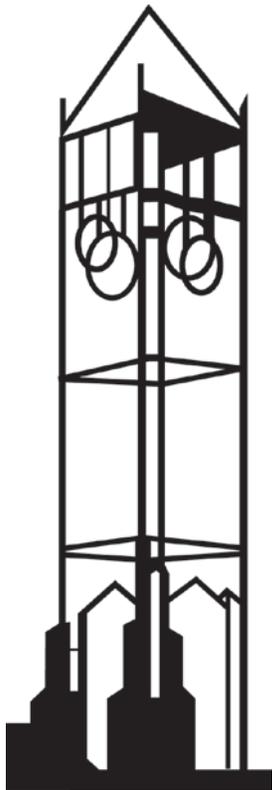


Agent-Based Modeling: The Right Mathematics for the Social Sciences?

Paul L. Borrill, Leigh Tesfatsion



Working Paper No. 10023
July 2010
Revised on November 2010

IOWA STATE UNIVERSITY
Department of Economics
Ames, Iowa, 50011-1070

Iowa State University does not discriminate on the basis of race, color, age, religion, national origin, sexual orientation, gender identity, sex, marital status, disability, or status as a U.S. veteran. Inquiries can be directed to the Director of Equal Opportunity and Diversity, 3680 Beardshear Hall, (515) 294-7612.

Agent-based modeling: The right mathematics for the social sciences?

Paul L. Borrill and Leigh Tesfatsion

Abstract:

This study provides a basic introduction to agent-based modeling (ABM) as a powerful blend of classical and constructive mathematics, with a primary focus on its applicability for social science research. The typical goals of ABM social science researchers are discussed along with the culture-dish nature of their computer experiments. The applicability of ABM for science more generally is also considered, with special attention to physics. Finally, two distinct types of ABM applications are summarized in order to illustrate concretely the duality of ABM: Real-world systems can not only be simulated with verisimilitude using ABM; they can also be efficiently and robustly designed and constructed on the basis of ABM principles.

1 Introduction

As in the physical sciences, theoretical modeling in the social sciences typically entails the specification and analysis of parameterized systems of differential equations. Many critical insights have been obtained by social scientists using this powerful classical mathematics approach.

Nevertheless, it is extremely difficult to capture physical, institutional, and behavioral aspects of social systems with empirical fidelity and still retain analytical

Paul Borrill
REPLICUS Software Corporation, 4287 Miranda Avenue, Palo Alto, CA 94306, USA, e-mail: paul@borrill.com,

Leigh Tesfatsion (Corresponding Author)
Professor of Economics, Mathematics, and Electrical & Computer Engineering, Heady 375, Iowa State University, Ames, IA, USA 50011-1070 USA, e-mail: tesfatsi@iastate.edu, <http://www.econ.iastate.edu/tesfatsi/>

Latest Revision: 29 June 2011. To appear in John Bryan Davis and D. Wade Hands, eds, *Elgar Companion to Recent Economic Methodology*, Edward Elgar Publishers, 2011.

tractability. Entities in social systems are neither infinitesimally small nor infinitely many, nor are their identities or behaviors necessarily indistinguishable from one another. Common simplifications, such as assumed homogeneous behaviors or the existence of single representative agents, are thus problematic. Moreover, the social sciences cannot separate observers from “the real world out there.” Rather, social scientists must consider multiple observers in a continual co-evolving interaction with each other and with their environment.

This leads us to question whether other forms of traditional mathematics, or even new forms of mathematics, might better serve the purposes of social scientists. In short, what is the “right” mathematics for the social sciences? Moreover, if a “right” mathematics exists for the social sciences, what are the implications for the physical sciences? And what can the social and physical sciences learn from each other?

As elaborated in Bridges (2009), constructive mathematics is distinguished from classical mathematics by the strict interpretation of “there exists” (\exists) as “we can construct.” Classical mathematicians accept the *law of the excluded middle (LEM)*: For any proposition P, either P is true or its negation is true; there is no middle ground that evades this decidability logic. Thus, classical mathematicians accept existence proofs based on proof by contradiction: If the negation of P is not true, then P must be true. In contrast, constructive mathematicians require a direct proof that P is true in the form of a computational procedure in order to rule out both the falseness *and the undecidability* of P. Constructive proofs can, in principle, be realized as computer programs. Constructive mathematics thus embodies the fundamental concepts of information, and the limitations on knowledge, implied by modern computability theory (Bridges 1999, Eberbach et al. 2004).

This distinction provides a dramatically different perspective on how we perceive models in our mind in relation to the real-world systems they are intended to represent. For example, social system modelers using classical mathematics typically assume (explicitly or implicitly) that all modeled decision makers share common knowledge about an objective reality, even if there is no constructive way in which these decision makers could attain this common knowledge. In contrast, social system modelers advocating a constructive mathematics approach have argued that the “reality” of each modeled decision maker ought to be limited to whatever that decision maker is able to compute (Velupillai 2010).

In this study we argue that *Agent-Based Modeling (ABM)* is an alternative and potentially more appropriate form of mathematics for the social sciences. Roughly, ABM is the computational modeling of systems as collections of autonomous interacting entities. As will be clarified in subsequent sections, ABM is a powerful blend of classical and constructive mathematical approaches.

Section 2 provides a basic introduction to the ABM methodology as an alternative form of mathematics, with a primary focus on its applicability for social science research.¹ Section 3 describes typical goals of ABM social science researchers and

¹ Detailed discussions of the ABM methodology in relation to the social sciences can be found in Epstein (2006), Gilbert (2007), Macy and Willer (2002), and Tesfatsion and Judd (2006). Annotated pointers to on-line ABM readings, research groups, and software resources can be found in Axelrod and Tesfatsion (2010) and Tesfatsion (2010a).

the “culture-dish” nature of their computer experiments. It also provides pointers to active ABM social science research areas. The applicability of ABM for science more generally is considered in Sections 4 and 5, with special attention to physics. This broader perspective highlights the conceptual similarity and applicability of agent representations across various scientific disciplines. Section 6 summarizes two ABM applications that illustrate concretely the duality of ABM: Real-world systems can not only be *simulated* with verisimilitude using ABM; they can also be efficiently and robustly *designed and constructed* on the basis of ABM principles. Concluding remarks are provided in Section 7.

2 What is ABM?

2.1 ABM as an Alternative Form of Mathematics

Social systems consist of heterogeneous communicating entities in an evolving network of relationships. One branch of mathematics that deals with networks of relationships is graph theory, and significant new perspectives and results are emerging from that field (Albert and Barabási 2002). Another is category theory, which specifies relationships (morphisms) among collections of objects as first-class citizens along with the objects themselves (MacLane 1998). Category theory is being promoted as the intellectual successor to set theory, which for generations has been considered the foundation of mathematics.

Graph theory and category theory are powerful tools supporting deductive reasoning in many sciences, including the social sciences. Indeed, as discussed by Laubenbacher et al. (2009, Section 6.3), certain classes of ABMs representable as finite dynamical system “objects” with appropriately defined types of morphisms can be shown to constitute a category. Nevertheless, the classical idea that we can deduce solutions (or “future states”) for systems *a priori*, purely from a study of their structural characteristics, is beginning to be overshadowed by the realization that many systems are computationally irreducible (Laughlin and Pines, 2000).

More precisely, for systems that are strongly interactive and/or highly sensitive to initial conditions, it is often not practical (or even possible) to predict their global outcomes in advance of actual implementation even when their laws of motion are known. Examples include John Conway’s Game of Life, Aristid Lindenmayer’s L-systems, Benoit Mandelbrot’s generated fractal sets, Stephen Wolfram’s elementary cellular automata (class four), and a long list of other foundational contributions too numerous to mention. Social systems appear to be subject to this form of computational irreducibility. Not only in practice, but now also in theory, we have come to realize that the only option we have to understand the global properties of many social systems of interest is to build and run computer models of these systems and observe what happens.

The computer modeling approach advocated here is the relatively young and still-developing methodology known as *Agent-Based Modeling (ABM)*. In this approach, systems are modeled as collections of autonomous interacting entities (“agents”) with encapsulated functionality that operate within a computational world.

As elaborated below, ABM combines constructive and classical modeling approaches. As is true for real people, agents can only acquire new data about their world constructively, through interactions. Nevertheless, again like real people, ABM agents can have *uncomputable beliefs* about their world that influence their interactions. These uncomputable beliefs can arise from inborn (initially configured) attributes, from communications received from other agents, and/or from the use of non-constructive methods (e.g., proof by contradiction) to interpret acquired data. These uncomputable beliefs enable agents to make creative leaps, to come up with new ideas about their world not currently supportable by measurements, observations, or logical extrapolations from existing information – in short, to be Smolin’s *seers* (Smolin, Chapter 18, 2002).

The import of this ABM blending of constructive and classical mathematics depends upon the purpose at hand. For example, for descriptive purposes, it permits human behavior to be captured with greater fidelity than simple algorithmic representations. For optimization purposes, it permits a deeper and more creative exploration of large domains, a melding of experience-tempered guesswork with step-by-step computation that could vastly extend the power of traditional finite search methods.

ABMs with no runtime interaction with external systems are Turing computable; hence, in principle, they can be equivalently expressed as finite systems of discrete-time recursive equations over finite state domains (Axtell, Section 2.1, 2003; Epstein, Chapter 2, 2006; Laubenbacher et al. 2009). However, ABMs can also be constructed to support persistent run-time interactions between computer agents and real-world entities via general types of input-output data streams. That is, ABMs can be *data-driven dynamic applications systems* (Darema 2005).

Data-driven ABMs are a form of interactive system called *Super-Turing Machines* by Eberbach et al. (2004) and Goldin et al. (2006). The latter authors argue that Super-Turing Machines constitute a new class of computation models capable of going beyond Turing Machines and algorithms. This claim is disputed in part by Prasse and Rittgen (1998). Nevertheless, the possibility of constructing ABMs as open systems with external runtime interactions and information flows has exciting practical ramifications even if the full computational and philosophical implications remain controversial (LeBaron and Tesfatsion 2008).

Finally, as will be illustrated concretely in Section 6.3, ABM is more than a modeling methodology; ABM principles can also be used to construct real-world systems. Examples include shopbots, automated Internet auctions, smart-grid electronic devices, and data storage systems. The resulting real-world systems can, in turn, be simulated by ABMs that mimic their basic architecture and constituent agent types. This natural duality affords researchers an opportunity, perhaps for the first time in history, to faithfully match the behavior of computer models to real de-

ployed systems. This should permit unprecedented predictive power, well beyond that achievable with either analytical models or traditional simulations.

We therefore view ABM as a bona fide member of the mathematician’s toolbox, suitable for the study of complex interactive processes in all scientific disciplines. Nevertheless, ABM is particularly congruent to social systems because of the ease of mapping agents to recognizable social entities and the natural hierarchical self-organization readily seen in social systems. We thus see the social sciences, including economics, as among the most promising areas for ABM application.

2.2 ABM Agents

Agents in ABM are entities that encapsulate data as well as methods that act on this data.² As illustrated in Fig. 1, ABM agents can represent a broad spectrum of entities ranging from passive physical materials governed by relatively simple dynamical methods, such as physical decay, to individual or group *decision-making agents (DMAgents)* with social capabilities.

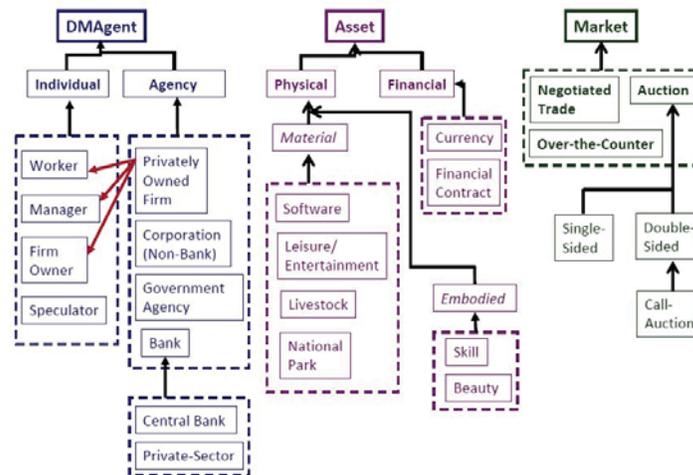


Fig. 1 Illustrative partial agent hierarchy for an economic ABM. Upward-pointing (black) arrows denote “is a” relationships and downward-pointing (red) arrows denote “has a” relationships.

² *Object-oriented programming (OOP)* represents “objects” as encapsulated bundles of data and methods. ABMs are now commonly implemented either directly in OOP languages or by means of toolkits based on OOP languages. Nevertheless, the current structure of these languages is overly rigid for many ABM purposes in terms of the permitted forms of encapsulation, distribution, inter-agent communication, and the envisioned degree of object autonomy (Jennings 2000). Although work-arounds are possible, ABM researchers would surely benefit from having a “new kind of programming” built from scratch with greater input from social scientists. Indeed, many initiatives along these lines are already under way (Teshfatsion 2010b).

ABM provides wide latitude when it comes to specifying the complexity and plasticity of agent representations. Analytical tractability is no longer a valid excuse for simplifications. Rather, for any particular purpose at hand, ABM researchers must decide the most appropriate level for their initial agent representations, the degree to which individual agent representations can undergo structural change over time, and the degree to which agent populations can evolve over time.

For example, an individual ABM agent can have a small number of simple fixed methods resulting in a relatively small range of if-then individually expressed behaviors within its world, similar to the fixed rules in cellular automata (von Neumann, 1951; Wolfram 2002). Just as the simple fixed rules of a chess game can produce an enormously large space of different games through player interactions, the simple fixed methods of individual agents within ABMs can produce unexpectedly rich global system behaviors through agent interactions.

Alternatively, individual ABM agents can have methods permitting more complex behaviors characteristic of people in real life. These behaviors can include: state-conditioned adaptive response (if this happens, what should I do?); anticipatory learning (if I do this, what will happen?); intertemporal planning; social communication; goal-directed learning leading to changes in state-conditioned responses; and reproduction (birth and death) leading to changes in the composition of agent populations.

The goals guiding goal-directed learning and action choice for ABM agents can be inborn or evolved. Moreover, these goals can be open-ended in form: for example, “maintain high power status,” or “avoid bankruptcy,” or “earn as much as possible” for oneself or for a group. An ABM agent with open-ended goals is more aptly described as a tracker rather than as an optimizer or satisficer, and its goal-seeking methods have no natural finite termination point apart from the termination of the agent itself.

The data and methods of each ABM agent are *encapsulated* in the sense that their form and content can be hidden from other agents. An agent communicates with other agents only through its *public interface*, the subset of its data and methods that other agents are permitted to see.

Agent encapsulation gives ABMs a striking resemblance to real-world systems. Information hiding (state containment) results in uncertainty in agent interactions, in the sense that agents can never be entirely certain how other agents will behave. Even if an agent is acting in accordance with a fixed private behavioral method, it can appear as a “different entity” in different interactions at different times due to the differences in its expressed behaviors induced in these interactions.

Moreover, in any ABM that is *closed* – that is, without external runtime interactions – agent encapsulation enforces the real-world constraint that all calculations must be carried out by the agents that actually reside within the ABM world. Free-floating procedures and restrictions influencing ABM world outcomes, such as global continuity or equilibrium conditions externally imposed across agents, are not permitted. Conversely, the procedures and restrictions encapsulated in the methods of a particular ABM agent can only be implemented using the particular resources available to that agent. An ABM agent that exhausts its resources is

constrained in its future ability to act effectively within its world. Thus, relative to traditional equation-based modeling, agent encapsulation in ABM permits a more realistic representation of real-world systems composed of interacting distributed entities with limited information, limited possible responses, limited material resources, and limited computational capabilities.

A key aspect of decision-making agents (DMAgents) in ABM is their increased autonomy relative to the decision makers appearing in analytical social science models. This increased autonomy arises from agent encapsulation. DMAgents can self-activate and self-determine their actions on the basis of hidden internal data and methods. These methods can include *pseudo-random number generators (PRNGs)*³ permitting randomizations of behaviors and decisions. For example, DMAgents might use “coin flips” to decide among equally preferred actions or action delays, mixed strategies in game situations to avoid exploitable predictability, and mutations (random perturbations) of normal routines to explore new possibilities. Moreover, the data and methods of DMAgents can change over time as they interact within their world and learn from these interactions.

Indeed, DMAgents in ABMs should be able to pass the following *constructive replacement test*: Given any DMAgent interacting within an ABM through its public interface, it should be feasible to replace this DMAgent with a person that interacts through this same public interface. This leads naturally to the idea of a *constructive Turing test*: Would a person interacting within an experimental framework involving a mixture of human and computer-agent participants be able to discern which were human and which were computer-generated? (Barr et al. 2008).

These proposed tests reflect the dual ability of ABMs to both represent and synthesize real-world systems. Already in ABM we are seeing a blurring of the lines between direct human control of computer-agent representatives (avatars), modeling of human behaviors via computer agents, and self-directed behaviors by autonomous computer agents. These advances in human-computer interfaces could ultimately revolutionize the theory and practice of social science.

2.3 ABM Horizontal and Hierarchical Organization

ABM agents are *distributed* in the sense that each agent experiences its own explicit or implicit locality within its larger world. Even if ABM agents are initially specified to be structurally identical, they can change or evolve over time to have widely varying data and methods supporting persistent cross-sectional behavioral diversity.

On the other hand, ABM agents are *connected* in the sense that each agent is embedded in a network of links representing some form of interaction (communication, trades,...) with other agents. The form and strength of these interaction networks can evolve over time through necessity (e.g., the death of agents) as well as through choice and chance.

³ Alternatively, “true” random data can be streamed directly into ABM experiments in place of PRNG-generated data, a possibility that raises interesting philosophical questions.

ABMs are horizontally scalable. An ABM initially developed on a single computer with a small number of interacting agents to facilitate debugging and compilation can subsequently be expanded to include larger numbers of interacting agents distributed across multiple machines on a network with the additional memory and CPU cores used to accelerate the simulation (Axtell 2003). Indeed, in principle, “cloudbursting” an ABM onto cloud-computing infrastructures would allow scaling out an ABM to any arbitrary size.⁴

In addition, ABM facilitates the modeling of hierarchical constructions. ABM agents can include other agents as members. Examples include household agents with multiple family members, community agents with multiple households as members, and nation agents with multiple communities as members. Moreover, member agents can reason about the larger agent of which they are a constituent part. Consequently, ABM enables the modeling of systems encompassing nested self-referential subsystems as commonly seen in the real world.

ABM also permits a more realistic modeling of *changes* in horizontal and hierarchical organization (Chang and Harrington 2006). For example, real-world firms often modify their internal organization over time to better compete with other firms. The standard economic modeling of firms by means of parameterized equations is therefore problematic. In contrast, in ABM a firm can be represented as an agent that includes other agents as members. As depicted in Fig. 1, these member agents might include “workers,” “managers,” and a “firm owner.” The horizontal extent and hierarchical organization of these member agents can then evolve naturally over time through hiring, firing, layoffs, resignations, and internal reorganization decisions.

The ability of ABM to model fluidity in horizontal and hierarchical organization in turn facilitates the modeling of complex real-world innovation processes (Chen and Chie 2007; Dawid 2006; Gilbert et al. 2001). Consider Google. A key aspect of Google’s competitive position is the continual development of new search algorithms, new ways of advertising, and so forth. This innovation occurs in at least two distinct ways. Google continually hires new employees that bring new ideas into Google from other places. Also, existing Google employees are encouraged through the local rules of their collective to generate new ideas. Indeed, Google has claimed that half of its new products come from ideas generated by engineers during the 20% worktime during which they are free to work on whatever they wish. The structure and behavior of Google thus exhibits fundamental change over time even as Google continues to maintain its identity within the world as “Google”.

These types of innovation processes are not easily modeled in terms of parameterized equations. ABM on the other hand is particularly well suited for the modeling of shifting landscapes involving changes in the behavioral methods of individual agents as well as evolutionary changes in the composition of agent populations.

⁴ *Cloud computing* is Internet-based computing. *Cloudbursting* permits a computational load to dynamically expand into additional resources in other private or public clouds as needed.

3 ABM as an Exploratory Research Tool

3.1 ABM Research Goals

Normative ABM research commonly addresses two basic “what if” questions with regard to system changes. First, do intended consequences actually arise? Second, under what conditions might a system change give rise to unintended consequences?

For example, would a change in the pricing rule for a market generate more efficient outcomes over time, as desired? Or would the change in fact lead to inefficiency as strategic participants within the system waste resources in their attempts to exploit the changed pricing rule for their own advantage? Would a change in a standard operating procedure or a human resource policy lead to improved profits for a firm, as planned? Or would the change in rules in fact reduce the firm’s profits due to unanticipated incompatibilities with the corporate culture?

A common goal of descriptive ABM research is to provide possible “generative explanations” for observed empirical regularities (Epstein 2006). Can an observed empirical phenomenon be reliably generated by a particular form of ABM starting from particular forms of initial conditions? For example, can the empirically observed thick tails for stock return distributions be reliably reproduced within ABM stock market frameworks that include suitably heterogeneous trader agents?

Other forms of ABM research involve qualitative insight. Under what conditions might a system give rise to unanticipated behaviors that lead to deeper intuitions about its nature? The quintessential example here is the venerable yet still unresolved concern of economists such as Adam Smith (1776) and Friedrich von Hayek (1948) to understand the surprising ability of traders in decentralized market economies to self-organize into resilient trade networks.

The ideal goal of qualitative ABM research is to characterize the complete dynamic landscape for a system. This includes the behaviors exhibited at any equilibria that might exist as well as the behaviors exhibited in the basins of attraction associated with these equilibria. It also includes a characterization of any phase transitions exhibited by a system in response to changes in its structural form or scale. Examples include the transition of water from steam to liquid to ice in response to reduced temperature, the transition from unpredictable to predictable outcomes in minority games as the ratio of the number of resolvable past play histories to the number of agents increases (Farmer et al. 2005), and the bifurcation of chaotic systems into foliations of distinct possible long-run behaviors in response to small perturbations in their parameter values (Devaney and Keen 1989).

Whatever the exact nature of their objectives, ABM researchers must also address challenging model verification and empirical validation issues (Tesfatsion 2010c). *Verification* concerns consistency with objectives: Does an ABM do what a researcher intends, or is there some form of logical or conceptual programming error? *Empirical validation* concerns consistency with empirical reality: Does an ABM appropriately capture the salient characteristics of a real-world system of interest, and does it provide outcomes that cohere with empirical observations?

3.2 *ABM Research as Culture-Dish Experimentation*

ABM researchers use controlled computer experiments to investigate how large-scale effects arise from the micro-level interactions of dispersed autonomous agents. In principle, as in wetware culture-dish experimentation, the only intervention permitted by ABM researchers is the setting of initial experimental conditions.

In carrying out an ABM experiment, an ABM researcher typically implements the following eight steps in sequence:

- *Step One:* Develop an experimental design for the systematic exploration of a theoretical issue of interest.
- *Step Two:* Construct a computer world (“culture dish”) consisting of a collection of constituent agents appropriate for the study of this theoretical issue.
- *Step Three:* Configure the computer world in accordance with the experimental design.
- *Step Four:* Compile and run the computer world with no further external interference and record world outcomes of interest.
- *Step Five:* Repeat this “same” computer experiment multiple times for multiple PRNG seed values to generate an *ensemble of runs* from which sample distributions for recorded world outcomes can be derived.
- *Step Six:* Repeatedly iterate steps three through five until the full range of configurations specified under the experimental design has been explored.
- *Step Seven:* Analyze the resulting sample distributions for recorded world outcomes and summarize their theoretical implications.
- *Step Eight:* Use these theoretical summaries to form hypotheses (conjectures) that can be brought to historical or real-time data for testing and empirical validation.

The intended meaning of some of these steps is more fully explained below.

In Step One, for example, a researcher might be interested in exploring whether differences in learning have systematic effects on the formation of stock price bubbles. To examine this issue, the researcher might develop an experimental design in which traders with variously specified mixes of learning capabilities engage in a sequence of stock market trades.

With regard to Step Two, the initial data and methods of the computer world determine the initial physical realm (e.g., spatial landscape) within which the constituent agents interact as well as simulation controls such as simulation stopping rules and non-perturbational instrumentation devices (e.g., pause buttons). The initial data and methods of the constituent agents determine the initial sequencing of their actions as well as the initial pattern of their interactions. For scientific investigations of real-world phenomena, these constructions should correspond to actual or proposed real-world counterparts and should reflect the physical, institutional, and behavioral aspects of these real-world counterparts.

The configuration in Step Three typically includes the initial number of constituent agents, the specification of fixed agent attributes (e.g., learning capabilities), initial settings for variable agent attributes (e.g., available resources), and the

initial spatial or social network determining permissible channels for initial agent interactions. It also includes the specification of seed values for any pseudo-random number generators (PRNGs) appearing in agent methods. The *(initial) state of the world* then consists of the (initially configured) data and methods for the world and for each of its constituent agents.

In Step Four the initial state of the world is permitted to evolve, driven solely by agent interactions. The only further role permitted for the external ABM researcher in this evolutionary process is the non-perturbational observation and recording of world outcomes of interest.

In Step Seven the theoretical summaries of experimental findings can often be enhanced by graphical visualizations. For example, grey-scale or color heatmaps can be used to display large data sets to facilitate discovery of any emergent global patterns.

3.3 ABM Research Areas

The range of ABM research is now extensive. Only a few areas of study relevant for social scientists are noted here to indicate the versatility of the methodology.

Within economics, ABM research areas that have been particularly active in recent years include agricultural and environmental economics, automated markets, business and management, electricity markets, financial economics, industrial organization, labor markets, macroeconomics, political economy, and economic network formation. A list of pointers to ABM research sites in these areas can be accessed at Tesfatsion (2010d).

Within the social sciences more generally, highly active ABM research areas include emergence of collective behavior, evolution of cooperation and trust, innovation, institutional design, learning, norms, social influence, and social network formation. Pointers to selected work in these areas can be accessed at Axelrod and Tesfatsion (2010).

Finally, ABM interdisciplinary application areas relevant for social scientists include ecological systems (Grimm and Railsback 2005), epidemiology (Epstein, 2006, Chapter 12), health care management (Huang et al. 1995), information storage and management (Borrill 2005, 2008), land use (Parker et al. 2003), military planning (Cioppa et al. 2004), transportation systems (Nagel and Wagner 2006), and urban planning (Batty 2005).

4 Alternative Modeling Modalities

4.1 *God's-Eye-View or Local-Observer-View?*

Science has traditionally been explored by means of analytical models based on classical mathematical principles. These models have been used to amplify human intuition about the way our world works. Computer models have frequently been seen as a poor substitute, a tool for determining approximate solutions to intractable analytical models.

Yet there is another side to this story. In analytical modeling, as well as in computer modeling used as an approximation tool, systems are typically represented from a *God's-Eye-View (GEV)*. The mathematician or programmer presides over the modeled world like some form of Laplace's demon, able in principle to discern the entire course of world events based on a complete understanding of the initial state of the world as well as its laws of motion.

Recently, however, it has been proven that Laplace was wrong to claim the future can be predicted without error given sufficient knowledge of the present, even in a classical non-chaotic universe (Binder 2008; Wolpert 2008). The capabilities of physical inference devices are inherently limited, not because of chaotic dynamics or quantum mechanical indeterminism, but rather due to a "Cantor diagonalization" demonstration that at least some portion of knowledge will always remain unavailable to any one inference device.

Constructive mathematics is well-matched to this reality because it relies solely upon a *Local-Observer-View (LOV)*. The data that a constructive mathematician can acquire about a system is limited to what the mathematician can obtain by means of computations.⁵ As developed by Errett Bishop, constructive mathematics is based on intuitionistic logic (Bridges 1999), a deductive system⁶ \mathcal{D} that does not include the law of the excluded middle ($P \vee \neg P$) among its rules of inference. Intuitionistic logic permits three logical possibilities for a proposition: true, false, or undecidable. A proposition is *true* (or *false*) relative to \mathcal{D} if its truth (or falsity) can be established by a computation within \mathcal{D} . A proposition whose truth or falsity cannot be established by a computation within \mathcal{D} is said to be *undecidable* relative to \mathcal{D} . In addition, the logical status of a proposition can be *open* relative to \mathcal{D} in the sense that its classification within \mathcal{D} has not been established. The uncountably large class of undecidable and open propositions (problems) relative to \mathcal{D} dominates the attention of many current constructive mathematicians (Ambos-Spies and Fejer, 2006).

⁵ According to Soare (1996), "a *computation* is a process whereby we proceed from initially given objects, called inputs, according to a fixed set of rules, called a program, procedure, or algorithm, through a series of steps and arrive at the end of these steps with a final result, called the output."

⁶ Given a formal language \mathcal{L} consisting of formulas (propositions) well-constructed in accordance with some prescribed syntax, a *deductive system* expressed in \mathcal{L} consists of axioms $L \subseteq \mathcal{L}$ together with a set of binary relations (rules of inference) on L that can be used to derive the theorems (conclusions) of the system (Dalen 2008).

ABM supports a LOV modeling approach in the sense that the “reality” of each ABM agent is confined to the network of agents within which it interacts. An ABM agent starts with a configuration of data and methods constituting its initial understanding of its world. The agent can then migrate from one part of its world to another, redefining its locality within this world by adding and deleting links with other agents; yet it is always restricted to interactions with other agents only one link away. Communication in ABM is thus a percolation process, and it is only through such percolation processes that new data about the world is acquired and global properties emerge.

On the other hand, although ABM agents acquire new data constructively through interactions, they can be configured so that portions of their initial data represent GEV (uncomputable) assertions about their world. They can also believe GEV assertions communicated to them by other agents, and they can have methods for interpreting data that entail the use of non-constructive deductions (e.g., proof by contradiction). Thus, as is true for real people, the content and timing of the constructive actions that ABM agents take within their world can be influenced by uncomputable beliefs.

4.2 Time and Asynchronicity in ABM

An important theoretical and practical concern for ABM researchers is how to specify the relative timing of agent interactions and the methods by which agents update their internal states based on these interactions. Careless treatments of these timing issues can induce undesirable artifacts in simulation outcomes, hide important potential system behaviors, and even result in a complete inability to generate empirically relevant results (Axtell 2001; Newth and Cornforth 2009; Tosic 2005).

In particular, considerable care must be exercised when running ABM simulations on conventional computer architectures – for example, on existing high performance computing (HPC) platforms – or even on commodity (multiple-vendor open-standard) hardware platforms that use a non-uniform memory architecture (NUMA). These basic computer hardware configurations can impose hidden constraints on temporal relationships among agents that a modeler does not intend, beyond simple variations in the behavior of the cache/memory hierarchy from one simulation run to the next.

For example, key areas of a computing infrastructure in which unintended *distributed simultaneity assumptions* can arise include: pre-emptive schedulers in the operating system when running multiple processes on a single processor; multi-threading within each process; communication through a shared memory bus on a symmetric multiprocessor; a common clock driving multiple cores at once on the same chip; and a single isochronous global clock in a single-instruction multiple-data (SIMD) graphic processing unit used for agent modeling in an array. If this listing leaves the reader with the impression that it is almost impossible not to have some potentially problematic simultaneity assumptions lurking somewhere in their

computational infrastructure, then a crucial point will have been effectively conveyed.

More fundamentally, resolution of timing issues forces ABM researchers to think deeply about the nature of the reality they are supposedly modeling. Conventional modeling of dynamical systems is frequently based on an implicit assumption that there exists a global (absolute or Newtonian) clock. This clock permits the perfectly synchronized updating of all components of the system state vector at each successive time step, where a time step is represented as an instant or interval along a real time line. The empirical artificiality of global clocks permitting perfect synchronization is abundantly clear. Most real systems are massively asynchronous in nature, and any implied simultaneity across distributed systems has long since been proven false in physics. Yet to what extent, and in what manner, should ABM researchers work to achieve a more empirically compelling modeling of time and true asynchronicity?

Unfortunately, even the best available theories from physicists (Barbour 2001, 2009; Markopoulou 2000; Rovelli 2008; Smolin 2002) and philosophers (Markosian 2010; Price 1996) provide little definitive insight into the nature of time; like the quantum measurement problem, it continues to be an enigma. Computer scientists appear to be even further behind in their understanding, typically endorsing a pre-relativistic and pre-quantum Newtonian concept of a background time (Herlihy and Shavit, Chapter 2.1, 2008).

This leaves ABM researchers in a conundrum. If the computer platforms available for ABM simulations, as well as the very theoretical foundations on which ABM simulations are built, embody empirically questionable notions of time and simultaneity, what can an ABM researcher do?

This is where the fundamental role of “interactions” in ABM comes into play. Each ABM agent can have an entirely independent sense of time that bears no relation to the flow or passage of time perceived by other agents. Nevertheless, the agents can still experience “change” as a result of the interactions between them, and they can observe an ordering of events (generally unique to each observer). The changes experienced by each agent can be accumulated in the information retained by the agent, and also in the instrumentation provided to observe and record system outcomes.

In particular, running ABMs as genuinely asynchronous interactive systems could permit them to reflect the forms of *self-synchronization* observed in many natural phenomena (Strogatz 2003) and in spatial game situations (Newth and Cornforth 2009). For example, consider an ABM consisting of a collection of fully independent agents in a weakly-asynchronous interaction network. It might be utterly intractable to “synchronously” explore the random sequential execution of agents with anything but a coarse resolution in time: imagine a 100 microsecond primary cycle, and consider how long it would take to explore combinatorial alternatives down to a 1 microsecond or even a 1 nanosecond resolution. However, running the ABM as a genuinely asynchronous simulation might permit self-synchronization to arise as a natural temporal phenomenon rather than as an artifact of a synchronous straight-jacket imposed to simplify implementation at the expense of empirical fi-

delity. Genuinely asynchronous modeling remains an important research topic for ABM researchers (Tosic 2005).

In summary, the explicit and implicit restrictions on timing and updating in the implementation of ABMs should be brought to the fore in the analysis and publication of ABM research results. This would permit peer researchers to examine the robustness of reported outcomes to seemingly inconsequential changes in these restrictions.

4.3 Simple or Complex Modeling of Human Behaviors?

A scientific model should be a condensed representation of some phenomenon: as simple as possible but no simpler. By stripping away unnecessary details, the modeler seeks to identify the salient conditions that enable the model to faithfully reflect the phenomenon under investigation. Stripping away superfluous detail allows the modeler to expose deeper principles driving the observed outcomes.

For descriptive studies, two approaches can be distinguished. The first approach is to start with simple behaviors, adding more complex behaviors only as needed to explain an observed regularity. The second approach is to start with the types of complex behaviors observed in natural settings, field experiments, or human-subject laboratory experiments, and then to progressively simplify the behaviors until the ability to generate an observed regularity is lost.

On the other hand, normative (goal-oriented) studies require different considerations. For such studies, agent behavior becomes an experimental treatment factor rather than a descriptive representation.

For example, suppose the goal is robust system performance. A researcher will then need to examine system performance under a wide range of possible scenarios, including stress-tests involving low-probability events having highly adverse social impacts. Ideally, the agent representation of human behavior under these alternative scenarios should span the full range of possibilities. What if people respond calmly? What if they panic? Moreover, what might be the longer-term effects on system performance if people learn from their experiences and modify their future behaviors?

Alternatively, suppose the goal is successful individual performance within a specified system setting. A researcher will then need to conduct a series of “what if” experiments involving alternative individual strategies, possibly including newly envisioned strategies never before tried in practice. Moreover, these individual strategies might have to involve learning and even learning-to-learn aspects to permit flexible response to the possible behaviors of other system participants.

In view of these considerations, ABM researchers are exploring a wide range of behavioral and learning specifications for the agents included in their models (Tsfatsion 2010e). These specifications take into account the behaviors actually observed in field and lab studies, but they also permit the study of new kinds of behaviors potentially suitable for achieving variously specified goals.

5 Is ABM the Right Mathematics for the Physical Sciences?

The traditional goal of physics has been to explain properties of matter that are relatively free of historical contingency (Farmer et al. 2005). The greatest scientific achievements of physicists have involved issues arising in the very large (gravity and cosmology) and in the very small (quantum level). “Observers” in these two extreme realms measure and assign states to the systems they observe. However, these observers have traditionally been interpreted as external moving spatial frames of reference rather than any form of conscious being localized in some space.

The social sciences occupy a middle ground between these two extreme realms. Social scientists, by definition, must address the problem of human interaction impinging on world events. These interactions leave traces on the human interactors themselves, changing their memories, their knowledge, their future interaction patterns, and their expressed behaviors in these future interactions. In consequence, social systems are intrinsically heterogeneous and path dependent. This has led some physical scientists to question the scientific status of the social “sciences.”

Nevertheless, many physicists have expressed the belief that physics cannot be considered complete until it provides a unified theory that encompasses all of nature, from general relativity to quantum physics. Understanding a natural world such as ours – a world that contains living entities – would require confronting the difficulties posed by having multiple observers imbedded within this world whose measurements are necessarily local and relative to each other and whose interactions can potentially alter its dynamic course.

In particular, physicists recognize the observer-dependent nature of cosmological knowledge (Döring and Isham 2010, Markopoulou 2002, Smolin 2002). For example, Smolin (2002, p. 32) writes:

“...we should not be surprised if both cosmology and social theory point us in the same direction. They are the two sciences that cannot be formulated sensibly unless we build into their foundations the simple fact that all possible observers are inside the system they study.”

Also, recent developments in quantum physics have focused on the modeling of interactions (information exchange) among relative entities rather than the measurement of absolute physical properties by some omniscient observer. In particular, according to the recently developed theory of *relational quantum mechanics (RQM)* there are no observer-independent states (Rovelli 1996). The distinctions between observer and observed, and between cause and effect, are replaced by a symmetric notion of mutual observers interacting with each other and leaving behind their informational footprint as irreversible processes that dissipate energy when information is erased.

RQM thus interprets quantum mechanics to be a theory describing the information one system can have about another, acquired only through interactions, rather than a theory providing a GEV perspective on the physical states of systems (Smerlak 2006, van Fraassen 2010). A further implication of RQM is insight into the non-commutative matrix mechanics characteristic of quantum theory, which implies that

the order of events defines the reality of what is measured. In RQM, the history (or order of events) measured by each observer *creates* the reality that this observer perceives as the evolution of the system.

These developments in physics support a key concept that appears to be mirrored in all of our fundamental laws of physics: there is no preferred frame of reference for any measurement. Rather, all frames are equally valid; and each observer will register a unique view of any measurement, including the order of observed events, from its own relative vantage point.

In addition, the implication of physical theories relying on “symmetry breaking” is that the properties of even very basic physical systems can be determined in part by historical conditions. For example, remarking on the multiple low-energy states of a polymer of amino acids, Laughlin et al. (2000, Section IV) note: “The property of having such a set of low-energy states is connected with the idea of replica symmetry breaking: Different copies of the same system may well fall into different long-lived states through accidents of detailed molecular motion.”

In summary, physics shares with the social sciences an interest in understanding the complicated interactions of entities composed of more elementary entities behaving in accordance with potentially simpler rules. Physics also shares with the social sciences the need to account for multiple observers with different perspectives on reality whose measurements necessarily entail perturbative interactions with persistent (information flow) traces. Finally, physical theories based on broken symmetries suggest that systems of interest to physicists can display what social scientists refer to as “path dependencies,” i.e., dependencies on historical conditions.

These aspects of physical reality are precisely the kind of complexity that ABM tools and ABM design principles have been developed to handle. ABM permits the modeling of asynchronous interactions among heterogeneous hierarchically-nested agents, each operating in terms of its own local coordinate system. Each ABM agent acquires data about its world through interactions with other agents and uses this data to build and maintain a coherent picture of a world centered about itself.

The question then arises whether the modeling problems faced by physicists attempting to derive unified theories of nature are in fact all that different from the modeling problems faced by social scientists. If not, ABM methods now being adopted for social science research might ultimately have a fundamental role to play within a larger scientific realm.

6 Two Illustrative ABM Applications

6.1 Overview

This section briefly describes two ABM applications to illustrate concretely the dual ability of ABM to both simulate and synthesize real-world systems. The first application involves the development and use of an ABM testbed for the perfor-

mance evaluation of restructured wholesale power markets. The second application involves the use of ABM design principles to develop a new type of architecture for the storage and management of information.

6.2 Electricity Market Design

In 2003 the U.S. Federal Energy Regulatory Commission (FERC) recommended the adoption of a common market design for U.S. wholesale power markets (FERC 2003). As indicated in Fig. 2, and elaborated in Joskow (2006), versions of this design have been implemented (or adopted for implementation) in North American energy regions in the Midwest, New England, New York, the Mid-Atlantic States, California, the Southwest, and Texas.

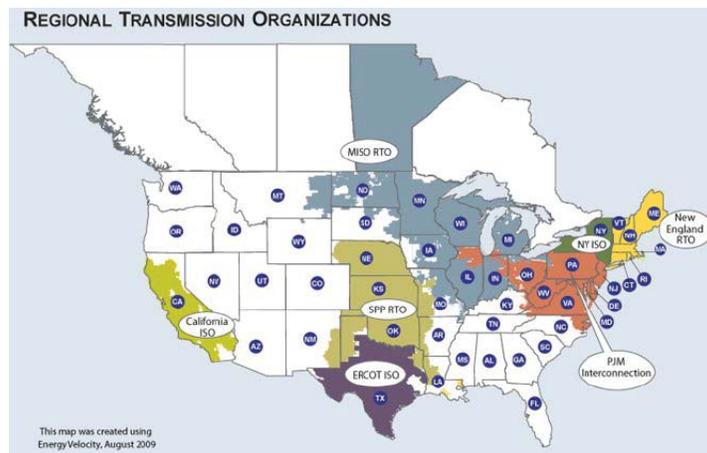


Fig. 2 North American energy regions that have adopted versions of FERC's wholesale power market design. Source: www.ferc.gov/industries/electric/indus-act/rto/rto-map.asp

A core feature of FERC's design is a reliance on *locational marginal pricing (LMP)* to manage transmission grid congestion. Under this pricing system, the price charged to wholesale buyers and received by wholesale sellers at a particular transmission grid bus location at a particular point in time is the least cost to the system of providing an additional increment of power at that bus location at that time.

For the past several years a group of researchers at Iowa State University has been developing and using an agent-based testbed AMES (Agent-based Modeling of Electricity Systems) to explore the performance characteristics of wholesale power markets operating under FERC's design (Tesfatsion 2010f). As indicated in Fig. 3, AMES models strategic trading among a dispersed collection of *Load-Serving Entities (LSEs)* who bid to buy power at wholesale and *Generation Companies (Gen-*

- **Traders**
 - Buyers = Load-Serving Entities (LSEs)
 - Sellers = Generating Companies (*GenCos*)
 - Learning capabilities
- **Independent System Operator (ISO)**
 - Day-ahead hourly scheduling via bid/offer-based DC optimal power flow (OPF)
 - System reliability assessments
- **Two-settlement process**
 - Day-ahead energy market (double auction, financial contracts)
 - Real-time energy market (settlement of differences)
- **AC transmission grid**
 - LSEs and *GenCos* located at various buses across the grid
 - Congestion managed via locational marginal pricing (LMP)

Fig. 3 Core features of FERC's wholesale power market design that have been incorporated into the AMES Wholesale Power Market Testbed.

Cos) who offer to sell power at wholesale. These trading activities take place within a two-settlement (real-time and day-ahead) energy market system administered by a not-for-profit *Independent System Operator (ISO)*, with congestion managed by LMP. The power flows resulting from these trading activities are constrained by the physical characteristics of an underlying AC transmission grid.

AMES experiments have been conducted using multi-period versions of a commonly used 5-bus ISO training case and a standard 30-bus IEEE test case. Key experimental treatment factors have included *GenCo* learning capabilities, the form of *GenCo* supply offers, the price-sensitivity of LSE demand bids, and price caps imposed by the ISO for the mitigation of market power.

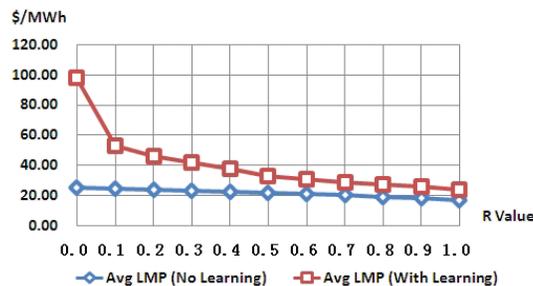


Fig. 4 Average LMP outcomes on day 1000 for the 5-bus test case with and without *GenCo* learning as the price-sensitivity of demand varies from 0% ($R=0.0$) to 100% ($R=1.0$).

One experimental finding has been the relative ease with which the *GenCos* can learn to exercise market power through economic and physical capacity withhold-

ing, even when LSE demand bids are 100% price sensitive (Li et al. 2010). As seen in Fig. 4, this capacity withholding results in a rise in LMP levels that is particularly dramatic for treatments in which the price-sensitivity of demand is low.

Another experimental finding involves “ISO net surplus” collections in day-ahead energy markets (Li and Tesfatsion 2011). *ISO net surplus* is determined each hour as the difference between the LMP payments received by the ISO from energy buyers (the LSEs) and the LMP payments distributed by the ISO to energy sellers (the GenCos). Congestion arising on a transmission grid in any hour necessarily results in separation between the LMPs at two or more transmission grid bus locations, which in turn necessarily results in a non-negative ISO net surplus collection for that hour.

As indicated by the experimental findings reported in Figs. 5 and 6, ISO net surplus and GenCo net earnings are simultaneously enhanced in circumstances *unfavorable* to market efficiency: namely, when demand exhibits low price sensitivity and GenCos have learning capabilities permitting them to exploit this lack of price sensitivity. The truly surprising finding here is just how substantial the net surplus collections of the not-for-profit ISO can be in these circumstances: similar in size to GenCo net earnings, and in some cases even exceeding GenCo net earnings.

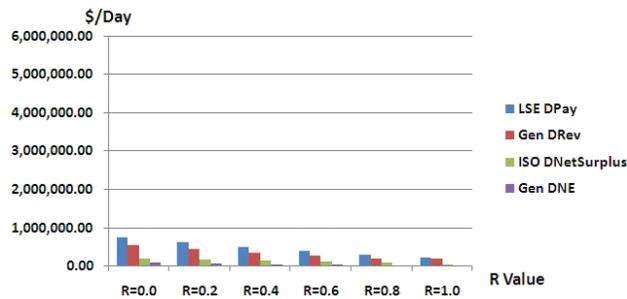


Fig. 5 No-Learning Treatment: LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings on day 1000 for the benchmark (no learning) 5-bus test case as the price-sensitivity of demand varies from 0% (R=0.0) to 100% (R=1.0).

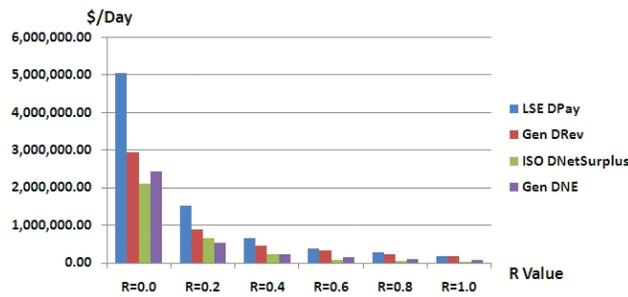


Fig. 6 Learning Treatment: Average LSE payments, GenCo revenues, ISO net surplus, and GenCo net earnings on day 1000 for the 5-bus test case with GenCo learning as the price-sensitivity of demand varies from 0% (R=0.0) to 100% (R=1.0).

Empirical investigations confirm that ISO net surplus collections in actual North American energy regions operating under LMP can indeed be substantial (Li and Tesfatsion 2010). For example, in 2008 these ISO net surplus collections ranged from US\$121 million in New England (ISO-NE) to US\$2.66 billion in the Mid-Atlantic States (PJM).

6.3 Agent-Based Storage and Management of Information

A real-world example where ABM principles enable a dramatic simplification in system design is in the storage and management of information (Borrill 2005).

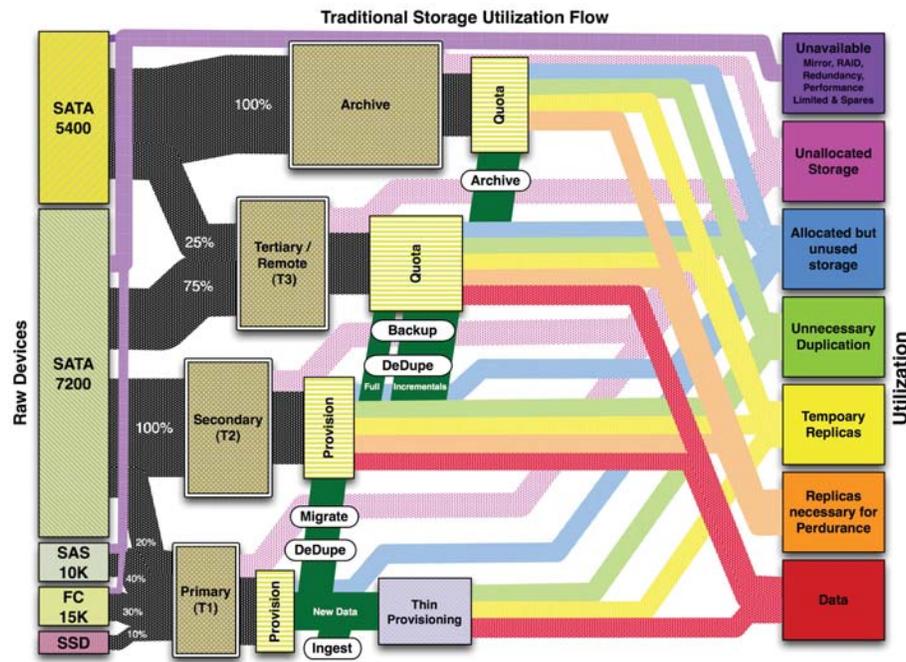


Fig. 7 A traditional enterprise information storage system. Raw resources on the left (storage devices) are utilized by subsystems that manage primary, secondary, tertiary and archival storage by means of various discrete functions that tend to be poorly integrated. Resulting resource utilization is shown on the right.

Enterprise information storage systems have traditionally been designed in a top-down manner by a human architect. This architect receives inputs in the form of system requirements, feasible technologies, and available products, from which a “blueprint” is created and presented to an *information technology (IT)* team for implementation.

Unfortunately we have now gone far beyond the ability of any individual expert or expert team to anticipate the evolving information storage requirements of a dynamic enterprise. As illustrated in Fig. 7, this has led in practice to the need to incrementally overlay existing storage systems with new functions, resulting in unwieldy and inefficient systems with hard-to-predict performance characteristics and brittle response to perturbations (failures, disasters, attacks).

The result is an on-going complexity crisis in the storage and IT industries. Chief information officers are frequently trapped between a rock and a hard place. They are compelled to employ and train an ever-increasing number of IT administrators to manage the overall operation of the many disparate subsystems comprising their information storage systems. Because of the centralized design of these systems, the administrators are naturally seduced into managing such systems from a single “crow’s nest” perspective, a tactic popularly referred to as “all behind one pane of glass.” The similarity between this centralized management design, the “client-server” architectural style of most software engineers, and the God’s-Eye-View (GEV) adopted in traditional systems modeling, is not coincidental.

The resulting price in complexity paid by this GEV design practice shows up when we try to scale these systems. Scaling increases the number, connectivity, and diversity of a system’s elements, and results in cognitive overload for the administrators in the crow’s nest whenever a non-routine failure or error occurs. This is particularly apparent during crises, such as equipment failures, natural disasters, and criminal/terrorist attacks (both physical and cyber).

A common defensive tactic has been to ensure there is only “one throat to choke,” meaning that a single vendor is made responsible for the entire information-storage operation in order to eliminate finger-pointing when things inevitably go wrong. However, this tactic involves major risk because the vendor is then in a monopoly position relative to the enterprise; it can dictate prices for its services and products, and it can take its time when developing needed upgrades, because the enterprise is a captive customer.

What this traditional approach to information storage has done, in effect, is to build human administrators into the systems while at the same time reducing their achievable reliability. Human beings are notoriously error-prone and often induce failure modes by their own inadvertent actions. Human beings are also bandwidth limited; they have limited cognitive ability to respond appropriately to large system perturbations occurring in short periods of time. Traditional information storage systems are thus prone to disaster, and recovery from these disasters – particularly for large systems – has proven to be very complex and time-consuming. Moreover, recovery can involve significant collateral damage to the businesses involved as well as to the local economy as commercial operations are curtailed until these and other IT systems are restored to their intended state of operation.

A primary motivation for considering the redesign of traditional information storage systems has been to find a way to dramatically reduce the time to recover these systems after a disaster. Recognizing that information storage systems inevitably evolve into complex systems encompassing many distributed interacting subsystems, researchers have converged on an alternative architecture based on ABM de-

sign principles to address the many “wicked” problems that occur in such environments (Borrill 2008).

This alternative ABM-based information storage system architecture is schematically depicted in Fig. 8. The architecture operates in a manner similar to a “network automaton” (Smith et al. 2009) in which resources grow and migrate as needed through an evolving network, driven by nearest-neighbor rendezvous and interactions. For the storage system, the individual (autonomous, encapsulated) storage cells represent the deployed units of hardware. Initially these storage cells are identical and substitutable. The system then auto-configures; storage cells form nearest-neighbor relationships with their peers and recursively build (and continue to evolve) the network topology. Given this physical infrastructure, it becomes trivial to deploy storage objects (files) over the infrastructure that are each treated as an agent – that is, as an abstract storage object encapsulating data, metadata, and if-then rules for (meta)data disposition conditional on the occurrence of certain events. Examples of such storage objects include regular files, virtual machine images, databases, and media such as pictures or movies.

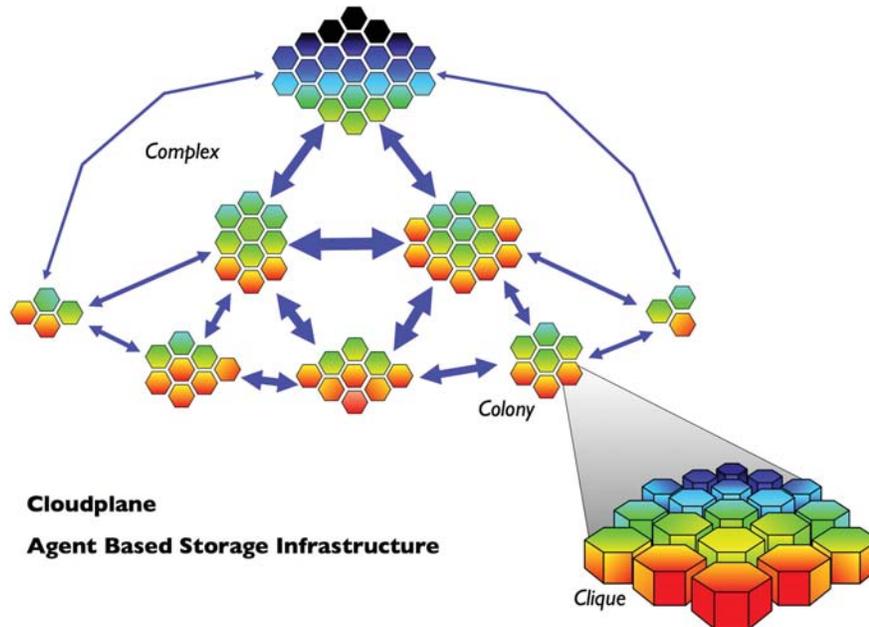


Fig. 8 An enterprise information storage system based on ABM design principles. The resulting system has a substantially simpler architecture yet displays greater robustness to perturbations and greater adaptivity to the evolving needs of geographically dispersed human users. The heatmap expresses one form of emergent behavior, namely, the most-recently-used data remain at the front of the network while the least-recently-used data migrate to the rear of the network.

The agents (storage objects) in the resulting ABM-based information storage system can freely roam, replicate, and evaporate down to some specifiable minimum number able to guarantee data persistence in the face of system perturbations. The relationship of an abstract file agent to its dynamically evolving set of peer replicas enables the ABM to respond in ways that no programmer or administrator would ever have the time, attention, or ingenuity to invent. Moreover, while the underlying methods that drive agent interactions can be exceedingly simple (ingest, push, pull, evaporate), they can also produce rich examples of emergent behaviors. For example, as indicated by the heatmap depiction in Fig. 8, data can migrate automatically across a system's resources (even if geographically distributed) on the basis of latest time of use, resulting in the creation of a self-tiered network.

Another advantage of information storage systems built in accordance with ABM principles is that their relatively simple architecture facilitates performance testing under normal lifecycle-management scenarios as well as under emergency disaster-recovery scenarios. In particular, it is relatively straightforward to copy the real system configuration into an ABM simulator that models the same resources, agents, and network topology (extracted from instrumentation in the real system), using the same rules as the real system.

This simulator can then be used to conduct various "what-if" experiments that would be prohibitively expensive or disruptive to carry out on the real system. For example, the simulator might be used to investigate the direct effects of system disturbances on bandwidth, capacity, and recovery time, and to provide guidance for improving the robustness of the system against such disturbances. It might also be used to improve system efficiency by identifying ways to reduce connectivity or resource usage while at the same time maintaining essentially the same service capabilities at levels visible to human users.

Furthermore, disaster preparedness in personnel training, emergency procedure planning, and other social phenomena under disaster scenarios could easily be built into the simulation. This would permit the testing of personnel responses to unusual events, including their necessary interactions with the data storage system as they initially evacuate their posts and as they later return to these posts to reinstate service to their customers.

7 Conclusion

This article has taken readers on a brief guided tour of ABM. Although applicability for the social sciences has been stressed, we have also touched on intriguing connections with ongoing methodological developments in constructive mathematics, computability theory, cosmology, and quantum physics.

It is perhaps unsettling to realize that the foundations of all of these disciplines are currently in ferment. The uncomfortable truth is that our universe is not the perfectly ordered and deterministic place envisioned by Laplace and Hilbert but the rich, unruly, and largely uncomputable world of Kurt Gödel, Alonzo Church, Alan

Turing, Errett Bishop, and David Deutsch. What a wondrous development it would be if a methodology such as ABM, specifically designed for the study of complex interactive systems, were to enable common light to be shed across these seemingly disparate disciplines.

Acknowledgements

The authors thank Rob Axtell, Julian Barbour, Daniel Goroff, Matheus Grasselli, Fotini Markopoulou, Simone Severini, and Lee Smolin for insightful discussions and kind generosity in responding to questions. The collaboration of Michael Honig and Kenneth Steiglitz on Fig. 1 is also gratefully acknowledged.

References

1. Albert, Réka, and Albert-László Barabási (2002), "Statistical Mechanics of Complex Networks," *Reviews of Modern Physics* 74, January, 47–97.
2. Ambos-Spies, Klaus, and Peter A. Fejer (2006), "Degrees of Unsolvability," Unpublished Manuscript, March.
3. Axelrod, Robert, and Leigh Tesfatsion (2010), *On-Line Guide for Newcomers to Agent-Based Modeling in the Social Sciences*, www.econ.iastate.edu/tesfatsi/abmread.htm.
4. Axtell, Robert L. (2003), "Economics as Distributed Computation," pp. 3–23 in Hiroshi Deguchi, Keiki Takadama and Takao Terano, eds, *Meeting the Challenge of Social Problems via Agent-Based Simulation*, Springer, Tokyo, Japan.
5. Axtell, Robert L. (2001), "Effects of Interaction Topology and Activation Regime in Several Multi-Agent Systems," pp. 33–48 in Scott Moss and Paul Davidsson, eds., *Multi-Agent Based Simulation*, Lecture Notes in Computer Science, Volume 1979/2001, Springer-Verlag, Berlin/Heidelberg.
6. Barbour, Julian (2001), *The End of Time: The Next Revolution in Physics*, Oxford University Press, Oxford, UK.
7. Barbour, Julian (2009), *The Nature of Time*, March. <http://arxiv.org/abs/0903.3489>
8. Barr, Jason M., Troy Tassier, Leanne J. Ussher, Blake LeBaron, Shu-Heng Chen, and Shyam Sunder (2008), "The Future of Agent-Based Research in Economics: A Panel Discussion," *Eastern Economic Journal* 34, 550–565.
9. Batty, Michael (2005), *Cities and Complexity: Understanding Cities with Cellular Automata, Agent-Based Models, and Fractals*, MIT Press, Cambridge, MA.
10. Binder, Philippe (2008), "Theories of Almost Everything," *Nature* 455, 884–885.
11. Borrill, Paul (2005), "Autopoietic File Systems: From Architectural Theory to Practical Implications," Vanguard Conference, San Francisco, January.
12. Borrill, Paul (2008), "Smart Data and Wicked Problems," Vanguard Conference, February, <http://www.replicus.com/REPLICUS/Resources.html>
13. Bridges, Douglas (2009), "Constructive Mathematics", *The Stanford Encyclopedia of Philosophy* (Summer 2009 Edition), Edward N. Zalta (ed.), plato.stanford.edu/archives/sum2009/entries/mathematics-constructive
14. Bridges, Douglas (1999), "Constructive Mathematics: A Foundation for Computable Analysis," *Theoretical Computer Science* 219, 95–109.
15. Chang, Myong-Hun, and Joseph E. Harrington, Jr. (2006), "Agent-Based Models of Organizations," Chapter 26 (pp. 1273–1337) in Leigh Tesfatsion and Kenneth L. Judd, eds, *op. cit.*

16. Chen, Shu-Heng, and Bin-Tzong Chie (2007), "Modularity, Product Innovation, and Consumer Satisfaction: An Agent-Based Approach," pp. 1053-1062 in Hujun Yin, Peter Tino, Emilio Corchado, Will Bryne, and Xin Yao, eds., *Intelligent Data Engineering and Automated Learning – IDEAL 2007*, Springer-Verlag, Berlin/Heidelberg.
17. Cioppa, Thomas M., Thomas W. Lucas, and Susan M. Sanchez (2004), "Military Applications of Agent-Based Simulations," *Proceedings of the 36th Winter Simulation Conference*, New Orleans, Louisiana, 171–180.
18. Dalen, Dirk (2008), *Logic and Structure*, 4th Edition, Springer-Verlag, Berlin/Heidelberg.
19. Darema, Frederica (2005), "Grid Computing and Beyond: The Context of Dynamic Data Driven Applications Systems," *Proceedings of the IEEE: Special Issue on Grid Computing* 93(3), 692–697.
20. Dawid, Herbert (2006), "Agent-Based Models of Innovation and Technological Change," Chapter 25 (pp. 1235–1272) in Leigh Tesfatsion and Kenneth L. Judd, eds, *op. cit.*
21. Devaney, Robert L., and Linda Keen, eds., (1989), *Chaos and Fractals: The Mathematics Behind the Computer Graphics*, American Mathematical Society, Providence, RI.
22. Döring, Andreas, and Chris J. Isham (2010), "'What is a Thing?': Topos Theory in the Foundations of Physics," in Robert Coecke, ed., *New Structures for Physics*, Lecture Notes in Physics, Vol. 813, Springer, 753–937.
23. Eberbach, Eugene, Dina Goldin, and Peter Wegner (2004), "Turing's Ideas and Models of Computation," pp. 159–194, in Christof Teuscher, ed., *Alan Turing: Life and Legacy of a Great Thinker*, Springer-Verlag, Berlin/Heidelberg.
24. Epstein, Joshua M. (2006), *Generative Social Science: Studies in Agent-Based Computational Modeling*, Princeton University Press, Princeton, NJ.
25. Farmer, J. Doyne, Martin Shubik, and Eric Smith (2005), "Economics: The Next Physical Science?" *Physics Today*, September, 37–42.
26. FERC (2003), *Notice of White Paper*, U.S. Federal Energy Regulatory Commission, Washington, D.C., April.
27. Gilbert, Nigel (2007), *Agent-Based Models*, Volume 153, Series: Quantitative Applications in the Social Sciences, SAGE Publications, Inc., London, UK.
28. Gilbert, Nigel, Andreas Pyka, and Petra Ahrweiler (2001), "Innovation Networks - A Simulation Approach," *Journal of Artificial Societies and Social Simulation* 4(3). www.soc.surrey.ac.uk/JASSS/4/3/8.html
29. Goldin, Dina Q., Scott A. Smolka, and Peter Wegner, eds, (2006), *Interactive Computation: The New Paradigm*, Springer-Verlag, Berlin/Heidelberg.
30. Grimm, Volker, and Steven R. Railsback (2005), *Individual-Based Modeling and Ecology*, Princeton University Press, Princeton, NJ.
31. Herlihy, Maurice, and Nir Shavit (2008), *The Art of Multiprocessor Programming*, Morgan Kaufmann, Burlington, MA.
32. Huang, Jun, Nick R. Jennings, and John Fox (1995), "An Agent-Based Approach to Health Care Management," *Applied Artificial Intelligence* 9(4), 401–420.
33. Jennings, Nicholas R. (2000), "On Agent-Based Software Engineering," *Artificial Intelligence* 117, 277–296.
34. Joskow, Paul (2006), "Markets for power in the United States: An interim assessment," *The Energy Journal* 27(1), 1–36.
35. Laubenbacher, Reinhard, Abdul S. Jarrah, Henning Mortveit, and S.S. Ravi (2009), "A Mathematical Formalism for Agent-Based Modeling," pp. 160–176 in Robert A. Meyers, ed., *Encyclopedia of Complexity and Systems Science*, Springer-Verlag, Berlin/Heidelberg.
36. Laughlin, Robert B., and David Pines (2000), "The Theory of Everything," *Proceedings of the National Academy of Sciences, U.S.A.* 97(1), 28–31. <http://www.pnas.org/content/97/1/28.abstract>
37. Laughlin, Robert B., David Pines, Joerg Schmalian, Branko P. Stojković, and Peter Wolynes (2000), "The Middle Way," *Proceedings of the National Academy of Sciences, U.S.A.* 97(1), 32–37.

38. LeBaron, Blake, and Leigh Tesfatsion (2008), "Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents," *American Economic Review Papers & Proceedings*, 98(2), 246–250.
39. Li, Hongyan, Junjie Sun, and Leigh Tesfatsion (2010), "Testing Institutional Arrangements via Agent-Based Modeling: A U.S. Electricity Market Application", in Herbert Dawid and Willi Semmler, eds, *Computational Methods in Economic Dynamics*, Springer-Verlag, Berlin/Heidelberg.
40. Li, Hongyan, and Leigh Tesfatsion (2011), "ISO Net Surplus Collection and Allocation in Wholesale Power Markets under Locational Marginal Pricing," *IEEE Transactions on Power Systems* 26(2), 627–641.
41. MacLane, Saunders (1998), *Categories for the Working Mathematician*, Second Edition, Springer-Verlag, New York, NY.
42. Macy, Michael W., and Robert Willer (2002), "From Factors to Actors: Computational Sociology and Agent-Based Modeling," *Annual Review of Sociology* 28, 143–166.
43. Markopoulou, Fotini (2000), "The Internal Description of a Causal Set: What the Universe Looks Like from the Inside," *Communications in Mathematical Physics* 211(3), 559–583. <http://arxiv.org/abs/gr-qc/9811053>.
44. Markosian, Ned (2010), *Time*, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/time/>
45. Nagel, Kai, and Peter Wagner (2006), *Traffic Flow: Approaches to Modelling and Control*, John Wiley & Sons, New York, NY.
46. Newth, David, and David Cornforth (2009), "Asynchronous Spatial Evolutionary Games," *BioSystems* 95(2), 120–129. <http://www.ncbi.nlm.nih.gov/pubmed/18926874>
47. Parker, Dawn C., Steven M. Manson, Marco A. Janssen, Matthew J. Hoffman, and Peter Deadman (2003), "Multi-Agent Systems for the Simulation of Land-Use and Land-Cover Change: A Review," *Annals of the Association of American Geographers* 93(2), 314–337.
48. Prasse, Michael, and Peter Rittgen (1998), "Why Church's Thesis Still Holds: Some Notes on Peter Wegner's Tracts on Interaction and Computability," *The Computer Journal* 41, 357–362.
49. Price, Hew (1996), *Time's Arrow & Archimedes' Point*, Oxford University Press, Oxford, UK. <http://www.usyd.edu.au/time/price/TAAP.html>
50. Rovelli, Carlo (1996), "Relational Quantum Mechanics," *International Journal of Theoretical Physics* 35, 1637–1678.
51. Rovelli, Carlo (2008), "Forget Time." <http://www.fqxi.org/community/forum/topic/237>
52. Smerlak, Matteo (2006), "Relational Quantum Mechanics," Internship Report, École normale supérieure de Lyon, September 17.
53. Smith, Adam (1776), *An Inquiry into the Nature and Causes of the Wealth of Nations*, reprinted in 1976 by William B. Todd, ed., Glasgow Edition of the Works and Correspondence of Adam Smith, Volume 1, Oxford University Press, Oxford, UK.
54. Smith, David M.D., Jukka-Pekka Onnela, Chiu Fan Lee, Mark Fricker, and Neil F. Johnson (2009), "Network Automata: Coupling Structure and Function in Real-World Networks," [arXiv:physics/0701307v3](http://arxiv.org/abs/physics/0701307v3)[physics.soc-ph]
55. Smolin, Lee (2002), *Three Roads to Quantum Gravity*, Basic Books, New York, NY.
56. Soare, Robert I. (1996), "Computability and Recursion," *Bulletin of Symbolic Logic* 2, 284–321.
57. Strogatz, Steven (2003), *Sync: The Emerging Science of Spontaneous Order*, 1st Edition, Hyperion, New York, NY.
58. Tesfatsion, Leigh, and Kenneth L. Judd, eds, (2006), *Handbook of Computational Economics: Volume 2, Agent-Based Computational Economics*, Handbook in Economics Series, North-Holland/Elsevier, Amsterdam, The Netherlands.
59. Tesfatsion, Leigh (2010a), *Agent-Based Computational Economics (ACE) Homepage*, www.econ.iastate.edu/tesfatsi/ace.htm
60. Tesfatsion, Leigh (2010b), *ACE/CAS General Software and Toolkits*, www.econ.iastate.edu/tesfatsi/acecode.htm
61. Tesfatsion, Leigh (2010c), *Verification and Empirical Validation of Agent-Based Computational Models*, www.econ.iastate.edu/tesfatsi/EmpValid.htm

62. Tesfatsion, Leigh (2010d), *Agent-Based Computational Economics: Key Application Areas*, www.econ.iastate.edu/tesfatsi/aapplic.htm
63. Tesfatsion, Leigh (2010e), *Agent-Based Computational Economics Research Area: Learning and the Embodied Mind*, <http://www.econ.iastate.edu/tesfatsi/aemind.htm>
64. Tesfatsion, Leigh (2010f), *AMES Wholesale Power Market Testbed Homepage*, www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm
65. Tadic, Predrag T. (2005), "Cellular Automata for Distributed Computing: Models of Agent Interaction and Their Implications," *IEEE Transactions on Systems, Man and Cybernetics* 4, 3204–3209.
66. van Fraassen, Bas C. (2010), "Rovelli's World," *Foundations of Physics* 40(4), 390–417.
67. Velupillai, K. Vela (2010), *Computable Foundations for Economics*, Routledge, New York, NY.
68. von Hayek, Friedrich A. (1948), *Individualism and Economic Order*, University of Chicago Press, Chicago, IL.
69. von Neumann, John (1951), "The General and Logical Theory of Automata," pp. 1–41 in Lloyd A. Jeffress, ed., *Cerebral Mechanisms in Behavior-The Hixon Symposium*, John Wiley & Sons, New York, NY.
70. Wolfram, Stephen (2002), *A New Kind of Science*, Wolfram Media, Inc., Champaign, Illinois.
71. Wolpert, David H. (2008), "Physical Limits of Inference," *Physica D* 237, 1257–1281.