Artificial Life

Intermediate article

Norman H Packard, Prediction Company, Santa Fe, New Mexico, USA
Mark A Bedau, Reed College, Portland, Oregon, USA

CONTENTS

Artificial life is the study of life and life-like processes through simulation and synthesis.

INTRODUCTION

Artificial life literally means 'life made by human artifice rather than by nature'. It has come to refer to a broad, interdisciplinary endeavor that uses the simulation and synthesis of life-like processes to achieve any of several possible ends: to model life, to develop applications using intuitions and methods taken from life, or even to create life. The aim of creating life in a purely technological context is sometimes called 'strong artificial life'.

Artificial life is of interest to biologists because artificial life models can shed light on biological phenomena. It is relevant to engineers because it offers methods to generate and control complex behaviors that are difficult to generate or control using traditional approaches. But artificial life also has many other facets involving *inter alia* various aspects of cognitive science, economics, art, and even ethics.

There is not a consensus, even among workers in the field, on exactly what artificial life is, and many of its central concepts and working hypotheses are controversial. As a consequence, the field itself is evolving from year to year. This article provides a snapshot and highlights some controversies.

HISTORY

The roots of artificial life are quite varied, and many of its central concepts arose in earlier intellectual movements.

John von Neumann implemented the first artificial life model (without referring to it as such) with his famous creation of a self-reproducing, computation-universal entity using cellular automata. At the time, the construction was surprising, since many had argued its impossibility, for example on the grounds that such an entity would need to contain a description of itself, and that description would also need to contain a description, *ad infinitum*. Von Neumann was pursuing many of the very issues that drive artificial life today, such as understanding the spontaneous generation and evolution of complex adaptive structures; and he approached these issues with the extremely abstract methodology that typifies contemporary artificial life. Even in the absence of modern computational tools, von Neumann made striking progress.

Cybernetics developed at about the same time as von Neumann's work on cellular automata, and he attended some of its formative meetings. Norbert Wiener is usually considered to be the originator of the field (Wiener, 1948). It brought two separate foci to the study of life processes: the use of information theory and a deep study of the self-regulatory processes (homeostases) considered essential to life. Information theory typifies the abstractness and material-independence of the approach often taken within both cybernetics and artificial life. Both fields are associated with an extremely wide range of studies, from mathematics to art. As a discipline, cybernetics has evolved in divergent directions; in Europe, academic departments of cybernetics study rather specific biological phenomena, whereas in America cybernetics has tended to merge into systems theory, which generally aims toward formal mathematical studies. Scientists from both cybernetics and systems theory contribute substantially to contemporary artificial life.

Biology (i.e. the study of actual life) has provided many of the roots of artificial life. The subfields of biology that have contributed most are microbiology and genetics, evolution theory, ecology, and development. To date there are two main ways that artificial life has drawn on biology: crystalizing intuitions about life from the study of life, and using and developing models that were originally devised to study a specific biological phenomenon. A notable example of the latter is Kauffman's use of random Boolean networks (Kauffman, 1969, 1993). Biology has also influenced the problems studied in artificial life, since artificial life's models provide definite answers to problems that are intractable by the traditional methods of mathematical biology. Mainstream biologists are increasingly participating in artificial life, and the methods and approaches pioneered in artificial life are increasingly accepted within biology.

The most heavily represented discipline among contemporary researchers in artificial life is computer science. One set of artificial life's roots in computer science is embedded in artificial intelligence (AI), because living systems exhibit simple but striking forms of intelligence. Like AI, artificial life aims to understand a natural phenomenon through computational models. But in sharp contrast to AI, at least as it was originally formulated, artificial life tends to use bottom-up models in which desired behavior emerges in a number of computational stages, instead of top-down models that aim to yield the desired behavior directly (as with expert systems). In this respect, artificial life shares much with the connectionist movement that has recently swept through artificial intelligence and cognitive science. Artificial life has a related set of roots in machine learning, inspired by the robust and flexible processes by which living systems generate complex useful structures. In particular, some machine learning algorithms such as the genetic algorithm (Holland, 1975) are now seen as examples of artificial life applications, even though they existed before the field was named. New areas of computer science (e.g., evolutionary programming, autonomous agents) have increasingly strong links to artificial life. (*See* **Artificial Intelligence, Philosophy of**)

Physics and mathematics have also had a strong influence on artificial life. Statistical mechanics and thermodynamics have always claimed relevance to life, since life's formation of structure is a local reversal of the second law of thermodynamics, made possible by the energy flowing through a living system. Prigogine's thermodynamics of dissipative structures is the most modern description of this view. Statistical mechanics is also used to analyze some of the models used in artificial life that are sufficiently simple and abstract, such as random Boolean networks. Dynamical systems theory has also had various contributions, such as its formulation of the generic behavior in dynamical systems. And physics and dynamical systems have together spawned the development of synergetics and the study of complex systems (Wolfram, 1994), which are closely allied with artificial life. One of artificial life's main influences from physics and mathematics has been an emphasis on studying model systems that are simple enough to have broad generality and to facilitate quantitative analysis.

The first conference on artificial life (Langton, 1989), where the term 'artificial life' was coined, gave recognition to artificial life as a field in its own right, although it had been preceded by a similar conference entitled 'Evolution, Games, and Learning' (Farmer *et al*., 1986). Since then there have been many conferences on artificial life, with strong contributions worldwide (e.g., Bedau *et al*., 2000). In addition to the scientific influences described above, research in artificial life has also come to include elements of chemistry, psychology, linguistics, economics, sociology, anthropology, and philosophy.

CONCEPTS AND METHODOLOGY

Most entities that exhibit lifelike behavior are complex systems – systems made up of many elements simultaneously interacting with each other. One way to understand the global behavior of a complex system is to model that behavior with a simple system of equations that describe how global variables interact. By contrast, the characteristic approach followed in artificial life is to construct lower-level models that themselves are complex systems and then to iterate the models and observe the resulting global behavior. Such lower-level models are sometimes called agent- or individual-based models, because the whole system's behavior is represented only indirectly and arises merely out of the interactions of a collection of directly represented parts ('agents' or 'individuals').

As complex systems change over time, each element changes according to its state and the state of those 'neighbors' with which it interacts. Complex systems typically lack any central control, though they may have boundary conditions. The elements of a complex system are often simple compared to the whole system, and the rules by which the elements interact are also often simple. The behavior of a complex system is simply the aggregate of the changes over time of all of the system's elements. In rare cases the behavior of a complex system may actually be mathematically derived from the rules governing the elements' behavior, but typically a complex system's behavior cannot be discerned short of empirically observing the emergent behavior of its constituent parts. The elements of a complex system may be connected in a regular way, such as on a Euclidean lattice, or in an irregular way, such as on a random network. Interactions between elements may also be without a fixed pattern, as in molecular dynamics of a chemical soup or interaction of autonomous agents. When adaptation is part of a complex system's dynamics, it is sometimes described as a complex adaptive system. Examples of complex systems include cellular automata, Boolean networks, and neural networks. Examples of complex adaptive systems include neural networks undergoing a learning process and populations of entities evolving by natural selection.

One of the simplest examples of a complex system is the so-called 'game of life' devised by the mathematician John Conway (Berlekamp *et al*., 1982). The game of life is a two-state two-dimensional cellular automaton with a trivial nearest-neighbor rule. You can think of this 'game' as taking place on a two-dimensional rectangular grid of cells, analogous to a huge checkerboard. Time advances in discrete steps, and a cell's state at a given time is determined by the states of its eight neighboring cells according to the following simple 'birth–death' rule: A 'dead' cell becomes 'alive' if and only if exactly three neighbors were just 'alive', and a 'living' cell 'dies' if and only if fewer than two or more than three neighbors were just 'alive'. From inspection of the birth–death rule, nothing particular can be discerned regarding how the whole system will behave. But when the system is simulated, a rich variety of complicated dynamics can be observed and a complex zoo of structures can be identified and classified (blinkers, gliders, glider guns, logic switching circuits, etc.). It is even possible to construct a universal Turing machine in the game of life and other cellular automata, by cunningly configuring the initial configuration of living cells. In such constructions gliders perform a role of passing signals, and analyzing the computational potential of cellular automata on the basis of glider interactions has become a major research thrust.

Those who model complex adaptive systems encounter a tension resulting from two seemingly conflicting aims. To make a model 'realistic' one is driven to include complicated realistic details about the elements, but to see and understand the emergent global behavior clearly one is driven to simplify the elements as much as possible. Even though complex adaptive systems include systems whose elements and dynamical rules are highly complicated, the spirit of most artificial life work is to look for the complexity in the emergent global behavior of the system, rather than to program the complexity directly into the elements.

Computation is used extensively in the field of artificial life, usually to simulate models to generate data for studying those models. Simulation is essential for the study of complex adaptive systems for it plays the role that observation and experiment play in more conventional science. Having no access to significant computational machinery, Conway and his students first studied the game of life by physically mapping out dynamics with go stones at teatime. Now thousands of evolutionary generations for millions of sites can be computed in short order with a conventional home computer. Computational ability to simulate large-scale complex systems is the single most crucial development that enabled the field of artificial life to flourish and distinguish itself from precursors (such as cybernetics or systems theory).

The dependence of artificial life on simulation has led to debate within the field over the ontological status of the simulations themselves. One version of strong artificial life holds that life may be created completely within a simulation, with its own virtual reality, yet with the same ontological status as the phenomenon of life in the real world. Some hold, however, that simulated, virtual reality cannot possibly have the same ontological status as the reality we experience. These point out that a simulated hurricane can never cause us to become wet. They also believe that if artificial life is to achieve the status of reality, it must include an element of embodiment, an extension into the real, non-simulated world enabling an interaction with that world. Believers in the reality of simulation point out that a simulation has its own embodiment within a computer, that a simulation is not an abstract formula specifying a program but the actual running of a program in a real physical medium using real physical resources. The belief that artificial life has its own bona fide reality is particularly strong among those who generate experimental data with simulations.

Both living systems and artificial life models are commonly said to exhibit emergent behavior – indeed, many consider emergent behavior to be a hallmark of life – but the notion of emergence remains ill-defined. There is general agreement that the term has a precise meaning in some contexts, most notably to refer to the resultant aggregate global behavior of complex systems. The higher-level structures produced in Conway's game of life provide a classic example of this kind of emergent behavior. In spite of clear examples like the game of life, there is no agreement regarding how one might most usefully define emergence. Some believe that emergence is merely a form of surprise. On this view, emergence exists only in the eye of the beholder and whether a phenomenon is emergent or not depends on the mindset of the observer. Others believe that there is an objective, observer-independent definition of emergence in terms of whether a phenomenon is derivable from the dynamical rules, even if it is often difficult to tell *a priori* what can be derived from the dynamical rules underlying complex systems. These difficulties lead some to argue that the term 'emergence' should simply be dropped from the vocabulary of artificial life. However, this advice is not widely heeded at present.

Complexity is another commonly recognized hallmark of life, and this notion has also so far eluded satisfactory definition. Apparently several different concepts are involved, such as structural complexity, interaction complexity, and temporal complexity. To some, it seems obvious that the biosphere is quite complex at present and that its complexity has increased on an evolutionary timescale. But the difficulties of defining complexity lead others to claim that life's present complexity and its increase over time are either illusory or a contingent artifact of our particular evolutionary history. Understanding complexity and its increase through the course of evolution are at the center of much research in artificial life. In fact, one of the field's main goals at present is to produce and then understand open-ended evolution, an ongoing evolutionary process with continually increasing complexity.

Darwin's view of evolution, with its emphasis on survival of the fittest, implied that the process of adaptation was the key to the creation of intelligent design through life's evolution. However, the role and significance of adaptation is controversial today. Some hold that adaptation is the main force driving the changes observed in evolution. Others maintain that most of evolution consists of non-adaptive changes that simply explore a complex space of morphological forms. Still others claim that much of the apparent intelligence of complex systems is a necessary result of certain complex system architectures. Artificial life may shed light on this debate by providing many diverse examples of evolutionary processes, with an attendant ability to analyze the details of those processes in a way that is impossible with the biosphere, because the analogous assaying of historical data is currently impractical and much of the historical data is simply unavailable.

Analysis of adaptation has led to the idea of a fitness landscape. Organisms (or agents in an artificial life model) are considered to be specified by a genome (or sometimes a set of model parameters). The interaction of the organism with other organisms as well as with its environment yields an overall fitness of the organism, which is often thought of as a real-valued function over the space of possible genomes (or model parameters). In various applications of evolutionary algorithms, such as the genetic algorithm, specifying a fitness function is an essential part of defining the problem. In such cases, adaptation is a form of optimization, 'hill climbing in the fitness landscape'. In artificial life models, however, fitness is often not specified explicitly, but is a property emerging from the interactions of an organism with its world.

The concept of a fitness landscape as an analytical device suffers various limitations. One is that a fitness landscape is generally an approximation; the fitness landscape itself can evolve when organisms in a population interact strongly with each other. Another reason is that on an evolutionary timescale, the space on which a fitness function is defined is changing with the advent of new elements to the genome or new model parameters for artificial organisms. Simulating agent-based artificial life models is a natural and feasible way to study these more general situations.

MODELS AND PHENOMENA

Generally, artificial life models choose a level of biological life to model. The lowest stratum may be thought of as analogous to the chemical level; higher stages include modeling of simple organisms such as bacteria, constituents of more complex organisms such as cells, complex organisms themselves, and varieties of complex organisms that can give rise to ecologies. One might consider a holy grail of artificial life to be the discovery of a single model that can span all these levels; so far the field has had difficulty producing a model that spans even one connected pair of levels.

The most primitive phenomenon explored by some artificial life models is self-organization. Such models study how structure may emerge from unstructured ensembles of initial conditions. Naturally, one aim is to discover the emergence of lifelike structure; some models explicitly aim to model the origin of life – such as chemical soups from which fundamental structures such as self-maintaining autocatalytic networks might be seen to emerge. Models for the immune system are another example of a lifelike process emerging from chemical interactions. Self-organization has also been studied in models for higher-level living structures, such as metabolisms and cell networks, with Boolean networks whose dynamics converge to different structures depending on model metaparameters (Kauffman, 1969, 1993).

A host of models target the organismic level, sometimes with significant interactions between organisms. These models typically allow changes in the organisms as part of the system's dynamics (e.g., through a genetic mechanism), and the most common goal of research using these models is to identify and elucidate structure that emerges in the ensuing evolutionary process. Some models fit in between the chemical level and the organismic level, aiming to understand development by modeling interacting cells. Other models are inter-organismic, in the sense that they aim explicitly to model interactions between different types of organisms or agents. These models often contain elements of game theory.

Many of the models studied in artificial life should be viewed as 'purely digital' models. Purely digital models drop any pretense of modeling any pre-existing biological structures; their elements are digital constructs having no direct biological reference. Such models seek to produce novel, purely digital instances of biological phenomena in their emergent behavior. Conway's game of life is a purely digital model at the physical or chemical level, embodying an extremely simple and unique form of 'chemical' interactions (the birth–death rule). The self-organization exhibited in the game of life is not a representation of chemical self-organization in the real world but a wholly novel instance of this phenomenon. Another chemical-level model is AlChemy (Fontana, 1992), which consists of a mixture of 'reacting chemical molecules' that are actually simple programs that produce new programs as output when one program is given as input to another program.

One example of a purely digital model on the 'organismic' level is Tierra (Ray, 1992), which consists of 'organisms' that are actually simple self-replicating computer programs populating an environment consisting of computer memory. Tierra was a mature version of earlier efforts of a model called Core Wars (Dewdney, 1984) and has been followed by more developed versions such as Avida (Adami and Brown, 1994). In Tierra, the world is a one-dimensional ring of computer memory, which may be populated with instructions that are much like idealized microprocessor assembly language instructions (e.g., copy, jump, conditional branch, etc.). The instructions are the microscopic components of the model, and the model's central processing unit (CPU) implements the instructions in memory, creating a chemistry from which structure in the model can emerge. The model is generally seeded with a primordial organism consisting of a group of instructions that can copy itself to another place in memory. The copying is accompanied by errors (mutations) that can enhance the functionality of the organisms.

The accomplishments and shortcomings of most artificial life models are exemplified by those of Tierra. On the side of accomplishments, Tierra shows clear evidence of evolution, and the resulting emergence of structure and organization that were not 'programmed' into the model explicitly. Careful analysis of the evolutionary results reveals computational features such as evolution of subroutines and versions of parasitism. On the negative side, the model shows only one level of emergence (e.g., the model must be seeded by a primordial organism; evolution of an unstructured soup has not yet produced an emergent viable organism). Secondly, the evolution of the digital organisms appears to 'level off', reaching a stage where increasingly insignificant innovations are absorbed into the population, instead of displaying the open-ended evolution of natural systems. Reasons for this limitation include (1) simplicity of the model's evolutionary driving force (the evolutionary value of replication with minimal CPU time), (2) structural limitations on the space of innovations possible, which create limitations on organism functionality, and (3) structural limitations on organisms' ability to interact with each other and their environment. Different artificial life models have different detailed reasons for the two limitations we have discussed in Tierra, but the limitations are generally prevalent.

Another important area of artificial life is not so much a modeling activity as much as an implementation activity. This work aims to produce hardware implementations of lifelike processes. Some of these implementations are practical physical devices. But some of this activity is primarily theoretical, motivated by the belief that the only way to confront the hard questions about how life occurs in the physical world is to study real physical systems. Again, there is an analogy with biological levels. The 'chemical' level is represented by work on evolvable hardware, often using programmable logic arrays (e.g., Breyer *et al.*, 1998). The 'organismic' level is represented by recent work in evolutionary robotics (e.g., Cliff *et al.*, 1993). An 'ecological' level might be represented by the Internet along with its interactions with all its users on computers distributed around the world.

Artificial life, like its antecedent, cybernetics, has a peculiarly broad cultural scope extending beyond cut and dried scientific progress. This breadth is best exemplified by the work of Karl Sims (Sims, 1991), who has coupled rich image-producing computational environments with interactions between those environments and people watching the images at an exhibit. The result is an evolutionary system that is not constrained to live within the confines of a particular model's framework, but rather that is a coupling of two evolutionary subsystems, one of which is natural (the audience). Sims' interactive evolutionary art has produced several visually striking results, and human interaction seems to give the evolutionary system an open-ended quality characteristic of natural evolution.

FUTURE DIRECTIONS

One broad direction artificial life will continue to take in the future is that of synthesis: the synthesis of significant biological phenomena either within the context of model simulation or hardware implementation. A grave difficulty facing progress in this area is the lack of any quantitative basis of comparison for many of the biological phenomena artificial life aims to model. An example of this difficulty is modeling open-ended evolution. How could we know when this is achieved? In general, measurable characterization of phenomena is a prerequisite to quantitative comparison, and much progress is needed in order to achieve this for many target phenomena.

Probably the largest goal of the field is to understand the nature of life itself. This will be furthered to some extent with the quantitative comparisons just mentioned, but there is also a broader goal of discerning what the boundaries of life are, and how the idea of life might be extended to phenomena beyond biological life. Is there a sense in which financial markets or sociotechnical networks are alive, independent of the lives of their biological constituents? Many in the field of artificial life believe that, if the concept of life is properly framed and understood, such questions may well have a precise affirmative answer.

References

Adami C and Brown CT (1994) Evolutionary learning in the 2D artificial life system 'Avida'. In: Brooks R and Maes P (eds) *Artificial Life*, vol. IV, pp. 377–381. Cambridge, MA: Bradford/MIT Press.

Bedau MA, McCaskill JS, Packard NH and Rasmussen S (eds) (2000) *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*. Cambridge, MA: MIT Press.

Berlekamp ER, Conway JH and Guy RK (1982) *Winning Ways*, vol. II. New York, NY: Academic Press.

Breyer J, Ackermann J and McCaskill JS (1998) Evolving reaction–diffusion ecosystems with self-assembling structures in thin films. *Artificial Life* **4**: 25–40.

Cliff D, Harvey I and Husbands P (1993) Explorations in evolutionary robotics. *Adaptive Behavior* **2**: 73–110.

Dewdney A (1984) In a game called core wars hostile programs engage in a battle of bits. *Scientific American* **250**: 14–23.

Farmer JD, Lapedes A, Packard NH and Wendroff B (eds) (1986) *Evolution, Games, and Learning: Models for Adaptation for Machines and Nature*. Amsterdam, Netherlands: North-Holland.

Fontana W (1992) Algorithmic Chemistry. In: Langton CG, Taylor CE, Farmer JD and Rasmussen S (eds) *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pp. 159–209. Reading, CA: Addison Wesley.

Holland JH (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press. [Revised and expanded edition (1992) Cambridge, MA: MIT Press.]

Kauffman SA (1969) Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology* **119**: 437–467.

Kauffman SA (1993) *The Origins of Order: Self-Organization and Selection in Evolution*. New York, NY: Oxford University Press.

Langton CG (ed.) (1989) *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*. Redwood City, CA: Addison Wesley.

Ray T (1992) An approach to the synthesis of life. In: Langton CG, Taylor CE, Farmer JD and Rasmussen S (eds) *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pp. 371–408. Reading, CA: Addison Wesley.

Sims K (1991) Artificial evolution for computer graphics. *Computer Graphics* **25**: 319–328. (ACM SIGGRAPH '91 Conference Proceedings, Las Vegas, Nevada, July 1991.)

Wiener N (1948) *Cybernetics, or Control and Communication in the Animal and the Machine*. New York, NY: Wiley.

Wolfram S (1994) *Cellular Automata and Complexity: Collected Papers*. Reading, MA: Addison Wesley.

Further Reading

Adami C (1998) *Introduction to Artificial Life*. New York, NY: Springer.

Adami C, Belew RK, Kitano H and Taylor CE (1998) *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*. Cambridge, MA: MIT Press.

Bedau MA (1997) Weak emergence. *Philosophical Perspectives* **11**: 375–399.

Bedau M, McCaskill J, Packard N *et al*. (2000) Open problems in artificial life. *Artificial Life* **6**: 363–376.

Boden MA (1996) *The Philosophy of Artificial Life*. Oxford, UK: Oxford University Press.

Brooks RA and Maes P (1994) *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press.

Floreano D, Nicoud J-D and Mondada F (1999) *Advances in Artificial Life: 5th European Conference, ECAL'99*. Berlin, Germany: Springer.

Husbands P and Harvey I (1997) *Fourth European Conference on Artificial Life*. Cambridge, MA: MIT Press.

Langton CG (ed.) (1995) *Artificial Life: An Overview*. Cambridge, MA: MIT Press.

Langton CG, Taylor CE, Farmer JD and Rasmussen S (eds) (1992) *Artificial Life II: Proceedings of the Workshop on Artificial Life*. Reading, CA: Addison Wesley.

Levy S (1992) *Artificial Life, The Quest for a New Creation*. New York, NY: Pantheon.

Koza JR (1992) *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.

Mitchell M (1996) *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press.

Varela FJ and Bourgine P (1992) *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*. Cambridge, MA: MIT Press.

Asset Market Experiments

Intermediate article

Daniel Friedman, University of California, Santa Cruz, California, USA
Hugh M Kelley, Indiana University, Bloomington, Indiana, USA

CONTENTS

Laboratory evidence indicates that prices of financial assets such as stocks and bonds respond to changes in the assets' fundamental value but are also sometimes distorted by investors' cognitive and other biases.

INTRODUCTION

Any society allocates some resources to current consumption and some to investment, to building a better future. Asset markets determine the extent and form of investment in modern economies. Non-market allocation procedures such as those once used in Communist countries clearly worked less well and became less prevalent in the late twentieth century. Asset markets now have global scope and significance.

By definition, asset markets are efficient when asset prices reflect all relevant information about investment opportunities. Theory shows that efficient asset markets lead society to choose only the most productive investment prospects, and to choose the best overall level of investment.

The efficient asset price is called *fundamental value*. Actual asset prices are set by fallible human investors in imperfect markets, and thus may contain other components, called *bubbles*, that can lead to inefficient resource allocation and impair future wellbeing.

Laboratory and field evidence sheds light on asset market efficiency. Asset markets sometimes compensate for investors' cognitive biases, but at other times they amplify them and produce large bubbles. Laboratory experiments help to test policies intended to improve asset market efficiency.

FUNDAMENTAL ASSET VALUE

An asset is anything that provides its owner with a stream of benefits over time. Its economic value is the monetary equivalent of the net benefits it provides. Valuation of a real asset (such as a house, a pizza delivery car or a microprocessor production facility) involves estimating prices for the services the asset generates and for the resources required to maintain its productivity. This article will focus