy and Kolmogorov complexity (algorithmic information). Another measure that I find particularly intriguing is depth (logical depth). Entropy is a measure of the distribution of probabilities of all potential states of a system. It measures uncertainty (Cover and Thomas 1991) and increases with the size of the system. For example, if the system consists of all binary strings of length n and the probabilities of the strings are equal, then the entropy (in bits) = n. Entropy is maximal when the probabilities of all possible states are the same, and unequal probabilities result in entropy less than maximal. When the entropy of a system is less than maximal it is sometimes useful to call the difference between the actual and the maximal entropy the negentropy. Negentropy is maximal when the probability of one state is one (and all other states have probability zero), and in this case, the negentropy equals the negative of the maximal entropy. Negentropy can be thought of as information stored in the system.

Kolmogorov complexity measures individual objects, and when applied to individual binary strings is defined as the length of the shortest input to a universal computer that will output that string (Li and Vitanyi 1997). Kolmogorov complexity of an object is maximal when the elements of that object are randomly arranged. So, Kolmogorov complexity measures the amount of randomness present. As one considers objects of different sizes, their Kolmogorov complexity may or may not reflect their actual size as some large objects can be computed from short input. An important property of Kolmogorov complexity is that it cannot be created in a deterministic computation. Thus, in a chain of deterministic computations, the Kolmogorov complexity of the final output may be less, but cannot be greater than the Kolmogorov complexity of the initial input.

Depth (originally defined as logical depth) measures the difficulty of computing an object (Bennett 1988). It is informally defined as the number of elementary steps required to compute an object from short input. Depth has a "slow growth" property, which means that it requires extensive computation to compute a deep object from shallow input. Because deep objects can be computed rapidly from appropriate deep input, depth can be thought of as stored computational effort. Deep objects are complicated in many ways.

The complexity of behavior is most easily studied by recording the behavior. The recording is then treated as an object that can be characterized in standard ways.

An important principle is that evolutionary computation (Figure 4.1) when provided access to random input is capable of generating increasing complexity measured by these three measures, or any other you may prefer. All that is required is for that particular aspect of complexity to be favored during the

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

selection process. The only serious caveat is that the information space in which the evolutionary process is working needs to be "correlated" (Kauffman 1993), which simply means that the space of possibilities in which evolution is occurring is structured so that some nearby structures are better at whatever is being selected than any in the current population. It is easy to create test systems that do not meet this requirement (i.e., are uncorrelated), but they seem to be rare in the real world.

Processes that produce outcomes from outputs (the rightmost arrows in Figure 4.1) have the potential to add additional complexity. Whenever those processes can be properly treated as computations, the nature of any additional complexity generated is restricted by the laws of information science. Thus, depth can only be increased slowly and Kolmogorov complexity can only be increased if the process accepts external input.

The Behaviors of Economic Systems

Returning to economics, it is reasonable to ask: Does accumulated purposeful information play a central role in economic systems? And, if so, what is the nature of that information and from where does it originate? Clearly, economies are systems that are characterized by organizations and individuals interacting with one another. Systems characterized by interacting adaptive agents (often called nodes) with memory, which act according to rules, are known as complex adaptive systems (CASs). In such systems, agents make decisions based on internal rules, past actions, and new input (Holland 2006; Miller and Page 2007; Page 2010). CASs include the most complex systems known and correspondingly exhibit an incredibly wide variety of behaviors. Economies are CASs. Some other examples include living organisms, ant colonies, ecosystems, corporations, and human communities. In all examples, the behavior and response of each node to actions of other nodes is rule based, and those rules need not be deterministic. In many CASs, agent rules evolve (in the meaning diagrammed in Figure 4.1). The nature of such systems changes fundamentally over time as the individual agent rules change. When agent rules evolve, the system is able to adapt over time in response to changing external conditions, even if those conditions have never before been encountered by the system.

When thinking about and analyzing an evolving CAS, it is important to understand the level at which selection is working in the system. Selection can act at the level of the entire system or at the level of individual agents, or simultaneously at both levels. To help keep things straight, Wilson (this volume) introduces two distinct meanings: CAS1 refers to a complex system that is adaptive as a whole; CAS2 refers to a complex system in which individual agents employ adaptive strategies. In CAS1, adaptations of individual agents are based on the performance of the whole system, whereas in CAS2, adaptations of individual agents to their local needs drive the behavior of the system.

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

In CAS2 it is easy for the "selfish" actions of individual agents to drive the system to states that are not optimal for the whole.

Behaviors of CASs vary widely in their details, but are often characterized by quasi stability when confronted with changing external events. When the external perturbations are too large, they either collapse entirely or transition rapidly to a new regime of quasi stability. Conditions that lead to major transitions can often be determined by modeling. Predicting the long-term future of CAS with evolving rules is problematical because evolutionary output is generally unpredictable; however, if selection is acting on the system as a whole (CAS1), then broad predictions can often be made. In particular, if stability is favored by selection, we can expect selected rule modifications to favor increasing stability. Surviving systems demonstrate by their very existence that they are resistant to the unexpected change of individual agents.

The characteristics and behavior of any CAS depend on the rules. Rules determine the nature of the agents, the nature of the interactions between agents, and responses of agents to those interactions. Rules are characterized by organized purposeful information and, as I argue, all such information ultimately originates in an evolutionary process. The nature of human agents is determined in large part by their DNA which evolves far too slowly to influence changes in economic systems. The nature of human-made institutions, however, can change much more quickly, and the thoughts and ideas in peoples' minds can change in seconds. Likewise, the rules of commercial interactions between people and between their institutions change significantly within the lifetime of individuals. It would be hard enough to understand the behaviors of a complex economy if the rules were static, but it is doubly difficult when the rules are constantly evolving.

Behaviors of Markets

Markets can be treated as CASs in which the rules that define the operation of the market evolve, agent strategies evolve, and many aspects of the external environment evolve. The evolution of market rules clearly operates at the level of CAS1, whereas the evolution of agent strategies operates at the level of CAS2. The evolution of the environment is usually not under control of the market, but may be under control of a higher level system (i.e., government). In a system, when selection occurs at the level of agents as well as at the level of the system as a whole, it is not uncommon for conflict to occur. The levels may work against one another. An example would be the creation of market rules that prohibit certain trading strategies.

When studying markets, questions naturally arise about how changing markets and/or trader rules will affect the behavior of the market. Several observations seem relevant here. First, as with ecosystems, there are many examples of traders who create novel strategies that cause them to go broke (go extinct). Second, unlike in nature, trader rules are, presumably, not usually modified

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

randomly. When things are not going well, most traders will modify their trading strategies in ways they think will help. Third, it is possible to model markets and systematically try out proposed new market rules and/or trading strategies without risking the actual market or a trader's economic survival. Because of the difficulty in predicting the future behavior of markets when new rules are introduced, modeling is an essential tool.

It is much easier to discern the rules at play in a financial or commodity market than in a biological ecosystem; accordingly, agent-based models of markets should have success predicting market behaviors and potential trader strategies (Farmer and Foley 2009; Bookstaber 2012). The success of this approach in ecology is well established (Grimm and Railsback 2005), but achievements have been limited. With both markets and ecosystems, the key to better models is the better identification of the rules in play.

Market complexity can refer to structure (organizing rules and number and complexity of agents) or to market (i.e., price) behavior. Most often it is behavior that people are attempting to analyze and understand. The motivation, of course, is that if we understood market behavior well enough, we would be able to predict future behaviors (and make lots of money, or take actions to head off disasters).

Because the behavior of each market is the output of a CAS, we should expect behavior that is not obvious. Complicating things further is the continuous evolution of the defining rules. This means that market rules, communication rules, and agent strategies change over time and that these changes will be unpredictable. As the rules change, markets will inevitably behave in new ways, and the new behaviors may not be obvious in advance. Of greatest concern are changes that destabilize markets.

As a general rule, it is difficult to model CAS analytically (equation based), and often such models require unrealistic assumptions (Farmer and Foley 2009). In such situations, an appealing strategy is to apply agent-based computer models (see Epstein and Chelen, this volume). Such models can closely mimic the market being modeled and can be used to test for unexpected market behaviors and try out the consequences of modified rules (Lebaron 2006).

Probably the biggest problem with all economic models is that market traders are people, or computers working for people. Thus, human irrationality is inevitable. Markets are also immersed in economies that generate events and respond to world events. External events to the markets, sometimes predictable and sometimes not, occur and affect market prices. If they can be foreseen, they too can be factored into agent-based models.

Final Thoughts

What value is there in understanding and appreciating the role of purposeful information in the creation of complexity in our world? Even if you accept

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.

that evolutionary computation underlies most things in our world, by allowing the creation of large bodies of purposeful information which, in turn, enables nearly all of the complexities we encounter, does this knowledge allow us to do anything useful? I provide three brief answers.

First, a clear understanding of how any system works must be a prerequisite for meaningful scientific study. Second, evolutionary systems can rarely, if ever, be modeled analytically. When the outputs of evolution are the inputs to a complex adaptive system, the only realistic hope is to set up agent-based models that test observed or hypothesized rule changes in advance. Third, evolutionary systems are inherently creative. As such we should expect the unexpected as these systems evolve. When systems, such as markets, rely on the outputs and outcomes of evolutionary processes, there is always the possibility for new structures and new behaviors to arise unexpectedly. Broad trends can sometimes be predicted, but the details are rarely predictable.

In summary, I do not believe that seeing and analyzing markets as evolutionary complex adaptive systems provides immediately superior analytical approaches. What it does do, however, is to provide a solid base on which to build future study.

From "Complexity and Evolution: Toward a New Synthesis for Economics,"
David S. Wilson and Alan Kirman, eds. 2016. Strüngmann Forum Reports, vol. 19,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03538-5.