

Cognition: The Black Box of Economics

W. Brian Arthur¹

Chapter 3 in *The Complexity Vision and the Teaching of Economics*,
David Colander, ed., Edward Elgar Publishing, Northampton, Mass, 2000.

In his autobiography Bertrand Russell tells us he dropped his interest in economics after half a year's study because he thought it was too simple. Max Planck dropped his involvement with economics because he thought it was too difficult. I went into economics because I'd been trained in mathematics and I thought, as Russell did, that economics looked easy. It took me several years to get from Russell's position to Planck's. Economics is inherently difficult.

In this paper I will explain one path by which I came to that view.

Whether one sees economics as inherently difficult or as simple depends on how one formulates economic problems. If one sets up a problem and assumes rationality of decision making, a well-defined solution normally follows. Economics here is simple: From the Problem follows the Solution. But how agents get from Problem to Solution is a black box; and whether indeed agents can arrive at the Solution cannot be guaranteed unless we look into this box. If we open this box economics suddenly becomes difficult.

Once in a while as economists, we do justify our assumed connection between problem and solution. In a well-known paper, John Rust (1987) tells the story of Harold Zurcher, the superintendent of maintenance at the Madison (Wisconsin) Metropolitan Bus Company. For 20 years Zurcher scheduled bus engine replacement of a large fleet of buses—a complicated problem that required him to balance two conflicting objectives: minimizing maintenance costs versus minimizing unexpected engine failures. Rust figured out the solution to this combinatorial optimization problem by stochastic dynamic programming, and matched that optimization against Zurcher's. He found a reasonably close fit. The point of Rust's article was that although this was an

enormously complicated problem, Harold Zurcher found the solution, and therefore at least in this case economists' assumption that individuals find optimal solutions to complex questions is not a bad assumption.

The Zurcher example leaves us with a broad question: Can the assumption that individuals find optimal solutions to economic problems be justified so that we can avoid studying the details of the decision process? In simple cases the answer is yes. In most cases however, it is no. Think of an ocean that contains all the well-defined problems that interest us in the economy, with ever more difficult problems at greater depths. Near the surface lie problems like tic-tac-toe. Below that are problems at the level of checkers, and deeper still are problems like chess and Go. We might know theoretically that a solution to Chess exists, say in mixed Nash strategy form, but we can't guarantee that human agents would arrive at it. So the problems that are solvable the way tic-tac-toe is solvable lie within two or three inches of the surface, and an ocean of problems deeper than these cannot be guaranteed of solution. We can add to these the many problems agents face, perhaps the majority they face, that are *not* well specified. Zurcher's problem lies on the boundary of what economics agents can accomplish by way of a "rational" solution. Deeper than this, economic "solutions" may not match reality or may not exist.

What happens at these deeper levels? Human decision makers do not back off from a problem because it is difficult or unspecified. We might say that when problems are too complicated to afford solutions or when they are not well-specified, agents face not a problem but a situation. They must deal with that situation; they must frame the problem, and that framing in many ways is the most important part of

the decision process. To consider that framing you have to consider what lies between the problem and the action taken. And between the problem and the action lies cognition. Between the problem and the solution there's a lot going on, and if one considers what is going on, economics becomes difficult. To paraphrase my question then: How do people make sense of a problem? How do individuals handle these more complicated problems? How do we really cognize?

In this talk I want to consider cognition as a cognitive psychologist might look at it, and apply the findings to thinking about two different issues: economic modeling and the education of graduate students.

Notions of the Mind

In economics we have a simple and old notion of mind. Mind is a container that holds data. The data are constantly updated by interaction with the world; and mind performs deductions based upon these data. All of this of course is implicit; in economics we don't talk about "mind." But we do view mind—or that which gives rise to ratiocination—as deduction upon collections of data sets. In economic theory this is reflected in treating beliefs about the world as expectations of variables conditioned upon current data (or sigma fields)—current information—and in formulating solutions based upon these. This is a shorthand, the sort of reasonable abstraction that any science makes that works well in many cases. But we need to get beyond it when we go deeper than two or three inches into the ocean of problems.

Let me look at mind and the cognitive process then from a deeper viewpoint—that of cognitive science. Imagine that at night you are reading a novel, say Haldór Laxness's *Independent People* and you're enjoying it. What is actually going on? Actually, that's complicated. The black and white marks on the page are focused onto the light sensors or pixels at the back of your retina. These sensory perceptions are transmitted to the rear part of your brain, and map into certain visual structures there. Somehow letters and words are parsed out, and somehow these fit together via an understanding of syntax. (Where I say "somehow" I mean that cognitive scientists do not know the exact mechanism of what is happening.) From syntax somehow "meaning" emerges. But what is meaning? Meaning in this case is a set of associations. You might read a sentence about rain: "Smoothly, smoothly it fell, over the whole shire, over the fallen marsh grass, over the troubled lake, the iron-gray gravel flats, the somber mountain above the croft, smudging out every prospect." These words trigger associa-

tions—associated memories really—and you form a picture, or a set of pictures. These associated memories and pictures in turn trigger what you might call "affect," or feelings. The feelings are often subtle, the kind of feelings of what it might be like to be in Laxness's world—the gloom of the rain, the dreariness of the gravel flats, the oppressiveness of the mountain, the smell of the croft in the dampness. These are subtle feelings, and these feelings actually are our intelligence, are part of our cognition. They're part of the meaning that we give to symbols. Reading and making sense of what is read consist of associated memories and associated feelings. How all this happens is not well understood by cognitive scientists; it's what French thinker Henri-Jean Martin call a mysterious alchemy.

Here's how the Princeton cognitive psychologist Julian Jaynes (1976) expresses this alchemy of mind: ²

"O, what a world of unseen visions and heard silences, this insubstantial country of the mind! What ineffable essences, these touchless rememberings and unshowable reveries! And the privacy of it all! A secret theater of speechless monologue and prevenient counsel, an invisible mansion of all moods, musings, and mysteries, an infinite resort of disappointments and discoveries. A whole kingdom where each of us reigns reclusively alone, questioning what we will, commanding what we can. A hidden hermitage where we may study out the troubled book of what we have done and yet may do. An introcosm that is more myself than anything I can find in a mirror. This consciousness that is myself of selves, that is everything, yet nothing at all – what is it? And where did it come from? And why?"

The point I want to make here is the meaning that's abstracted from the book is not *in* the book; it is in the mind. It's a point that starts to get recognized in philosophy in the 1700s by Kant, but isn't fully articulated until the 20th century. The meaning is not in the novel, the meaning is in one's mind. We construct meaning by the associations we make. If this seems strange, imagine a page in Dostoyevsky shown to a Russian reader and a non-Russian reader. Each gets exactly the same data, but the Russian has the associations to parse the Cyrillic script and make the written sense data come alive. The non-Russian sees exactly the same data; but his associations if he does not speak the language are nil and there is no meaning. Meaning therefore is imposed. It emerges by our imposing associations. It's not Dostoyevsky or the book *Independent People* that bring meaning to me—that's an illusion. It's *me* that brings meaning to *Independent People*. I'm making sense, I'm im-

posing associations, *I* impose meaning on what I'm seeing. Not just any old meaning, but the meaning that emerges from the associations the book makes with my neural memory.

Let me give you another example because I want to hammer on this point and derive a few things from it. There's a Yeats' poem that goes something like this: "Down by the salley gardens my love and I did meet; she passed the salley gardens on little snow-white feet. She bid me to take life easy, like the grass grows on the weirs, but I was young and foolish, and now am full of tears."³ These words will have different effects on different people—different meanings. Ask yourself what meaning you get out of weirs. For me this has enormous meaning because I and my friends played near weirs as children. (Weirs are little dams in a stream, usually covered with algae and some form of green trailing grass.) I also know what salley gardens are. But those who are not Irish will probably be affected differently. They may wonder: What are salley gardens anyway? Maybe Sally had a garden. Maybe there's such a thing as the Sally Gardens—maybe they exist on some estate near Dublin. In the absence of knowing what salley gardens are, you probably have an image of a garden well kept, surrounded by flowers and tended by keepers. But it's not that. The word in Gaelic is s-a-i-l-e-a-c-h, and it means "willow." So Yeats is near willows, and therefore likely near water. If there's a weir, the water is a stream or river. Once one has these associations, immediately the initial picture shifts. My point is that different meanings can be imposed on the same data. Different meanings that come from different associations.

Data—literary or economic—have no inherent meaning. They acquire meaning by our bringing meaning to them. And different people, with different experiences, will construct different meanings.

The Mind as a Fast-Pattern Completer

What conclusions does modern cognitive psychology draw from such examples? The first conclusion is that our brains are "associative engines" to use a phrase of Andy Clark, a philosopher and cognitive scientist from Washington University in St. Louis. (Clark, 1993). We're wonderful at association and in fact, in cognition, association is just about all we do. In association we impose intelligible patterns. To use another of Clark's labels, we are fast pattern completers. If I see a tail going around a corner, and it's a black swishy tail, I say, "There's a cat!" But it could be a small boy with

a tail on the end of a stick who's trying to fool me. But I don't do that. My mind is not built to do that. If I were strongly skeptical, I *could* do that, or if I saw some small boy playing pranks I could say, "Well, it's either a cat or a small boy." But in the absence of a small boy, all I'm really saying is, "Hey! I see a cat." But I didn't see a cat. I saw a black tail. A famous Bertrand Russell story makes the same point. A schoolboy, a parson and a mathematician are crossing from England into Scotland in a train. The schoolboy looks out and sees a black sheep and says, "Oh! Look! Sheep in Scotland are black!" The parson, who is learned, says, "No. Strictly speaking, all we can say is there is one sheep in Scotland that is black." The mathematician says, "No, still not correct. All we can really say is that we know that in Scotland there exists at least one sheep, at least one side of which is black."

Cognitive science repeatedly tells us that we don't think deductively as the mathematician did, we think associatively as the schoolboy did. And for a very good reason: Evolution has made it so. Our ability as humans a hundred thousand years ago to sniff the air and associate a fleeting humidity with the presence of water a few miles away had real survival value. Completing patterns fast, surmising the presence of water from the faintest of clues, helped us survive. Deductive logic did not; and in all but the most trivial of cases we do not use it at all. In fact, cognitive psychologists tell us that deductions themselves are primarily associative. I may say I can solve such-and-such a problem: it's a problem in spherical trigonometry. I then associate the problem with this framework. From there I associate structures and symbols with the sense data of the problem. And I proceed by such associations, stitching them together into a pattern. I'm not saying that association is all the human brain does, but cognitively, association is the main thing we do. And we do it fast. Our neural system searches fast over many associations before settling on one as a "meaning." Occasionally this process slows and we can see it in action, as with the 3-dimensional optical-illusion pictures that were popular a few years ago which appear flat and 2-dimensional until after staring for half a minute a 3-D picture "leaps out." So our brains process a large collection of associations into patterns—and a large set of metaphors which are merely more complicated associations with entailments. With metaphors we compare this to this and that to that, and if the comparison is good, we expect such and such to follow. Metaphor is a form of pattern association, and we process much information through metaphors. In sum, we have many different forms of associations: pictures; memories; metaphors; and theories, which are really elaborated metaphors. And this

collection, when it's fully operating, along with the rules for combining these (which are also associations) we call the mind.

Our minds then are extremely good at associating things, using metaphors, memories, structures, patterns, theories. In other words, the mind is not given. It's not an empty bucket for pouring data in. The mind itself is emergent. This idea is new in Western thinking but there's plenty of precedent for it in the East. The Neo-Confucian philosopher brothers Ch'eng Yi and Ch'eng Hao, writing during the Sung Dynasty about 900 years ago, both saw mind as emergent. They did not see the mind as a container, but rather as sets of ideas built one upon the other. The mind doesn't contain our ideas. It's these ideas—these associations—that instead contain the mind or constitute the mind. The mind is not fixed in any way; it consists in its associations and the apparatus to manipulate these. In this sense it's emergent. So strictly speaking I shouldn't say as I did earlier that meaning resides in the mind, because deep enough within cognitive philosophy the concept of mind itself dissolves. Meaning resides in associations our neural apparatus connects with the data presented. We are far now from seeing reasoning as deduction that takes place in a container of variables whose values are updated by "information." If reasoning is largely association, it depends on the past experiences of the reasoner. The framing of a situation, the "sense" made of it, are therefore dependent on the reasoner's history. And so is the outcome.

One final point about cognition. Sometimes we can say roughly that there is a "correct" meaning—a single, correct association. More often, in any situation of complication, there are multiple interpretations. We may hold one or we may hold many. Often, if we are trying to solve a puzzle, or to come to a decision such as the next move in a chess game, we make many hypothetical associations and search over these, perhaps retaining more than one until further evidence presents itself. In the black tail example, if I had indeed seen a small boy a few minutes earlier, I might hold in mind both "cat" and "prank" until further evidence arrived.

Modeling the Cognitive Process

All this is fine. But as economists how do we make use of it? How might we model the thinking process in problems that are complicated or ill-defined?

I would suggest the following, by way of distillation of the observations above: In problems of complication, as decision makers, economic agents look for ways to frame

the situation that faces them. They try to associate temporary internal models or patterns or hypotheses to frame the situation. And they work with these. They may single out one such pattern or model and carry out simplified deductions (at the level of Tic-tac-toe) on it, if they seek guidance for action. As further evidence from the environment comes in, they may strengthen or weaken their beliefs in their current models or hypotheses. They may also discard some when they cease to perform, and replace them as needed with new ones. In other words, where agents face problems of complication or ill-definition, they use clues from the situation to form hypothetical patterns, frameworks, associations. These hypothetical patterns fill the gaps in the agent's understanding.

Such a procedure enables us as humans to deal with complication: we construct plausible, simpler models that we can cope with. It enables us to deal with ill-definedness: where we have insufficient definition, we use working models fill the gap. Such behavior is inductive. It may look like ad-hoc and messy, but it is not antithetical to "reason," or to science for that matter. In fact, it's the way science itself operates and progresses.

More practically then, in a typical economics problem that plays out over time, we might set up a collection of agents, probably heterogeneous, and assume they make associations in the form of mental models, or hypotheses, or subjective beliefs. These beliefs might themselves take the form of simple mathematical expressions that can be used to describe or predict some variable or action; or of statistical hypotheses; or of condition/prediction rules ("If situation Q is observed/predict outcome or action D"). These will normally be subjective—they will differ among the agents. An agent may hold one in mind at a time, or several simultaneously, keeping track of the performance of each. When it comes time to make choices, the agent acts upon his currently most credible (or possibly most profitable) one. The others he keeps at the back of his mind, so to speak. As economists we will be tempted to say the agent rationally combines his several hypotheses. But cognitive psychology tells us we don't do this, we hold in mind many hypotheses at a time and act on the one currently most plausible. Once actions are taken the aggregative picture is updated, and agents update their confidence in each of their hypotheses.

This scheme I'm suggesting is of course also a simplification and abstraction. But it captures the idea that the agent is *imposing* meaning on the problem situation, or making sense of it by associating multiple frameworks, or belief structures, or hypotheses with it and allowing these to "compete." This is also a system in which learning takes

place. Agents "learn" which of their hypotheses work, and they "learn" also in the acts of discarding poorly performing hypotheses and generating new "ideas" to put in their place. Notice there is a built-in hysteresis: agents linger with their currently most believable hypothesis or belief model, but drop it when it no longer performs, in favor of a better one. A hypothesis or association or belief model is clung to not because it is "correct"—there is no way to know this—but rather because it has worked in the past and must cumulate a record of failure before it is worth discarding.

A key question remains. Where do the hypotheses or mental models come from? How are they generated? Behaviorally, this is a deep question in psychology, having to do with object representation, and pattern recognition. I will not go into it here. But there are some simple and practical options for modeling. Sometimes we might endow our agents with focal models—patterns or hypotheses that are obvious, simple and easily dealt with mentally. We might generate a "bank" of these and distribute them among the agents. Other times, given a suitable model-space, we might allow some similar intelligent search device such as the genetic algorithm to generate suitable models. The reader should note that whatever option is taken, the framework I've described is independent of the specific hypotheses or beliefs used, just as the consumer theory framework is independent of particular products chosen among.

Can such a scheme be put in practice in economics? The answer is yes. There is now a growing body of examples: the El Farol problem (Arthur 1994); the work of Tom Sargent (Sargent); the Santa Fe stock market study (Arthur, Holland, LeBaron, Palmer and Tayler, 1996). This type of study typically finds that "solutions"—patterns of beliefs and actions predicated upon these—need to be generated by computation because of the increased complication of heterogeneous beliefs. It also typically finds a richer world, a psychological world, where an ecology of beliefs about the problem in question emerges. Sometimes this ecology of hypotheses converges to some standard equilibrium of beliefs. More often it remains open-ended, always discovering new hypotheses, new ideas.

Cognition and Graduate Economic Education

Let me turn from modeling in economics to quite a different area that can benefit from the insights of cognitive science: the education of economists.

I want to start here by drawing attention to two ways in which we make sense: two types of association, not com-

pletely different and on a spectrum. Let me call one "theory" and the other "experience."

Theories are metaphors with entailments. If in 1705, Edmond Halley subscribed to Newton's gravitational theory and applied it to a comet that had previously appeared in 1531, 1607, and 1682, one entailment was that the comet would return in the year 1759. In using Newton's theory, Halley was making an association between a comet and the heavenly bodies Newton dealt with; and the entailments of the association allowed Halley to predict. I want to suggest that theories are *thin* associations: the theory fits if a narrow and precise set of conditions is fulfilled; and the entailments are also narrow and precise. Providing the theory fits correctly—is a good association—and is consistent within itself, then the entailments can be relied upon. Narrow fit, narrow entailments. Theories are in this way thin but powerful associations.

What about experience? Suppose I'm an executive sent to Korea, and I've never done business there. I arrive in Korea and I'm wondering how I shall act. I have no idea of how many times I should bow to my host, or if anybody bows to the host, or whether I should take my shoes off, or if I want to close the deal do I wait till the end of dinner or do I try to close the deal up front? But I do have a lot of experience in Japan and in China, and so I use these. In this case hundreds of pictures are going through my mind. This sort of association is more dream-like. It's richer. It covers a wider set of cases. It's suggestive of what will follow given what is. But it's much less accurate and less precise and less reliable than theory. So experience in the form of a wide collection of memories and pictures of situations—*thick* association—is also powerful. Its power lies in its width of coverage and its suggestiveness. Such experience is what we seek from human conversation and from taking in stories and novels and plays. We seek to draw into ourselves other people's experiences, to make their situations into our memory pictures that we can use later. In this way we construct and conjure a whole dream-like world where logic doesn't matter and precision doesn't matter, but where suggestiveness and coverage give power.

As I said earlier, these two types of association are not completely distinct; associations arrange themselves on a spectrum from narrowness and precision to width and suggestiveness.

What has this to do with graduate economic education? A lot. A great deal of education is the formation of associations; and the spectrum ranges from collections of narrow but precise theories on one side to wide but suggestive and imprecise pictures on the other. We need both types of asso-

ciation to function successfully as human beings.

In economics, graduate education at least in the first year or two consists in mastering 20 or 30 theoretical economic models—thin associations. These include the principal agent model, the overlapping generations model, the prisoner's dilemma model and so on. The idea is that these theoretical metaphors will later become useful associations. We hope that if the student is later employed say at the World Bank, she will be able to look at a situation and say, "This problem in African agriculture is partly a principal agent problem. It does have some overtones of overlapping generations, and it's also got this game theory component. So I can put together a hybrid version of the three models to get insight." All this is fine. It is fine that economics has recognized recurring structures that it has rendered into theories. We can hope and expect that a well-educated student will use these as association components later.

But models cannot be all that we teach. There's been a tendency in many graduate schools to increase teaching in theory at the expense of teaching in economic history and in case example. Students of course can still choose to study the experience-details of the economy; but they are aware that this may not enhance their graduate careers. In 1990 Colander and Klamer asked students how important having a "thorough knowledge of the economy" was to succeeding as an economist. Three percent thought it very important, and 68 percent thought it unimportant. Important was: "Being smart in the sense of being good at problem solving," and "excellence in mathematics." With this bias toward theory and away from experience, we eliminate the wider metaphors that come from history-experience—the thick associations. These allow students to put their models into perspective; they provide the vocabulary, so to speak, where theory provides the grammar; they provide a richness of thought and a breadth of association that theory cannot possibly match.

When a decision maker faces a situation of high complexity, say Bosnia in the mid-1990s, applying theory prematurely—a set of precise but narrowly applicable metaphors—can be dangerous. Let's say he is in the State Department looking at Bosnia and has been in graduate school in political science, doesn't have much experience and is full of theories. His reaction may be to shoehorn Bosnia into a pre-constructed framework. But in this situation it is better to wait and observe. And in observation to invoke a variable set of pictures on which he may conjure up a richer set of associations. Such free association comes from a study of history, not theory. "Well, it could be like a bit like the Bosnian crisis of 1908, but it's not unlike the situation under Turkish rule in 1831 when Husein seized power. On the

other hand, there are elements of the ethnic rivalries of 1875 that resulted in the Austro-Hungarians taking over." What's of use is to have thousands of such pictures from history, available for pondering and perusal. Eventually from such pondering and perusal—from dreamlike association—a composite set of hypotheses or composite picture might emerge. It's at this stage that theory might apply. Premature association without going through the richness of a wide set of pictures may be disastrous. Where I come from, Belfast—another complicated situation—we say: "If you're not confused, you don't know anything."

What about teaching the history of economic thought—another threatened discipline in economic education? From the cognitive point of view, the history of economic thought bestows on us an awareness of the associations we make. Without such awareness, associations can be unconscious and poorly suited to the case in hand. Consider the English painters who came to Australia in the late 1700s or early 1800s. These artists depicted trees in Australia as they would have depicted English trees; they were well trained in English art schools and knew how to paint trees. But in Australia the leaves of most trees—often eucalyptus trees—are thinner, and the sun shines through them. Trees there look different—lighter, more airy. It took a generation of Australian-born painters before the trees in paintings started to look like Australian trees. Before that European painters were unconsciously making European associations and imposing these upon Australia. Similarly Europeans depicted aboriginals in this early period as Europeans with dark skins. This is not to criticize artists. It is to be aware that the actions we take are built upon our unconscious associations. We need to be conscious of our associations and where they come from. We need to be suspicious of them. We need a Zen-like standing back and seeing from beginner's mind. We need an awareness that theories aren't exogenous—they were constructed by people with agendas from other times sometimes suited to the purposes of other times. We need knowledge of the history of economic thought to be fully aware of the associations we make in economics and their provenance.

So what in graduate economics education do we really want? We certainly need theory. As a theorist I'm all for theory. But we also need the rich pictures given by the study of history and institutions. We need both types of association: the theoretical, quantitative, precise frameworks and the dream-like, vivid pictures in their tens of thousands. To teach only theory is equivalent to training doctors by teaching only endocrinology and pathology, and not the wide diagnostics doctors learn on grand rounds. To operate only

with theory—think of driving a car—makes us beginners. It's not until we can seamlessly integrate theory and vivid pictures—theory and experience—that we become expert. I believe we are currently turning out students who lack those pictures. And in doing so, we're doing them a disservice.

Do Issues of Cognition Matter?

Perhaps in asking my fellow economists to think about the implications of cognition, I am asking for something useful but not necessary—a luxury? I don't believe so. Consider just one example. The Soviet Union in 1990-91 decided to go capitalist. And from us economists it got much advice. But our natural bias, given the current development of economics, was to concentrate upon a worthy, but imagined, general-equilibrium outcome where institutions would be in place and markets would work smoothly and incentives would be correct.

A cognitive view of economics might have balanced this ideal view with an awareness that Russians were not arriving

did they possess old structures both economic and political, they harbored from their 70 years of communism and earlier, czarist past old associations too—of what business means, of how one interacts with authorities, of how one organizes if one wants to make money, of what one does with economic power and wealth. More enlightened advice would have built upon an understanding of how these embedded structures and understandings would play out given the new possibilities. The subsequent history of Russia's experiment with capitalism showed that these matters of cognition had great importance.

Economic agents bring to their actions not just their preferences and endowments, but also their understandings—the associations and meanings they have derived from their history of previous actions and experiences. In many of the small, standard problems of economics, we can ignore this. In the larger issues of development and reconstruction, and in constructing an economics for problems of complication and ill-definition, we cannot. We need to take cognition seriously.

with empty minds to their version of capitalism. Not only

REFERENCES

Arthur, W. B. (1994), "Complexity in Economic Theory: Inductive Reasoning and Bounded Rationality," *American Economic Review*, **84** (2), 406-411.

Arthur W. B., J.H. Holland, B. LeBaron, R. Palmer, and P. Tayler, (1997), "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market," in *The Economy as an Evolving Complex System II*, Arthur, S. Durlauf, D. Lane, eds. Addison-Wesley: Reading, Mass.

Clark, A. (1993), *Associative Engines*, Cambridge: MIT Press.

Jaynes, J. (1976), *The Origin of Consciousness in the Breakdown of the Bicameral Mind*, Boston: Houghton Mifflin Company.

Laxness, H. (1935), *Independent People*.

Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica*, **55** (5), 999-1033.

Sargent, T.J. (1993), *Bounded Rationality in Macroeconomics*, New York: Oxford Univ. Press.

Notes

1. This paper originated as an after-dinner talk at the conference in Middlebury, Vermont that the book is based on. I have rewritten the talk but have tried to retain its informal flavor. I am grateful to David Colander for his editorial help in the rewriting.

2. Julian Jaynes: *The Origin of Consciousness in the Breakdown of the Bicameral Mind*.

3. The lines are lines 1, 2, 7, and 8 of an 8 line poem, "Down by the Salley Gardens," by W.B. Yeats. From *W.B. Yeats, the Poems*, ed. Richard J. Finneran, New York; Macmillan, 1983 (p. 20). "Down by the salley gardens my love and I did meet; / She passed the salley gardens on little snow-white feet. / She bid me take love easy, as the leaves grow on the tree, / But I, being young and foolish, with her would not agree. / In a field by the river my love and I did stand, / As on my leaning shoulder she laid her snow-white hand. / She bid me take life easy, as the grass grows on the weirs; / But I was young and foolish, and now am full of tears."