Short-Term Opportunities, Medium-Run Bottlenecks, and Long-Time Barriers to Progress in the Evolution of an Agent-Based Social Science

Chapter Sixteen

ROBERT L. AXTELL, GEORGE MASON UNIVERSITY

This chapter is extracted from *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, by Aaron B. Frank and Elizabeth M. Bartels, eds., RR-A1275-1, 2022 (available at www.rand.org/t/RRA1275-1).

Prepared for the Defense Advanced Research Projects Agency, Defense Sciences Office
Approved for public release; distribution unlimited

RAND
CORPORATION

For more information on this publication, visit **www.rand.org/t/RRA1275-1**.

Published by the RAND Corporation, Santa Monica, Calif.

CHAPTER SIXTEEN

Short-Term Opportunities, Medium-Run Bottlenecks, and Long-Time Barriers to Progress in the Evolution of an Agent-Based Social Science

Robert L. Axtell, George Mason University

In a prescient RAND Corporation working paper in the early 1950s, the late John Nash investigated the design of parallel computing machines.[1] This was prior to the widespread adoption of what has come to be known as the *von Neumann architecture*, which separates data and program storage from arithmetic execution in essentially all modern computers. Already in 1954, Nash saw that there were reasons why parallel execution models were important. He recognized that, because serial processing led to bottlenecks, it would be hard for non-parallel devices to be used to model the wide variety of scientifically interesting phenomena that manifest themselves in the real world. Although Nash was far ahead of his time in these investigations, much as he had pushed the research frontier forward in the theory of games a few years earlier, his ideas fell on somewhat deaf ears. The imperative at that time was to get digital computing off the ground for weather prediction and bomb-building;[2] little concern was given to conforming with the way that nature solves problems in parallel. At that time, the digital computer was thought of as a giant calculator—a tool for solving otherwise intractable equations, a great engine for ridding the sciences of tedious manual calculations, many

[1] John Nash, *Parallel Control*, Santa Monica, Calif.: RAND Corporation, RM-1361, 1954. John Forbes Nash, Jr., born in 1928 in Bluefield, West Virginia, held an undergraduate degree in mathematics from the Carnegie Institute of Technology and a Ph.D. in mathematics from Princeton University, where he did pioneering work in game theory, developing the main solution concept in that field that is named after him. He was also a former Massachusetts Institute of Technology faculty member and a former RAND staff member. He died in a car accident in 2015. His life is described in Sylvia Nasar, *A Beautiful Mind: A Biography of John Forbes Nash, Jr., Winner of the Nobel Prize in Economics, 1994*, New York: Simon & Schuster, 1998 (subsequently made into a movie).

[2] Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America*, Cambridge, Mass.: MIT Press, 1996.

of which were performed by human "calculators." Indeed, the first general-purpose computer language used to automate the process of generating assembly-level code to run on mainframe computers was FORTRAN (for "FORmula TRANslation"), a relatively simple language for translating formulas from mathematics and science textbooks into working code. Whether solving the equations of fluid mechanics for weather prediction or of quantum mechanics for fissioning the atom, early digital computers and their software were designed for converting classical and well-understood mathematical equations into computer code.

Now, almost 70 years later, there are many more ways to use digital computation than merely to solve equations. In this chapter, I discuss the use of Agent-Based Modeling (ABM) in the social sciences, focusing on the methodology for creating such models to faithfully represent actual social phenomena, grounded in empirical data, and not "toy" models used either pedagogically or as abstract thought experiments. Specifically, I critique the widespread use of so-called single-threaded execution for agent-based models and discuss several of its weaknesses. I then go through some of the many ways that multiple execution threads may be used in agent-based models, both for performance gains and for increased verisimilitude. With the growing availability of high-performance computing, the notion of large-scale agent-based models is becoming increasingly relevant. The related idea of full-scale agent-based models is also mentioned in this chapter, and its growing use in the foreseeable future is explored. I also attempt to illuminate bottlenecks that lie on the horizon, so far as they are known, and larger technical challenges that must be surmounted before any vision of "mirror world" or "digital twin" agent-based models is realizable.

Unlike most research, this chapter does not present specific results or describe particular models in any detail. It does discuss a modeling approach, but not in sufficient detail to be pedagogical. Rather, the goal here is to take a broad view of a relatively young and rapidly developing field and attempt to come to terms with its history, its early promise and potential, its recent progress, and its trajectory, both in the short term and over longer time horizons. In doing this, I am keenly aware that some of the claims made may be viewed as rather speculative, far from the world of mathematical theorems, statistical laws, and precise computational code that I normally inhabit. In rendering what are essentially judgments on the state of the art and editorializing about how I see things developing—where the problems, both big and small, are likely to lie and how they might be surmounted—I draw on nearly 30 years of experience with this novel methodology. After leaving graduate school in 1992, I was "present at the creation" in the early heyday of ABM, accomplished as it was on modest hardware and in what are by today's standards rather archaic programming environments. Even then, the fertility of the approach was apparent to many in the social sciences,[3] the computer sciences,[4] and ecology and other

[3] For instance, see W. Brian Arthur, "Inductive Reasoning and Bounded Rationality," *American Economic Review*, Vol. 84, No. 2, May 1994.

[4] For example, see John H. Holland and John H. Miller, "Artificial Adaptive Agents in Economic Theory," *American Economic Review*, Vol. 81, No. 2, May 1991.

nearby branches of biology,[5] including the then-new field of artificial life.[6] Much of this early work is associated with the Santa Fe Institute, once a hotbed for developments in this area. The loci of ABM development migrated long ago to universities (e.g., Oxford), research labs (e.g., Los Alamos, Argonne), and think tanks (e.g., Brookings); nascent commercial efforts are now underway as well (e.g., AnyLogic). The depth of activity in this broad area means that some version of ABM technology will be with us for a long time, especially in domains in which the role of *interaction* is critical for understanding system behavior overall. This is certainly true in many of the social sciences and appears increasingly true in several of the natural sciences,[7] and ABM is a methodology also of interest in the emerging field of *digital humanities*,[8] with many new approaches and exciting applications appearing in recent years. With continued community development of ABM, it is sure to take its place beside other important tools and techniques, such as statistics, econometrics, cognitive science, and neuroscience, on the road to greater scientific understanding of human social and behavioral processes.

Agent-Based Modeling as a Tool for Relaxing Unrealistic Assumptions in the Social Sciences

Many scientists have noted that, in certain ways, the social sciences are harder (i.e., more difficult) than the natural sciences.[9] Historically, controlled experiments were uncommon in both the behavioral and social sciences—a situation that has changed.[10] Social scientists typically had little high-quality data with which to develop and test theories, which has also changed, with the increasing availability of administrative data in digital form.

---

5 P. Hogeweg and B. Hesper, "Individual-Oriented Modelling in Ecology," *Mathematical and Computer Modelling*, Vol. 13, No. 6, 1990.

6 Rodney A. Brooks and Pattie Maes, eds., *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, Mass.: MIT Press, 1994; Christopher G. Langton, ed., *Santa Fe Institute Studies in the Sciences of Complexity*, Vol. VI, *Artificial Life*, Reading, Mass.: Addison-Wesley Publishing, 1989; Christopher G. Langton, Charles Taylor, J. Doyne Farmer, and Steen Rasmussen, *Santa Fe Institute Studies in the Sciences of Complexity*, Vol. X, *Artificial Life II*, Redwood City, Calif.: Addison-Wesley Publishing, 1992; and Christopher G. Langton, ed., *Santa Fe Institute Studies in the Sciences of Complexity*, Vol. XVII, *Artificial Life III*, Reading, Mass.: Addison-Wesley Publishing, 1994.

7 For example, see R. B. Laughlin, David Pines, Joerg Schmalian, Branko P. Stojkovic, and Peter Wolynes, "The Middle Way," *Proceedings of the National Academy of Sciences*, Vol. 97, No. 1, January 4, 2000.

8 Roy Rosenzweig, "Scarcity or Abundance? Preserving the Past in a Digital Era," *American Historical Review*, Vol. 108, No. 3, June 1, 2003.

9 For examples, see Paul Krugman, *Peddling Prosperity*: *Economic Sense and Nonsense in the Age of Diminished Expectations*, New York: W. W. Norton & Company, 1994; Charles A. Lave and James G. March, *An Introduction to Models in the Social Sciences*, New York: Harper & Row, 1975; and Herbert A. Simon, "Giving the Soft Sciences a Hard Sell," *Boston Globe*, May 3, 1987.

10 Vernon L. Smith, "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, Vol. 70, No. 2, April 1962.

In an early statement of the feasibility and usefulness of doing social science with computational agents,[11] it was suggested that ABM offered significant advantages over conventional mathematical and statistical approaches along at least four dimensions: (1) representing agent heterogeneity, (2) relaxing rationality specifications, (3) mediating interactions through social networks, and (4) permitting investigation of nonequilibrium dynamics. Whole essays have been written on each of these topics, so I will not belabor the issues in this chapter. Of course, models in the social sciences have many more than four dimensions, even if these four are foundational. Another 18 are displayed in Table 16.1, in which the normal representations of these dimensions that are typically taught to graduate students are stated (center column) alongside different, emerging perspectives that can be readily incorporated into agent-based models (right column). This table and variants of it have been described in detail elsewhere.[12] What I focus on here is that progress in the social sciences involves moving from the center column to the right, how this is being worked on by non-ABM approaches, how substantial barriers appear to block significant progress, and how ABM can facilitate the transition.

Scientific research is a kind of exploration, similar to expeditions to remote parts of the earth's surface, such as those undertaken in the past millennium. Both ventures involve life on the frontier—the boundary between what is known and what can only be guessed about—with different explorers holding distinct ideas about the most fertile directions in which to proceed. Some believe that the same tools and methods that have worked well to reach the frontier are the best approach for continued progress, while others point out the inevitable weaknesses of existing technologies and perspectives, the reasons why progress has paused at that location on the frontier, and the need for better techniques for assailing the new heights that apparently block further progress and need to be overcome.

This tension between progress through business-as-usual methods and progress through innovative approaches manifests in various ways. Kuhn distinguished *normal* science from *revolutionary* science as one such dichotomy.[13] Normal science takes received conceptions of the world and existing tools as given and elaborates on or deepens conventional understanding. Revolutionary science rejects one or more concepts that are foundational to the reigning view and comes to conclusions that are at odds with that view. Heliocentric astronomy and quantum mechanics are poster children for the latter. In one increasingly well-accepted

---

[11] Joshua M. Epstein and Robert Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, Cambridge, Mass.: MIT Press, 1996.

[12] Robert L. Axtell, "Hayek Enriched by Complexity Enriched by Hayek," *Advances in Austrian Economics*, Vol. 21, November 30, 2016b; Robert L. Axtell and J. Doyne Farmer, "Agent-Based Modeling in Economics and Finance: Past, Present, and Future," *Complexity Handbook*, Stuttgart, Germany: University of Hohenheim, August 1, 2018; Robert L. Axtell, Alan Kirman, Iain D. Couzin, Daniel Fricke, Thorsten Hens, Michael E. Hochberg, John E. Mayfield, Peter Schuster, and Rajiv Sethi, "Challenges of Integrating Complexity and Evolution into Economics," in David S. Wilson and Alan Kirman, eds., *Complexity and Evolution: Toward a New Synthesis for Economics*, Cambridge, Mass.: MIT Press, 2016.

[13] Thomas S. Kuhn, *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought*, New York: Vintage Books, 1957.

TABLE 16.1

Comparison of Conventional and Complex, Evolutionary Approaches to Model-Building in the Social Sciences

Model Feature	Conventional Conception	Complex, Evolutionary Approach
Number of agents	Representative (1, 2, *N*, infinite)	Many (preferably full-scale)
Diversity of agents	Homogeneous, a few types	Heterogeneous, idiosyncratic agents
Agent goals, objectives	Static, scalar-valued utility	Evolving, other-regarding
Agent behavior	Rational, maximizing	Purposive, adaptive, biased, myopic
Learning	Individual, fictitious play	Grounded in behavioral science or derived from artificial intelligence
Information	Centralized, possibly uncertain	Distributed, tacit
Beliefs	Coordinated for free	Uncoordinated, costly to coordinate
Interaction topology	Equal probability, well mixed	Social networks, fixed or changing
Markets	Auctioneer, global price vector	Decentralized, local prices
Firms and institutions	Unitary actors, production functions	Multiagent groups and organizations
Selection operators	Single level	Multilevel, group selection
Governance	Benevolent planner, median voter	Self-governance, incentive problems
Temporal structure	Static, impulse tests, one-shot	Dynamic, full transient paths
Source of dynamism	Exogenous, outside forces	Endogenous forces
Properties of dynamics	Smooth, differentiable	Irregular, volatile, heavy-tailed
Character of dynamics	Markovian, path is forgotten	Path-dependent, history matters
Solution concepts	Equilibrium at agent level	Macro steady state (stationarity)
Multilevel character	Neglected, dual fallacies	Intrinsic, macrolevel emerges
Methodology	Deductive, mathematical	Abductive, computational
Ontology	Representative agent, *max U*	Ecology of interacting agents
Data	Samples, aggregate	Microdata, big data
Policy stance	Designed from the top down	Evolved from the bottom up

NOTE: *max U* = maximum utility; *N* = population size.

perspective on the philosophy of science, models *mediate* between theory and data.[14] Social scientists build models using the vocabulary in the leftmost column of Table 16.1. The normal science, or the hard core,[15] taught to graduate students,[16] is in the middle column. These specifications are sometimes justified on the grounds of mathematical tractability—that relaxing such specifications, whether in the direction of the third column or otherwise, makes the resulting models impossible or at least much more difficult to solve.

One way to interpret the third column is as desiderata for more-realistic kinds of social science models. Some of the entries have a basis in experiment (e.g., other-regarding preferences).[17] Others are simply less stylized than the conventional specification (e.g., social networks in lieu of equal probability of interaction). Although it may be possible to use conventional mathematical or statistical methods to relax one of the standard conceptions holding the others fixed—as is essentially the case in Jackson's work on social networks in economics[18]—how to move from the center column to the right column, in general, using normal approaches, is not understood.

The potentially revolutionary power of ABM is that it is now possible to move from the center column to the right using computational agents. For essentially each row of Table 16.1, there are examples of agent-based models that relax the normal specification in the manner described there. I will not go through the myriad examples;[19] instead, I will simply point out that this ability makes ABM a game-changing technology for the social sciences if it can be realized to create better models of actual social phenomena, grounded in both the behavior of people and the emergent properties of institutions and social aggregates. If the current ABM method had come along a generation earlier, before the widespread availability of big data, it might have developed into a fertile tool for creating plausible, theoretical, toy models of human social reality. But its maturation today, just as large-scale data collection and dissemination has become commonplace, represents a truly unusual situation, perhaps a once-in-a-lifetime innovation in methodology that, driven by advances in computing machinery, can transform the social sciences.

---

14 Mary S. Morgan and Margaret Morrison, *Models as Mediators: Perspectives on Natural and Social Science*, New York: Cambridge University Press, 1999.

15 Imre Lakatos, "Falsification and the Methodology of Scientific Research Programmes," in Imre Lakatos and Alan Musgrave, eds., *Criticism and the Growth of Knowledge*, New York: Cambridge University Press, 1970.

16 For an example in economics, see Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory*, New York: Oxford University Press, 1995.

17 David J. Cooper and John H. Kagel, "Other-Regarding Preferences: A Selective Survey of Experimental Results," in John H. Kagel and Alvin E. Roth, *The Handbook of Experimental Economics*, Vol. 2, Princeton, N.J.: Princeton University Press, 2015.

18 Matthew O. Jackson, *Social and Economic Networks*, Princeton, N.J.: Princeton University Press, 2008.

19 For an overview article in economics, see Axtell and Farmer, 2018.

The Growing Availability of Data and the Evolution of Agent-Based Models Toward Larger Scales

In a famous short story, the Argentine writer Jorge Luis Borges described a land in which people had developed the ability to create maps with so much detail as to be completely faithful representations of the real world.[20] But it turned out that such maps were essentially useless because they were as comprehensive as the real world. The idea of a map is to abstract the most-useful or most-interesting features of the world to build a representation of it that emphasizes the parts that are the most relevant to some particular purpose. So-called mirror worlds,[21] or digital twins,[22] might be considered the modern, high-tech realizations of Borges's fictional maps—high-fidelity representations of actual places, often containing an amount of detail limited only by the time and energy of the creators. (That is, however much specificity is present in a digital twin, it is almost always possible to add more if it is useful to do so.) Although some might find Borges's critique relevant to digital twins, there is one enormous difference between a full-scale map and a high-fidelity computer model. The latter typically represents the evolution of its twin over time, and this process can typically be run much faster than real time, perhaps a thousand or a million times faster, so such models can be extremely useful in developing an understanding of the real world mapped at large scales, whether for policy or other purposes.

Regardless of one's position vis-à-vis large-scale models, it is not possible to build high-verisimilitude representations of the real world without detailed data about it. Shorn of actual data, such models might be relevant to social life somewhere but are unlikely to be relevant to humans on earth in the 21st century. The happy situation in which we find ourselves today is that, as a consequence of the information technology revolution—i.e., the widespread (although not universal) digitization of administrative records; the online gathering and archiving of user data by such companies as Google, Facebook, and Twitter; and so on—incredible amounts of data that are relevant to social and behavioral science have become available over the past two decades, with much more likely to become available in the

[20] Jorge Luis Borges, *Labyrinths*, New York: New Directions, 1962.

[21] David Gelernter, *Mirror Worlds, or: The Day Software Puts the Universe in a Shoebox . . . How It Will Happen and What It Will Mean*, New York: Oxford University Press, 1992.

[22] See Michael Grieves and John Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in Franz-Josef Kahlen, Shannon Flumerfelt, and Anabela Alves, eds., *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*, Switzerland: Springer Nature, 2017; and, in this report, Chapters Fourteen and Fifteen (Ben Connable, "Authentically Describing and Forecasting Human Behavior for Policy Analysis: A Review and a Path Forward," in Aaron B. Frank and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, Santa Monica, Calif.: RAND Corporation, RR-A1275-1, 2022; and Zev Winkelman, "Using Technology to Improve the Agility of Force Generation Processes," in Aaron B. Frank and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, Santa Monica, Calif.: RAND Corporation, RR-A1275-1, 2022).

coming decade. Across both the social and natural sciences, some have called this the era of big data,[23] although not without naysayers.[24]

For an example of big data derived from administrative sources, consider government tax records on businesses. With tax forms filed at least annually—and more frequently for large companies—such records represent detailed descriptions of at least the basic financial aspects of firm operations, often with links to nonfinancial characteristics, such as employee records that can be further linked to individual tax returns or other administrative data, such as retirement accounts. Such tax records are primarily self-reports; only the largest firms are audited by independent accounting firms, and only a tiny fraction are reviewed by government tax authorities. It is also the case that these records are subject to various strategic behaviors, such as certain kinds of costs being reported in particular years to swell losses or profits according to the vagaries of business cycles and other macroeconomic conditions. Yet such information has proven invaluable to researchers in developing a quantitative understanding of cross-sectional characteristics of businesses,[25] how they interact with other firms,[26] and so on.

As a specific example of such data, consider firm sizes, whether measured by employees or by receipts (i.e., revenue). Data on the sizes of U.S. companies from tax filings have been available for a couple of decades. Early work with such data revealed firm sizes to be extremely skewed *over the entire range of sizes*, from businesses with a single worker up to Walmart, with more than 1 million employees.[27] When such analyses were repeated for other countries, comparable results were obtained,[28] suggesting that a gross empirical regularity existed that was

[23] Susan Athey, "Beyond Prediction: Using Big Data for Policy Problems," *Science*, Vol. 355, No. 6324, February 3, 2017; and Liran Einav and Jonathan Levin, "Economics in the Age of Big Data," *Science*, Vol. 346, No. 6210, November 7, 2014.

[24] Edward Tenner, *The Efficiency Paradox: What Big Data Can't Do*, New York: Alfred A. Knopf, 2018.

[25] On firm growth rates, see, for example, Michael H. R. Stanley, Luís A. N. Amaral, Sergey V. Buldyrev, Shlomo Havlin, Heiko Leschhorn, Philipp Maass, Michael A. Salinger, and H. Eugene Stanley, "Scaling Behaviour in the Growth of Companies," *Nature*, Vol. 379, No. 29, February 1996.

[26] On buyer-supplier networks, for example, see Enghin Atalay, Ali Hortaçsu, James Roberts, and Chad Syverson, "Network Structure of Production," *Proceedings of the National Academy of Sciences*, Vol. 108, No. 13, 2011.

[27] Robert L. Axtell, "Zipf Distribution of U.S. Firm Sizes," *Science*, Vol. 293, No. 5536, September 7, 2001b.

[28] For examples, see Yoshi Fujiwara, Corrado Di Guilmi, Hideaki Aoyama, Mauro Gallegati, and Wataru Souma, "Do Pareto-Zipf and Gibrat Laws Hold True? An Analysis with European Firms," *Physica A: Statistical Mechanics and Its Applications*, Vol. 335, No. 1, April 1, 2004; Eduardo Gaffeo, Mauro Gallegati, and Antonio Palestrini, "On the Size Distribution of Firms: Additional Evidence from the G7 Countries," *Physica A: Statistical Mechanics and Its Applications*, Vol. 324, No. 1, June 1, 2003; Jinzhong Guo, Qi Xu, Qinghua Chen, and Yougui Wang, "Firm Size Distribution and Mobility of the Top 500 Firms in China, the United States and the World," *Physica A: Statistical Mechanics and Its Applications*, Vol. 392, No. 13, July 1, 2013; Sang Hoon Kang, Zhuhua Jiang, Chongcheul Cheong, and Seong-Min Yoon, "Changes of Firm Size Distribution: The Case of Korea," *Physica A: Statistical Mechanics and Its Applications*, Vol. 390, No. 2, January 15, 2011; and Jianhua Zhang, Qinghua Chen, and Yougui Wang, "Zipf Distribution in Top Chinese Firms and an Economic Explanation," *Physica A: Statistical Mechanics and Its Applications*, Vol. 388, No. 10, May 15, 2009.

previously only hazily understood among the very largest firms.[29] Today, in addition to data on firm sizes, data exist on firm productivities (measured variously), ages and lifetimes, growth rates conditional on sizes and ages, various firm financials, some firm networks, and the location of firms in space. Essentially all of these data are resolved down to the level of individual firms, although for reporting purposes they are always binned to protect company privacy.

Armed with such microdata, researchers can build agent-based models that have close connections to each real-world actor. This is not the norm today, but several examples of this approach exist, and the increasing availability of microdata suggests that more efforts of this type are on the short-term horizon. For example, consider my model, based on firm-level data from administrative sources, of the U.S. private sector.[30] Each year over the past two decades, between 5 million and 6 million U.S. firms have employed a total of between 100 million and 120 million workers annually. Data on all of these firms and employees are available, appropriately anonymized to ensure privacy. The creation of a family of agent-based models to facilitate the study of the formation, operation, and evolution of such firms began at small scale when modest computing power was available,[31] but it has grown such that the entire private sector can now be represented at full scale.[32] Such models involve hundreds of millions of worker agents, who interact directly in production operations within millions of firms, with all agents and firms represented as unique software objects. These models can be estimated from empirical data so as to closely reproduce not just the myriad statistical properties of American businesses but also their dynamics, involving monthly job-to-job transitions by millions of workers and the formation and dissolution of tens of thousands of firms.[33] Other large-scale agent-based models involve traffic models resolved at the level of every vehicle in a city and epidemiological models written in terms of all of the susceptible, infected, and recovered or removed individuals in specific geographical regions.[34]

---

[29] Herbert A. Simon and Charles P. Bonini, "The Size Distribution of Business Firms," *American Economic Review*, Vol. 48, No. 4, September 1958.

[30] Robert L. Axtell, "120 Million Agents Self-Organize into 6 Million Firms: A Model of the U.S. Private Sector," in J. Thangarajah, K. Tuyls, C. Jonker, and S. Marsella, eds., *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*, Singapore, May 9, 2016a.

[31] Robert L. Axtell, "The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions," working paper, Washington, D.C.: Brookings Institution, 1999.

[32] Robert L. Axtell, "Endogenous Firm Dynamics and Labor Flows via Heterogeneous Agents," in Cars Hommes and Blake LeBaron, eds., *Handbook of Computational Economics*, Vol. IV, *Heterogeneous Agent Modeling*, Amsterdam: Elsevier, North-Holland, 2018.

[33] Robert L. Axtell and Omar A. Guerrero, *Dynamics of Firms from the Bottom Up: Data, Theories, and Models*, Cambridge, Mass.: MIT Press, forthcoming.

[34] For examples of traffic models, see Chris Barrett and Richard Beckman, *TRANSIMS: Portland Case Study Report*, Vol. I, *Introduction and Overview*, Los Alamos, N.M.: Los Alamos National Laboratory, 1995; and Kai Nagel and Christopher L. Barrett, "Using Microsimulation Feedback for Trip Adaptation for Realistic Traffic in Dallas," *International Journal of Modern Physics C*, Vol. 8, No. 3, 1997. For exam-

Full-scale agent-based models grounded in microdata need not be large scale, such as when a social process involves a relatively small number of people. Take a fishery, for instance, management of which might involve a very large number of fish but relatively few fishers.[35] For example, the POSEIDON model is an agent-based model that was initially created to better understand the effects of alternative policies on the ground-fishery off the North American West Coast.[36] All of the trawlers that operate there have had a variety of data-gathering facilities on board for several years, such as Global Positioning System devices, electronic logbooks, and even regulatory personnel, giving a quasi-comprehensive picture of the activities of the roughly 100 vessels working there. Using these data, the POSEIDON agent-based model incorporates a reasonably general behavioral model, particularized to that fishery, to closely reproduce the actual actions of the fishing fleet.[37] This model is being applied to other fisheries, including ones for which fewer data are available (e.g., in Indonesia). The key advance of such agent-based models over mathematical models in resource economics is the ability to represent heterogeneous behavior of boundedly rational people (i.e., fishers) interacting through social networks away from equilibrium.[38] This enables researchers to study policies that better preserve resources while producing economically viable yields.

ples of epidemiological models, see Alberto Aleta, David Martín-Corral, Ana Pastore y Piontti, Marco Ajelli, Maria Litvinova, Matteo Chinazzi, Natalie E. Dean, M. Elizabeth Halloran, Ira M. Longini, Jr., Stefano Merler, et al., "Modelling the Impact of Testing, Contact Tracing and Household Quarantine on Second Waves of COVID-19," *Nature Human Behaviour*, Vol. 4, No. 9, September 2020; Andrew Crooks and Ates Hailegiorgis, "Disease Modeling Within Refugee Camps: A Multi-Agent Systems Approach," in R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds., *2013 Winter Simulations Conference (WSC)*, Washington, D.C.: IEEE, 2013; and Stephen Eubank, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltán Toroczkai, and Nan Wang, "Modelling Disease Outbreaks in Realistic Urban Social Networks," *Nature*, Vol. 429, May 13, 2004.

[35] For example, see Stephen Lewis Scott, *Computational Modeling for Marine Resource Management*, dissertation, Fairfax, Va.: George Mason University, 2016.

[36] Richard M. Bailey, Ernesto Carrella, Robert Axtell, Matthew G. Burgess, Reniel B. Cabral, Michael Drexler, Chris Dorsett, Jens Koed Madsen, Andreas Merkl, and Steven Saul, "A Computational Approach to Managing Coupled Human–Environmental Systems: The POSEIDON Model of Ocean Fisheries," *Sustainability Science*, Vol. 14, No. 2, 2019.

[37] Ernesto Carrella, Steven Saul, Kristin Marshall, Matthew G. Burgess, Reniel B. Cabral, Richard M. Bailey, Chris Dorsett, Michael Drexler, Jens Koed Madsen, and Andreas Merkl, "Simple Adaptive Rules Describe Fishing Behaviour Better Than Perfect Rationality in the US West Coast Groundfish Fishery," *Ecological Economics*, Vol. 169, March 2020.

[38] For examples, see Colin W. Clark, *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*, 2nd ed., New York: Wiley Inter-Science, 2005; and Colin W. Clark, *The Worldwide Crisis in Fisheries: Economic Models and Human Behavior*, New York: Cambridge University Press, 2006.

The Parallel Social World and Single-Threaded Versus Multi-Threaded Agent-Based Modeling

In the real world, people are quasi-autonomous and take actions in accordance with their own goals and objectives, whatever those might be. A variety of norms, institutions, and technologies exist in human societies for partially synchronizing human activities, as when a church announces that it will hold services on a particular day and time or when an office or store posts its business hours. Other types of synchronization are more ephemeral, as when two vehicles meet at an intersection governed by a stop sign and the one arriving earlier gets to proceed first, by mutual agreement with an established social norm. But for large parts of their lives, people act asynchronously, doing what they want to do when and where it occurs to them to do so.

Perhaps somewhat peculiarly, given how real populations behave, almost all agent-based models do *not* model human behavior as occurring asynchronously. The simple reason is that the computer hardware and, to a lesser extent, software on which such models live are based on the digital computer architecture of von Neumann, in which there is a single central processing unit (CPU) and data are stored in random access memory (RAM); the appropriate data are copied into the CPU when needed and sent back to be stored once they are no longer needed. In essence, the flow of control in such computer architectures is serial in nature, with limited opportunities for parallel execution. Furthermore, because the hardware has these properties, the computer languages that have grown up to use such processors have mostly facilitated the creation of programs that run serially, on a single thread of execution, not in parallel. Given that all of the major ABM programming environments—e.g., Repast,[39] Multi-Agent Simulation of Neighborhoods or Networks (MASON),[40] and NetLogo[41]—generate code that is essentially single-threaded, probably 90–95 percent of agent-based models do violence to the real world of parallel, asynchronous interactions among agents.

However, probably far less than 90 percent of the clock cycles spent on agent-based models are serial, because many of the biggest models execute, at least to some extent, in parallel, using programming paradigms that permit multi-threaded code to run on modern, multicore hardware. Creating such code is typically not easy, because parallel programming is almost as much art as science today. Writing parallel code to solve problems in the natural sciences is often very hard given that such problems might not naturally decompose into neatly separable pieces that can each be deployed on a separate thread or processor. In the social sciences, the situation can be easier because people interact with only a few others at

---

[39] Michael J. North and Charles M. Macal, *Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation*, New York: Oxford University Press, 2007. Repast high-performance computing does permit parallel execution of agent-based models.

[40] Sean Luke, *Multiagent Simulation and the MASON Library*, Fairfax, Va.: George Mason University, 2015. A variant known as D-MASON facilitates distributed execution, especially explicitly spatial agent-based models.

[41] Uri Wilensky and William Rand, *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with NetLogo*, Cambridge, Mass.: MIT Press, 2015.

a time, making agent-based models naturally parallelizable. Typically, achieving parallelism with agent-based models is done by simply breaking the population of agents up into pieces that are appropriate to the architecture, letting those agents that are on the same thread interact, and then periodically bringing all the agents back together to be redivided into somewhat different groups to run on new threads. By writing agent-based models in this way, it is often possible to parallelize agent execution to achieve nearly linear speedup with the number of processors or cores or threads available. The trends in microprocessor design, development, and production are shown in Figure 16.1, which depicts an exponentially growing number of transistors and processors, plateauing processor performance in sequential operations, flatlining processor speeds, and increasing numbers of cores per processor. With 48 core CPUs available (although at premium cost), and with designs for many more cores per chip on company drawing boards, the era of parallel computing hardware is clearly upon us. But, at the same time, the operating systems and programming languages—in short, the software—running on this hardware typically were not designed for such parallel environments. This suggests that at least some rethinking of the situation is in order.

Specifically, although several add-ons to conventional serial programing languages exist to facilitate parallel execution, such as OpenMP and MPI,[42] these extensions are perhaps best suited for parallelizing models that were previously single-threaded. Newer software libraries for parallelization, although more flexible and more tailored to the hardware level,[43] are capable of providing better speedups but do not really change the basic parallel programming paradigm. Furthermore, although moving execution off bottlenecked CPUs and onto graphical processing units (GPUs) has shown promise,[44] doing so poses unique problems for agent-based models, such as working around synchronous updating. Comparing several contemporary parallel languages and frameworks shows a wide variety of performance variations, even on a relatively simple agent-based model.[45]

---

[42] Rohit Chandra, Leonardo Dagum, Dave Kohr, Dror Maydan, Jeff McDonald, and Ramesh Menon, *Parallel Programming in OpenMP*, San Francisco, Calif.: Morgan Kaufmann Publishers, 2001; Barbara Chapman, Gabriele Jost, and Ruud van der Pas, *Using OpenMP: Portable Shared Memory Parallel Programming*, Cambridge, Mass.: MIT Press, 2007; and Michael J. Quinn, *Parallel Programming in C with MPI and OpenMP*, New York: McGraw-Hill, 2003.

[43] For instance, see James Reinders, *Intel Threading Building Blocks: Outfitting C++ for Multi-Core Processor Parallelism*, Sebastopol, Calif.: O'Reilly, 2007.

[44] R. M. D'Souza, M. Lysenko, and K. Rahmani, "SugarScape on Steroids: Simulation over a Million Agents at Interactive Rates," presented at AGENT 2007 Conference on Complex Interaction and Social Emergence, Evanston, Ill., 2007; and Mariam Kiran, Paul Richmond, Mike Holcombe, Lee Shawn Chin, David Worth, and Chris Greenough, "FLAME: Simulating Large Populations of Agents on Parallel Hardware Architectures," in Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, eds., *AAMAS '10: Proceedings of the 9th International Conference on Autonomous Agent and Multiagent Systems*, Toronto, Canada, 2010.

[45] Stefan McCabe, Dale Brearcliffe, Peter Froncek, Marta Hansen, Vince Kane, Davoud Taghawi-Hejad, and Robert L. Axtell, "A Comparison of Languages and Frameworks for the Parallelization of a Simple

FIGURE 16.1

Trends in Computer Hardware over the Past 40 Years

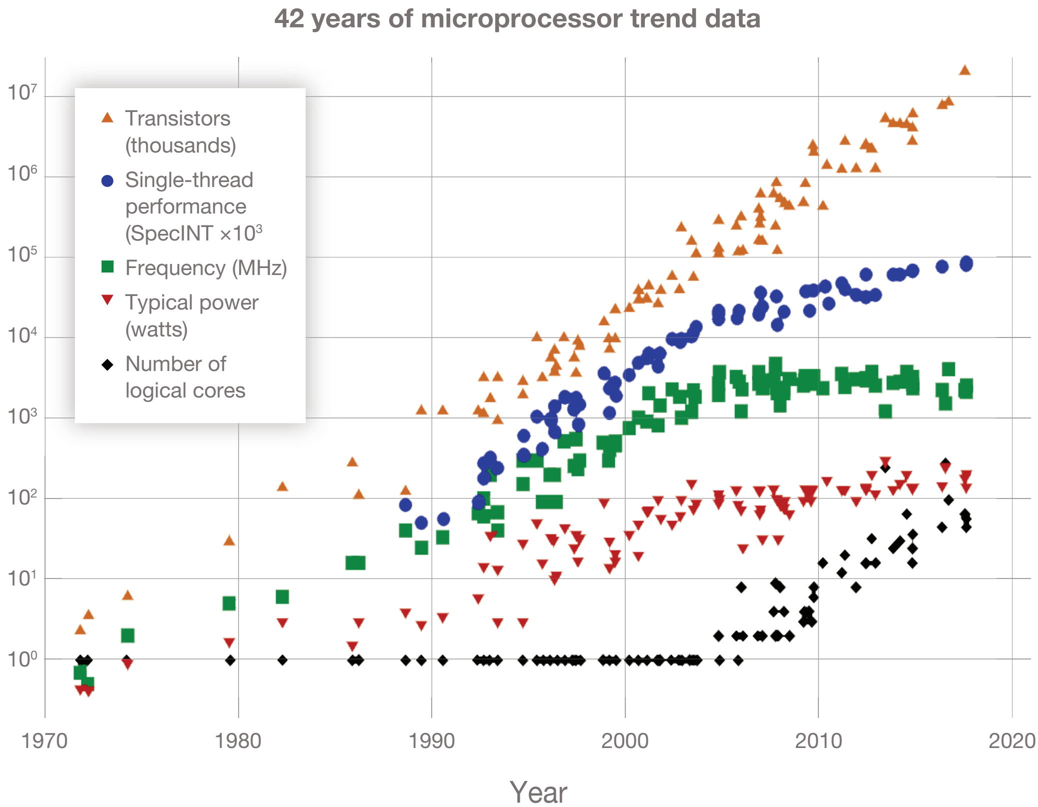

SOURCE: Karl Rupp, "Microprocessor Trend Data," data repository, last updated July 16, 2020 (CC BY 4.0).
NOTE: Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten. New plot and data collected for 2010–2017 by K. Rupp.

All of these considerations suggest the desirability of making the creation of parallel agent-based models easier, their deployment on extant hardware simpler, and their migration to hardware with greater parallel capabilities—i.e., a larger number of cores in the short term—more straightforward. Some new thinking and novel technologies to accomplish all of this have appeared on the technological horizon with the promise of facilitating the growth of parallel, multi-threaded agent-based models. However, there might be systematic, structural bottlenecks to the widespread adoption of such innovations.

In the next two sections, I speculate on the capabilities of these new ideas, with an eye toward the opportunities that they provide and the limitations that are apparent at this time. I try to look beyond existing programming paradigms and hardware architectures to think

Agent Model," working paper, Fairfax, Va.: Department of Computational and Data Sciences, George Mason University, forthcoming.

about what a better future might look like for the parallel execution of agent-based models in the short term, when few systematic changes can be made, and in the long term, when perhaps the overall structure of languages and hardware can be evolved to accommodate the needs of future agent-based models that are likely to be very large scale.

Large-Scale Models Realized with Threading: Opportunities

There are several distinct motivations for creating agent-based models that are multi-threaded, most of which are implicit in the previous section and which I make explicit here. These motivations represent opportunities to create new *flavors* of agent-based models, the full implications of which are not fully understood today.

The First Motivation

The first motivation for parallel execution with threads is (as discussed in the previous section) that the real world of quasi-autonomous individuals is full of parallel activity, and to represent behavior in any other way is potentially problematic. Such concerns were on display in the early days of ABM, when ostensibly significant results about the well-known prisoner's dilemma model from game theory were reported by Nowak and May for interactions on a spatial landscape, executed in parallel but with perfectly synchronous updating.[46] Their attempt to use parallel updating was laudable, but Huberman and Glance quickly demonstrated that the synchrony they employed generated patterns that were not robust to the relaxation of the updating mode toward asynchrony.[47] In essence, Nowak and May's results about high levels of cooperation were computational artifacts and quickly gave way to the expected results of pure defection with even a small amount of asynchrony. Although Nowak and May defended their results on the grounds that many biological processes are largely synchronous, it is now understood that their main purported results are a classic example of a "brittle outcome" resulting from the microscopic specification of agent interactions—in this case, agent updating.

At the opposite extreme of parallel updating is the serial execution model that is the norm in agent-based models today. This is a peculiar representation of human behavior, but one with some basis in human (or modeler) cognition. Often when contemplating our own actions, we take those of others as unchanging or fixed, perhaps through some sense of what constitutes typical human interactions. For example, in the grocery store, we expect the checkout person to scan our selections, say how much we owe, and offer us a receipt. We do not expect that person to throw our groceries in the garbage can or drink the milk that we have put on the conveyor belt. Nor do we expect other customers to be eating food in the aisles or throwing

[46] Martin A. Nowak and Robert M. May, "Evolutionary Games and Spatial Chaos," *Nature*, Vol. 359, October 29, 1992.

[47] Bernardo A. Huberman and Natalie S. Glance, "Evolutionary Games and Computer Simulations," *Proceedings of the National Academy of Sciences*, Vol. 90, No. 16, August 1993.

items onto the floor. We imagine that we will successfully purchase food at the grocery store because everyone will engage in conventional behavior and, to a first approximation, their behavior does not interfere with ours.

This is also how rules of agent behavior are often constructed. We think about a typical agent (object) and write code (methods) for the actions (messages) that it will take with (send to) other agents and the ways that it may change either its own state (instance variables) or that of the environment, depending on how other agents behave. In thinking about such interactions, whether with other shoppers at the grocery store or with other agents in an agent-based model, we typically take the rest of the world as fixed or in some stationary state. My decision calculus at the store would be quite different if I knew that an arsonist was at my house setting it on fire or that the President was about to declare nuclear war on an adversary. In the same way, the rule writer, coder, or programmer abstracts from the agents that it is not interacting with, taking their behavior as given and not interfering with its own actions. This may or may not be a good way to think about truly parallel social worlds, but it is a passable approach to social cognition.[48]

To be completely clear, single-threaded agent-based models represent human behavior by freezing the actions of all agents that are not being updated. In looking at typical visualizations of agent-based models, this successive freezing and unfreezing is not obvious because the execution speed is so fast, causing many agent states to change in a short amount of time, giving the false impression of parallel interactions when only one agent, or perhaps a few, is actually active at once.

The Second Motivation

Another rationale for parallel execution of agent-based models is much more pragmatic. Many large-scale models are simply not feasible unless they can be sped up by running on parallel hardware. Such models have so many agents or so much cognition per agent, or both, that single-threaded execution would take too much wall time to make them practical. In this case, many processors and multiple cores are manna from heaven, and the question becomes simply how best to leverage the technology. There exist many parallel programming schemes, some alluded to already, and I will not go into detail on them here. Suffice it to say that there is a growing variety of approaches, each with its own set of advantages and disadvantages, and the situation is rapidly evolving. Agent-based modelers gain from the technologies created in other fields for parallelization.

The Third Motivation

Another reason for multi-threaded agent-based models is close to the previous one but distinct from it. Imagine a model of some social phenomenon that is written at full scale and

[48] Martin Davies and Tony Stone, *Mental Simulation*, Cambridge, Mass.: Blackwell Publishers, 1995.

produces large amounts of output that can be directly compared with the real world, because the model is on the same scale. The physicist Richard Feynman noted that "[t]he same equations must have the same solutions."[49] By analogy, an agent-based model and its real-world twin should generate the same data. In my experience with the 120 million–agent model of firm dynamics, it is simply much easier to estimate parameters of a model, particularly a large model for which each run is computationally expensive, when the model output is directly comparable with the actual target data. Of course, models are always abstractions, to greater or lesser extents, so model output will never exactly coincide with data, but to try to accomplish this with single-threaded models running at small scale adds layers of difficulty.

The Fourth Motivation

A fourth incentive for parallel execution, partially related to the first, is that multi-threaded agent-based models may yield different results than their serial counterparts. Conceptually, it is easy to see how this *might* be the case, because serial agent updating is so stilted and so computationally extreme, with most agents spending most of their lives just waiting around to be active. It is as if each agent gets 15 minutes of fame over its entire life, to execute all of its most important tasks, and is frozen in amber for the rest of the time. For a concrete example of parallel activation yielding different results, consider the so-called zero-intelligence (ZI) trader model of Gode and Sunder,[50] in which financial market agents use a heavy dose of randomness in their own behavior yet produce overall market behavior that is broadly in line with data.[51] Cliff has examined the relative performance of several distinct types of ZI traders,[52] finding patterns of stochastic dominance of certain types over others. What is interesting is that he gets *different* results in the single-threaded version of his code than in multi-threaded versions. He has advanced a rationale for why this is so, tied up with the details of so-called order books in financial markets, which I will not delve into here. The main point is that

[49] Richard P. Feynman, Robert B. Leighton, and Matthew Sands, *The Feynman Lectures on Physics*, Vol. II, *Mainly Electromagnetism and Matter*, Reading, Mass.: Addison-Wesley Publishing Company, 1964, p. 12–1.

[50] Dhananjay K. Gode and Shyam Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, Vol. 101, No. 1, February 1993; and Dhananjay K. Gode and Shyam Sunder, "What Makes Markets Allocationally Efficient?" *Quarterly Journal of Economics*, Vol. 112, No. 2, May 1997.

[51] See, for example, J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko, "The Predictive Power of Zero Intelligence in Financial Markets," *Proceedings of the National Academy of Sciences*, Vol. 102, No. 6, February 8, 2005.

[52] Dave Cliff, "Methodological Mess-Ups in Modelling Markets with Minimal-Intelligence Agents," presentation at the First Conference on Zero/Minimal Intelligence Agents, New Haven, Conn.: Yale University, October 23, 2020.

different results can be produced through multi-threading—that certain results from single-threaded agent-based models might be artifacts of the way they are coded.[53]

The Fifth Motivation

Lastly, a not often articulated rationale for large-scale models realized in threaded environments is simply "why not," given that the current technological zeitgeist involves ever expanding computing resources. Now is a transition time, but, at some point in the not too distant future, it will be the norm to create large models that are executed on parallel machines, and it is high time to get started on this enterprise. A different way to say this is that researchers will have much more-capable computing resources in the future than at present, so it is important to figure out now the basic principles of large-scale agent computing and to elucidate them.

Toward Ubiquitous Large-Scale, Multi-Threaded Agent-Based Models: Bottlenecks

Among the many metrics that guide progress in parallel ABM research, certainly one of the most important is the degree of speedup achieved, as when a model that is run on a 16-core processor runs ten times faster in its multi-threaded implementation than when single-threaded. The holy grail is *linear speedup*—that is, doubling the number of processors, threads, or cores halves the wall time required for model execution.[54] Workstations with shared memory and $O(100)$ cores are now available, with distributed memory clusters with $O(1{,}000)$ cores often available at universities, while the Bridges system at the Pittsburgh Supercomputing Center has $O(10{,}000)$ cores. Clearly, if speedups of 100 to 10,000 times can be achieved, then very large-scale models can be pursued. This is the situation today; by 2025, machines will be two to four times more capable, while perhaps by 2030 it will be by another factor of 10.

However, as the search for linear speedups of agent-based models unfolds, there are some dark clouds on the horizon. Various problems plague the conventional threading model in general and in its specific application to agent-based models. I will go through some of these, treating them rather briefly while pointing to the literature for details.

Efficient parallelization often entails being able to partition a problem into more or less independent pieces that each run on their own threads, cores, or processors. When the

[53] Researchers do not have a general theory for how and why such differences appear in single-threaded versus multi-threaded execution, an important lacuna in the ABM paradigm.

[54] It is sometimes possible to achieve superlinear speedups, as when a model executed on 100 threads runs *more than* 100 times as fast. Such phenomena are usually due to cache behavior, and my experience is that, although real for toy models, the presence of superlinear model performance is rarely encountered in large-scale agent-based models. For example, it is present in certain parallel codes for the ZI model but has not been found in any variants of my 120 million–agent *firms* model. See McCabe et al., forthcoming.

threads need to interact, whether to share data, check a condition, or simply synchronize, significant slowdown often occurs. One reason for this is that, on modern hardware, there is no (easy) way to guarantee thread execution order or timing, about which I will say more later; the upshot is that a running thread might have its execution blocked for many reasons, even for checking whether another thread is running only to find that it is not. For agent-based models, there is no easy escape from such problems, and they manifest in specific ways that are worth mentioning to expose the nature of the underlying problems.

One important rationale for building an agent-based model in the first place is to replace uniform-interaction (mean-field) models, which are easy to work with mathematically but highly unrealistic, with social networks of some kind. This has the beneficial effect on ABM execution that, when it is an agent's turn to act, that agent need look at only a subset (typically small) of the agent population. The character of agent-to-agent interactions—i.e., primarily local and close-knit[55]—means that few data on far-flung agents have to be brought into an agent's decisionmaking calculus. The trick is to put all of the agents who are most closely interacting onto the same thread or process because interthread or interprocess communication is usually costly, as mentioned already. Modeling agent-to-agent interactions is the key to many and perhaps even most agent-based models; thus, getting the most–densely interactive agents onto the same thread usually results in large performance improvements. Stated the opposite way, when agents cannot be partitioned into groups with dense intragroup interactions, because either the actual interactions are not known ahead of time or the interactions are not really clustered but are more homogeneously distributed throughout the entire population, then there often will be very little speedup associated with moving to multi-threaded models.

For a concrete example, consider a spatial agent-based model in which the agents interact only with their physical neighbors. In such a model, tessellating the space into regions and putting each of these on its own thread usually results in significant speedups.[56] However, when agents either move or interact across the region boundaries and, therefore, across threads, it might be the case that significantly less speedup is realized.

A related problem involves load-balancing across threads. In common fork-join parallelism, execution can move beyond the "join" only once the last thread terminates.[57] If the threads are not loaded comparably, then there will be significant idle time as many threads wait for a few to complete their tasks, reducing performance.[58] Load-balancing can be tricky with agent-based models because naïve approaches, such as putting comparable numbers

[55] Gabriel E. Kreindler and H. Peyton Young, "Fast Convergence in Evolutionary Equilibrium Selection," *Games and Economic Behavior*, Vol. 80, July 2013.

[56] This is the approach taken by the *D-MASON* extension of the basic *MASON ABM* framework in Java.

[57] For a discussion of the fork-join algorithm as an early and central form of parallel programing, see Linus Nyman and Mikael Laakso, "Notes on the History of Fork and Join," *IEEE Annals of the History of Computing*, Vol. 38, No. 3, July–September 2016.

[58] Load-stealing by idle threads, as is possible with Intel's Threading Building Blocks library, appears to be a good way to deal with this problem.

of agents on threads, will lead to poor performance when agent execution time is highly variable. For a real example, consider the production phase that happens each month in my 120 million–agent firms model,[59] during which all firms make output. Because the firms that grow up in this model are highly heterogeneous, with most having just a few worker agents and a few having hundreds of thousands or even 1 million employees, and because the execution time needed to produce output in a firm is proportional to the number of employees the firm has, load-balancing cannot be done simply by partitioning the firm population or the employee population; load-balancing must evaluate both populations. Poor load-balancing in that model can slow execution down by a factor of 10.

When threads are used to take advantage of all processing cores, many of the difficulties associated with writing good parallel ABM code have to do with the fact that threads running on CPUs can be interrupted at any time to perform some other, perhaps more machine- or system-critical, task. The difficulties posed by existing threading technologies have been extensively discussed in the computer engineering literature.[60] Interrupts make it impossible for modern CPUs to guarantee when any specific thread will execute, and this makes the fork-join model an imperfect paradigm for ABM parallelism. In the next section, I will discuss what alternative technologies might look like. Here, I conclude with observations regarding the constraints imposed by threading technologies on agent-based models, and the bottlenecks that need to be dealt with to achieve higher levels of speedup.

Consider an agent-based model in which the number of agents is much larger than the number of cores so that each thread will manage the execution of many agents, as is common today: e.g., 120 million agents on $O(100)$ cores. Parallelism of this type can often lead to good speedups but does not solve the problem of unknown artifacts being impressed into what are essentially $O(100)$ single-threaded execution streams. That is, there are still more than 1 million agents executing on each core in single-threaded fashion. On top of this, to avoid generating further computational artifacts, it is important to regularly remix the agents onto different threads so that micro-correlations, such as agent *i* always moving before agent *j* and so on, do not occur.[61] Furthermore, it seems that, to write efficient parallel code, it is necessary to be able to say, at least by run-time, which agents or agent groups are going to be the biggest users of clock cycles. If it is not possible to do so, then it is practically impossible to do any kind of even approximate load-balancing, thus jeopardizing the goals of parallel agent-based models.

---

59 Axtell, 2016a; Axtell, 2018.

60 For example, see Edward A. Lee, *The Problem with Threads*, Berkeley, Calif.: Electrical Engineering and Computer Sciences, University of California at Berkeley, January 10, 2006.

61 Robert L. Axtell, "Effects of Interaction Topology and Activation Regime in Several Multi-Agent Systems," in Scott Moss and Paul Davidsson, eds., *Multi-Agent-Based Simulation*, Vol. 1979, Heidelberg, Germany: Springer Verlag, 2001a; and Kenneth W. Comer, *Who Goes First? An Examination of the Impact of Activation on Outcome Behavior in Agent-Based Models*, dissertation, Fairfax, Va.: George Mason University, 2014.

Technological Advances That Would Accelerate the Development of Agent-Based Modeling Technologies

More cores, faster memory, bigger caches, better parallel programming paradigms—all of these technologies are visible from the research frontier outpost that researchers now occupy. All will make agent-based models more capable by permitting larger models that run faster and are easier to program and debug. Some approaches to "debottlenecking" are also on the horizon, and there can be little doubt that many of the problems described in the previous section are likely to find either solutions or satisfactory workarounds as ABM tools and codebases mature. Some of the challenges will be harder to resolve than others and might require significant innovations, but progress is both inevitable and, from someone who has struggled trying to bend many extant computing technologies toward the needs of ABM over the years, very welcome.

Advances in Hardware

There will come a day when each agent runs on its own core, but that day is a long way off, at least for large models. This would not be a panacea anyway, given that the raison d'être for ABM in the first place is to model interactions; perhaps two agents per core is the lowest agent density worth pursuing, at least as long as interthread communication is slower—and more hazardous from a data race perspective—than intrathread operations.

From the highest-level perspective, what is needed on the hardware side are noninterruptible processors, or at least cores, so that execution streams (e.g., threads) would have guaranteed run-time. This would solve many problems of thread synchronization and would also make life easier in other ways. First, fork-join processes would have repeatable, reliable behavior. Second, debugging would be much easier because execution would be (more) deterministic. Third, load-balancing could be accomplished empirically, as when running an agent-based model many times shows which processes, routines, or agent groups need more time than others. Imagine a world of many core processors in which a few cores are dedicated to the operating system but most are available for the models and cannot be interrupted by the operating system. Specifically, some kind of CPU-GPU hybrid, in which the CPU precomputes agent partitions while the model runs on the GPU, might work for agent-based models.[62] Having 10,000 to 100,000 lightweight cores operating asynchronously could go far toward having a "supercomputer on a board" for agent-based models.

Less speculatively, innovations in other parts of the machine would also serve the ABM community well. In certain respects, an agent-based model is a giant database of agent state information, maintained in RAM and updated according to the agent objects' rules of behavior. Advances in database technologies, involving persistence and various kinds of error-checking, might help provide guarantees of the repeatability of ABM execution. Improvements in display technologies would also be of importance for progress with agent-based

[62] Intel's Phi coprocessor board was a step in this direction but is now discontinued.

models. Today's biggest 6K displays have something short of 50 million pixels. Visualizing the state of a model with 100 million agents is not possible until resolutions go up by a factor of 5 or 10 and will probably not be useful until a factor of 20 or 50 times larger is reached. Finally, because agent-based models are large databases, fast storage technologies for huge state spaces are also of interest. Recording the full evolution of 100 million agents, each 1 kilobyte, running for 10,000 periods would require at least a petabyte of storage per run, an exabyte for many runs, and a zettabyte for many model instances in a model family, assuming that the full evolution of agent states is useful to record.

Advances in Software: Agent Languages, Operating Systems, and Software Engineering

Although it might be natural to think of operating systems as more foundational than programming languages, standardized languages came first, and many early mainframe installations had highly localized, essentially bespoke operating systems. In contemplating the kinds of changes in software ecosystems that would accommodate large-scale, parallel-executed agent-based models, it is not clear whether agent-oriented operating systems or specialized programming languages would come first or which would provide the greater performance boost. Certainly, if new kinds of uninterruptible hardware were to become available, such as is discussed in the previous subsection, new operating systems would be needed. However, it is also unclear whether these would or should be agent-oriented in any meaningful way.

In the past decade, there have been various proposals for developing specialized operating systems based on agents.[63] This approach is related to self-aware computing systems.[64] Agent-oriented software engineering has been around for some time.[65]

Ideas about agent-oriented programming languages have been discussed in the artificial intelligence community for long enough to have reached a certain level of maturity.[66] Origi-

[63] For example, see Javier Palanca Cámara, Marti Navarro, Estefania Argente, Ana Garcia-Fornes, and Vicente Julián, "Modeling an Operating System Based on Agents," in Emilio Corchado, Václav Snášel, Ajith Abraham, Michal Wozniak, Manuel Graña, and Sung-Bae Cho, eds., *Hybrid Artificial Intelligent Systems: 7th International Conference, HAIS 2012 Proceedings*, Pt. 1, Salamanca, Spain, March 28–30, 2012.

[64] Javier Cámara, Kirstie L. Bellman, Jeffrey O. Kephart, Marco Autili, Nelly Bencomo, Ada Diaconescu, Holger Giese, Sebastian Götz, Paola Inverardi, Samuel Kounev, and Massimo Tivoli, "Self-Aware Computing Systems: Related Concepts and Research Areas," in Samuel Kounev, Jeffrey O. Kephart, Aleksandar Milenkoski, and Xiaoyun Zhu, eds., *Self-Aware Computing Systems*, Cham, Switzerland: Springer Nature, 2017.

[65] Paolo Ciancarini and Michael J. Wooldridge, eds., *Agent-Oriented Software Engineering*, Vol. 1957, New York: Springer, 2000; and Nicholas R. Jennings, Francisco J. Garijo, and Magnus Boman, "Agent-Oriented Software Engineering," presented at the Multi-Agent System Engineering: 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Valencia, Spain, July 1999.

[66] For examples, see Yoav Shoham, "Agent-Oriented Programming," *Artificial Intelligence*, Vol. 60, No. 1, March 1993; and Yoav Shoham, "An Overview of Agent-Oriented Programming," in Jeffrey M. Bradshaw, ed., *Software Agents*, Menlo Park, Calif.: American Association for Artificial Intelligence/MIT Press, 1997.

nally envisioned as a programming paradigm analogous to object-oriented programming (OOP), agent-oriented programming has not evolved or matured as expected, despite a proliferation of tools.[67] Analogous, domain-specific programming paradigms, such as market-oriented programming,[68] remain nascent.

In the short term, it might be the creation of specialized, disciplinary-specific software libraries that will make the biggest impact on the creation of ABM codes. Typically, when a researcher begins to build a new agent-based model, they harvest code from previous, similar models, often their own. Given that there are several domains that are heavily worked by agent-based models, such as traffic, finance, ecology, and epidemiology, the creation of code libraries in these areas would greatly facilitate the more rapid creation of such models. There is currently little incentive for researchers to condense their code into such libraries, and little support from funding agencies to create software tools, so it is unclear how such efforts might successfully unfold, although there seems to be more momentum in certain areas than in others.[69]

Closely related to the creation of software libraries is the idea of community software, which is used in the climate modeling community.[70] There, the husbanding of code resources—i.e., programs and data—at the Geophysical Fluid Dynamics Laboratory and the National Center for Atmospheric Research has led to the progressive evolution of both more-comprehensive and more-accurate models. A similar situation exists at the National Weather Service. No comparable efforts are underway in the social sciences, although the need for such efforts is clear.[71]

---

67 For instance, see Rafael H. Bordini, Jomi F. Hübner, and Renata Vieira, "Jason and the Golden Fleece of Agent-Oriented Programming," in Rafael H. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah Seghrouchni, eds., *Multi-Agent Programming: Languages, Platforms and Applications*, New York: Springer Science+Business Media, Inc., 2005; and Sebastian Rodriguez, Nicolas Gaud, and Stéphane Galland, "SARL: A General-Purpose Agent-Oriented Programming Language," presented at the 2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), Warsaw, Poland, August 11–14, 2014.

68 Michael P. Wellman, "Market-Oriented Programming: Some Early Lessons," in Scott H. Clearwater, ed., *Market-Based Control: A Paradigm for Distributed Resource Allocation*, Singapore: World Scientific, 1996; and Michael P. Wellman, Amy Greenwald, and Peter Stone, *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*, Cambridge, Mass.: MIT Press, 2007.

69 See, for example, J. Doyne Farmer and Duncan Foley, "The Economy Needs Agent-Based Modeling," *Nature*, Vol. 460, No. 6, August 2009.

70 David M. Higdon, Robert L. Axtell, Venkatramani Balaji, Lawrence E. Buja, Katherine V. Calvin, Kathleen M. Carley, Rebecca Castaño, Ronald R. Coifman, Omar Ghattas, James A. Hansen, et al., *From Maps to Models: Augmenting the Nation's Geospatial Intelligence Capabilities*, Washington, D.C.: The National Academies Press, 2016.

71 For example, when the 2007–2009 financial crisis hit, U.S. government agencies had limited data and modeling capabilities available, leading to a variety of decisions, such as those concerning the size of bailouts needed, that were made by back-of-the-envelope calculations. In the wake of this situation, a National Institute of Finance was proposed to aggregate relevant data and models that would be useful when the next crisis hit, to be funded in one proposal at least at the $1 billion level. Although this effort ultimately failed, the Dodd-Frank Act funded the creation of the Office of Financial Research for just such data-gathering efforts.

Toward Automated Synthesis of Agent-Based Models

One way in which agent-based models are still something of an art form requiring domain-specific knowledge is in the specification of agent rules of behavior. In the context of a micro-to-macro modeling paradigm, such rules play cornerstone roles in most agent-based models. Sometimes there are microdata that provide detailed guidance for what such rules should look like, perhaps data gleaned from experiments with human subjects,[72] but aggregate data are more commonly available, and, in these cases, the agent rules must be inferred according to which specifications are sufficient to "hit" the target data. In calibrating or estimating models in this way, ABM is quite similar to standard statistics and econometrics, in which parameters are inferred from data. Conventional techniques, such as estimation by simulation,[73] are readily applicable to agent-based models. However, because agent-based models can be large and computationally expensive, it is generally costly in terms of time and effort to obtain well-fitting parameters, and much less is known about the identifiability of model parameters in agent-based models than in conventional mathematical models.

In light of these difficulties, ideas have been developed of late about how to use machine learning and other automated techniques to create and calibrate agent-based models, thus leading to progress in several areas. Specifically, for agent-based models in finance, an area in which copious amounts of data are available, machine learning has been used to estimate nonlinear relationships between inputs and outputs in some well-known models.[74] Relatedly, deep learning has been used to develop a better understanding of agent-based models and has also been used within agent-based models.[75] A recent idea with many interesting implications (e.g., the amount of volatility in financial markets) involves specifying the behavioral repertoire for agents as machine learning;[76] that is, individual agents use machine learning to figure out how to behave.

[72] For example, see John Duffy, "Agent-Based Models and Human Subject Experiments," in Leigh Tesfatsion and Kenneth L. Judd, eds., *Handbook of Computational Economics*, Vol. 2, *Agent-Based Methods*, New York: North-Holland, 2006.

[73] Daniel McFadden and Paul A. Ruud, "Estimation by Simulation," *Review of Economics and Statistics*, Vol. LXXVI, No. 4, November 1994.

[74] William A. Brock and Cars H. Hommes, "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model," *Journal of Economic Dynamics and Control*, Vol. 22, Nos. 8–9, July 1998; Giorgio Fagiolo and Giovanni Dosi, "Exploitation, Exploration and Innovation in a Model of Endogenous Growth with Locally Interacting Agents," *Structural Change and Economic Dynamics*, Vol. 14, No. 3, 2003; and Francesco Lamperti, Andrea Roventini, and Amir Sani, "Agent-Based Model Calibration Using Machine Learning Surrogates," *Journal of Economic Dynamics & Control*, Vol. 90, May 2018.

[75] Sander van der Hoog, *Deep Learning in (and of) Agent-Based Models: A Prospectus*, arXiv.org, June 20, 2017.

[76] Christophre Georges and Javier Pereira, "Market Stability with Machine Learning Agents," *Journal of Economic Dynamics & Control*, Vol. 122, January 2021.

More generally, evolutionary programming and related heuristic techniques have been used to solve the inverse problem of determining agent behaviors from data.[77] Among the issues that arise with such approaches—overfitting, the value of minimal models—the general question of when data are sufficient to determine model parameters—model identification, again—is crucial. In general, there will be many configurations of microsystems that are compatible with macrodata.

For example, a city has certain gross characteristics independent of the arrangement of the dry cleaning and convenience stores on the corner of State and Main Streets. Under the assumption that the main goal of most agent-based models is *not* to explain or predict microconfigurations, it is reasonable to expect micro-founded models to be underdetermined in the sense that the data, no matter how detailed, will not usually have enough context and historical depth to permit the digital world to evolve in exactly the same way as the real world. There are combinatorially enormous numbers of microstates that will generally be consistent with the data to be explained or predicted, as in statistical mechanics, in which the specific locations and velocities of *O*(Avogadro's number) particles are irrelevant if only the temperature of the room is the key statistic. Therefore, it is reasonable to expect that a host of distinct inverse methods for inferring rule systems for agent-based models will be broadly equivalent, not in the sense of yielding the same rules but in the ability to produce the same kinds of output. Questions about how to define and make use of equivalence classes of such rule systems or even the inverse methods themselves are important frontier issues on which additional research is needed, a topic that probably should be considered as high priority within the ABM research community.

Putting the Pieces Together: The Promise of Large-Scale Agent-Based Models Formulated from Big Data and Executing in Parallel

It is my hope to have communicated some of the excitement around ABM, in terms of both what can be achieved today and the developments on the horizon. The field has progressed immensely from its artificial life and OOP beginnings, running on the first generation of microcomputer hardware, with little or no data and a single thread of execution. If the coming decades experience comparable evolutionary developments, then by 2030 whole new classes of social science models should be possible, to say nothing of 2040 or 2050.

---

[77] Chathika Gunaratne, *Evolutionary Model Discovery: Automating Causal Inference for Generative Models of Human Social Behavior*, dissertation, Orlando, Fla.: University of Central Florida, 2019; and Tuong Manh Vu, Charlotte Probst, Joshua M. Epstein, Alan Brennan, Mark Strong, and Robin C. Purshouse, "Toward Inverse Generative Social Science Using Multi-Objective Genetic Programming," presented at the Genetic and Evolutionary Computation Conference (GECCO '19), Prague, Czech Republic, July 13–17, 2019.

Models are increasingly seen as essential to both positive scientific understanding and the orderly creation, implementation, and execution of public policy.[78] We are quickly moving beyond the idea that data alone are sufficient. While data analysis can expose relationships and suggest explanations, causal models that are capable of generating artificial data having the same properties as real-world data can provide much richer explanations for why the data have the structure they do. Social scientists now possess historically unprecedented amounts of high-quality data and the requisite computing power to process them, which has led to the appearance of wholly new, qualitatively different types of models, some of which I have mentioned already.

Since the mid-19th century, the idea of representing all individuals who are active in social processes in models has been a cornerstone of economics and finance,[79] spilling out into certain domains of other fields as well. However, from a practical point of view, this *methodological individualist* perspective was stillborn in essentially every area investigated because of the inability to render models at full scale with the real world. This was especially so given that the techniques that worked so well for physics (e.g., statistical mechanics) did not readily translate to economics because of heterogeneity, network effects, and so on.[80]

I think that it is not widely understood that only now, some 150 years beyond the birth of methodological individualism, are social scientists in a position to fully realize such models, in which millions or even billions of agents take parallel, asynchronous actions using data that they glean from the world in pursuit of their own self-interests. Interestingly, the ABM method that is the only real way to realize such models today also trades in concepts of methodological holism and pluralism, recognizing the important role of emergent phenomena within systems of interacting agents.[81] The newfound capabilities—the combination of ABM as a methodology, new computing opportunities, and new data—will manifest themselves in many ways over the next decades, so that by 2050 entirely new and more comprehensive agent-based models will have taken their places in scientific and policy circles; this will supplant certain mathematical and statistical models that are today based on limited, aggregate, and infrequently updated (e.g., quarterly) data that abstract from the vast heterogeneity that is present in the real world.

---

78 Higdon et al., 2016.

79 Carl Menger, *Investigations into the Method of the Social Sciences*, trans. Francis J. Nock, Grover City, Pa.: Libertarian Press, Inc., [1871] 1985; and Léon Walras, *Éléments d'Économie Politique Pure, Édition Définitive, Revue et Augmentée*, Paris: Pichon et Durand-Anzias, 1926.

80 Philip Mirowski, *More Heat than Light: Economics as Social Physics, Physics as Nature's Economics*, New York: Cambridge University Press, 1989.

81 See, for example, Robert L. Axtell, "What Economic Agents Do: How Cognition and Interaction Lead to Emergence and Complexity," *Review of Austrian Economics*, Vol. 20, No. 2, 2007; Robin Clark and Steven O. Kimbrough, "The Spontaneous Emergence of Language Variation from a Homogeneous Population," presented at the Computational Social Science Society of America, Santa Fe, N.M., 2015; and Nigel Gilbert, "Varieties of Emergence," in Charles Macal and David Sallach, eds., *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, Chicago: University of Chicago, October 2002.

Speculations on the kinds of new agent-based models that will appear in economics in the short to medium terms have been made elsewhere;[82] these include macroeconomic models written in terms of hundreds of millions of interacting agents, international trade models grounded in data on *all* individual firms that engage in export-import behavior, and economics models that are tied much more closely to the people and conditions in developing countries as opposed to regression models that are based on misspecified relationships derived from the properties of already developed nations.

In this chapter, I extend such speculations beyond economics proper to questions of political economy broadly construed, the defense and advancement of U.S. national security, and what can be done to better understand the latter through ABM.

One class of agent-based models that seems quite clearly on the horizon is policy-relevant models that are based on quasi-comprehensive administrative data and that can be used to explore alternative policy decisions. Such a model on a small scale was discussed—the POSEIDON model for fishery management. A larger model is the full-scale agent-based model of the U.S. housing market bubble circa the mid-1990s to the late 2000s, its bursting, and subsequent economic consequences, specifically as they played out in the Washington, D.C., area.[83] There are several efforts underway today to create agent-based models that are suitable for use as macroeconomic policymaking tools,[84] analogous to the role played by so-called dynamic stochastic general equilibrium models in central banks.[85] Given the accelerating developments in this field in the past decade, the coming one is sure to see systematic progress toward large-scale macroeconomic agent-based models that are grounded in microdata.

Agent-based models that are relevant to the world's great common-pool resource problems have begun to appear.[86] These will become larger—approaching a global scale—as data

---

82 Axtell and Guerrero, forthcoming, Chapter Twenty.

83 John Geanakoplos, Robert Axtell, J. Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang, "Getting at Systemic Risk via an Agent-Based Model of the Housing Market," *American Economic Review*, Vol. 102, No. 3, May 2012.

84 Herbert Dawid and Domenico Delli Gatti, "Agent-Based Macroeconomics," in Cars Hommes and Blake LeBaron, eds., *Handbook of Computational Economics*, Vol. IV, *Heterogeneous Agent Modeling*, Amsterdam: Elsevier, North-Holland, 2018.

85 See David Colander, Peter Howitt, Alan Kirman, Axel Leijonhufvud, and Perry Mehrling, "Beyond DSGE Models: Toward an Empirically Based Macroeconomics," *American Economic Review*, Vol. 98, No. 2, May 2008; Giorgio Fagiolo and Andrea Roventini, "Macroeconomic Policy in DSGE and Agent-Based Models," working paper, Paris, France: EconomiX, 2012; and Frank Smets and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, Vol. 97, No. 3, June 2007.

86 On common-pool resource problems, see Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*, New York: Cambridge University Press, 1990; Elinor Ostrom, Joanna Burger, Christopher B. Field, Richard B. Norgaard, and David Policansky, "Revisiting the Commons: Local Lessons, Global Challenges," *Science*, Vol. 284, No. 5412, April 9, 1999; and Elinor Ostrom, Roy Gardner, and James Walker, *Rules, Games, and Common-Pool Resources*, Ann Arbor, Mich.: University of Michigan Press, 1994. For examples of agent-based models that are relevant to these problems, see Marcin Czupryna,

increasingly roll in. Agent-based models of governance,[87] international migration,[88] and other multicountry phenomena will also grow in scale and scope with better data, perhaps with data acquired from remote sensing.[89]

In international relations, the role for models will grow to the extent that they outperform traditional adversary forecasting approaches, grounded as they are mainly in leader psychology and statistical assessments of resource availability.[90] Compare international relations to weather forecasting. Data on weather in the United States have been systematically gathered for a century, but it was not until approximately 1980 that computational weather models could outperform people who were experts on the historical data.[91] As better data on international political actors combine computationally with richer behavioral models, those countries that can synthesize agent-based models that predict the actions of neighboring countries will thrive in the international system. From this perspective, it is not a question of whether such models will appear but rather *when* the research establishment in some country will invest sufficient resources to make such a paradigm viable, after which other countries will play catch-up.

Thus, large-scale, high-fidelity agent-based models of real-world social phenomena are coming soon, rendered for both scientific (positive) and policy (normative) purposes. Such agent-based models will be in a perpetual state of evolution given the real-time flow of high-quality, high-frequency data. The results will be available in visual form and will perhaps even be sent to decisionmakers' phones, iPads, and other screens. In some domains in which ABM is being used that are partially relevant to the social sciences, such as forest fire manage-

Christian L. E. Franzke, Sascha Hokamp, and Jürgen Scheffran, "An Agent-Based Approach to Integrated Assessment Modelling of Climate Change," *Journal of Artificial Societies and Social Simulation*, Vol. 23, No. 3, 2020; M. D. Gerst, P. Wang, A. Roventini, G. Fagiolo, G. Dosi, R. B. Howarth, and M. E. Borsuk, "Agent-Based Modeling of Climate Policy: An Introduction to the ENGAGE Multi-Level Model Framework," *Environmental Modelling & Software*, Vol. 44, June 2013; Francesco Lamperti, Antoine Mandel, Mauro Napoletano, Alessandro Sapio, Andrea Roventini, Tomas Balint, and Igor Khorenzhenko, "Towards Agent-Based Integrated Assessment Models: Examples, Challenges, and Future Developments," *Regional Environmental Change*, Vol. 19, No. 3, 2019; and Varun Rai and Adam Douglas Henry, "Agent-Based Modelling of Consumer Energy Choices," *Nature Climate Change*, Vol. 6, No. 6, June 2016.

87 Nils B. Weidmann and Idean Salehyan, "Violence and Ethnic Segregation: A Computational Model Applied to Baghdad," *International Studies Quarterly*, Vol. 57, No. 1, March 2013.

88 Ali Mansoor and Bryce Quillin, *Migration and Remittances: Eastern Europe and the Former Soviet Union*, Washington, D.C.: World Bank, 2006.

89 Basudeb Bhatta, *Analysis of Urban Growth and Sprawl from Remote Sensing Data*, Berlin: Springer-Verlag, 2010; Ron Mahabir, Peggy Agouris, Anthony Stefanidis, Arie Croitoru, and Andrew T. Crooks, "Detecting and Mapping Slums Using Open Data: A Case Study in Kenya," *International Journal of Digital Earth*, Vol. 13, No. 6, 2020; and Seyed M. Mussavi Rizi, Maciej M. Latek, and Armando Geller, "Merging Remote Sensing Data and Population Surveys in Large, Empirical Multiagent Models: The Case of the Afghan Drug Industry," presented at the Third World Congress on Social Simulation, Kassel, Germany, 2010.

90 Ian S. Lustick and Philip E. Tetlock, "The Simulation Manifesto: The Limits of Brute-Force Empiricism in Geopolitical Forecasting," *Futures & Foresight*, Vol. 3, No. 2, June 2021.

91 J. Doyne Farmer, personal communication, from a conversation with Ed Lorenz at the National Center for Atmospheric Research.

ment and related areas of disaster mitigation and amelioration, models with these properties are already on display. Extension of such developments across all of the social sciences will lead to *computationally enabled policy* and should result in better management of the social and natural worlds, with large rewards for the country that can accomplish this first.

Concluding Thoughts

ABM has been around from the late 1980s to the early 1990s. Beginning with small-scale, abstract, and otherwise toy models, the method has now grown through advances in both hardware and software. As a result, ABM is now capable of rendering social phenomena at full scale—every person, every institution represented—and in deeply empirical ways, making systematic use of both microdata, when available, and aggregate data, either as input specifications or as target outputs. Overall, for those who look closely at the state of the technology, ABM has the potential to revolutionize the social sciences by facilitating the relaxation of unrealistic behavioral and structural specifications in conventional models. With continued advances in computing power, ever larger models based on more and more data, executing on increasingly parallel machines in less and less wall time, will progressively become a reality. In almost any business-as-usual scenario, such advances will occur and be game-changing in many fields.

Yet there are also opportunities for more-rapid, even more-guided, evolution of the ABM paradigm through specific hardware and software innovations, as I have sketched out. Different hardware—essentially uninterruptible processor cores that provide execution guarantees—could dramatically reduce run-times for large-scale agent-based models. Furthermore, the technical development of new parallel programming languages, frameworks, and software libraries would do much to advance ABM software development in the short term, while coupled hardware and software developments—perhaps in the guise of specialized boards or nodes for large-scale agent computation—could radically accelerate the entire paradigm over longer timescales. U.S. research institutions relevant to the social sciences are not well positioned to support such efforts, such as the creation of community agent-based models relevant to specific disciplines. It will probably take a farsighted administrator or an act of Congress to create the nucleus for such institutional support as is needed for the rapid expansion of perhaps the most revolutionary social science methodology since the appearance, in the waning days of World War II, of von Neumann's and Morgenstern's research on game theory,[92] work which subsequently served as the basis for much Cold War strategic theorizing.[93]

---

[92] John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton, N.J.: Princeton University Press, 1944.

[93] R. Duncan Luce and Howard Raiffa, *Games and Decisions: Introduction and Critical Survey*, New York: John Wiley & Sons, 1957; Thomas C. Schelling, *The Strategy of Conflict*, Cambridge, Mass.: Harvard University Press, 1960; Martin Shubik, "Bibliography on Simulation, Gaming, Artificial Intelligence and Allied Topics," *Journal of the American Statistical Association*, Vol. 55, No. 292, December 1960; and Martin Shubik, ed., *Game Theory and Related Approaches to Social Behavior*, New York: John Wiley & Sons, Inc., 1964.

The international system will remain highly complex and dynamic, populated by peer, near-peer, and asymmetric competitors, as well as by allies and partners. Meeting the demand for a new generation of analytic capabilities to better understand, engage, and compete over the short and long terms represents an application area that the emerging suite of tools, techniques, and technologies surrounding ABM is poised to fill. ABM will move beyond conventional game theory as the method of choice for representing, understanding, explaining, and forecasting the behavior of U.S. friends and foes alike. It is already being used all over the world for a variety of purposes.

The United States has been a leader in the application of ABM and has a large head start on other nations. However, many leading scientists in those countries also see the value in this new way of building models and are moving quickly to harness it for their own purposes—to better manage their own economies and strategic assets, but also, inevitably, to provide security from adversarial agents while fostering partnerships to shape institutions to their benefit. In the coming world in which each nation builds models of every other nation, which in turn model each of the nations modeling them, how the ecosystem of models interacts with policymakers' decisions is almost certainly impossible to forecast. Surely it will be a complex system with many surprising, emergent properties, and it will likely surprise us unless we engage with it fully.

Acknowledgments

I thank the editors of this report for their editorial support and encouragement.

Abbreviations

ABM	Agent-Based Modeling
CPU	central processing unit
GPU	graphical processing unit
MASON	Multi-Agent Simulation of Neighborhoods or Networks
OOP	object-oriented programming
RAM	random access memory
ZI	zero intelligence

References

Aleta, Alberto, David Martín-Corral, Ana Pastore y Piontti, Marco Ajelli, Maria Litvinova, Matteo Chinazzi, Natalie E. Dean, M. Elizabeth Halloran, Ira M. Longini, Jr., Stefano Merler, et al., "Modelling the Impact of Testing, Contact Tracing and Household Quarantine on Second Waves of COVID-19," *Nature Human Behaviour*, Vol. 4, No. 9, September 2020, pp. 964–971.

Arthur, W. Brian, "Inductive Reasoning and Bounded Rationality," *American Economic Review*, Vol. 84, No. 2, May 1994, pp. 406–411.

Atalay, Enghin, Ali Hortaçsu, James Roberts, and Chad Syverson, "Network Structure of Production," *Proceedings of the National Academy of Sciences*, Vol. 108, No. 13, 2011, pp. 5199–5202.

Athey, Susan, "Beyond Prediction: Using Big Data for Policy Problems," *Science*, Vol. 355, No. 6324, February 3, 2017, pp. 483–485.

Axtell, Robert L., "The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions," working paper, Washington, D.C.: Brookings Institution, 1999.

Axtell, Robert L., "Effects of Interaction Topology and Activation Regime in Several Multi-Agent Systems," in Scott Moss and Paul Davidsson, eds., *Multi-Agent-Based Simulation*, Vol. 1979, Heidelberg, Germany: Springer Verlag, 2001a, pp. 33–48.

Axtell, Robert L., "Zipf Distribution of U.S. Firm Sizes," *Science*, Vol. 293, No. 5536, September 7, 2001b, pp. 1818–1820.

Axtell, Robert L., "What Economic Agents Do: How Cognition and Interaction Lead to Emergence and Complexity," *Review of Austrian Economics*, Vol. 20, No. 2, 2007, pp. 105–122.

Axtell, Robert L., "120 Million Agents Self-Organize into 6 Million Firms: A Model of the U.S. Private Sector," in J. Thangarajah, K. Tuyls, C. Jonker, and S. Marsella, eds., *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*, Singapore, May 9, 2016a, pp. 806–816.

Axtell, Robert L., "Hayek Enriched by Complexity Enriched by Hayek," *Advances in Austrian Economics*, Vol. 21, November 30, 2016b, pp. 63–121.

Axtell, Robert L., "Endogenous Firm Dynamics and Labor Flows via Heterogeneous Agents," in Cars Hommes and Blake LeBaron, eds., *Handbook of Computational Economics*, Vol. IV, *Heterogeneous Agent Modeling*, Amsterdam: Elsevier, North-Holland, 2018, pp. 157–213.

Axtell, Robert L., and J. Doyne Farmer, "Agent-Based Modeling in Economics and Finance: Past, Present, and Future," *Complexity Handbook*, Stuttgart, Germany: University of Hohenheim, August 1, 2018.

Axtell, Robert L., and Omar A. Guerrero, *Dynamics of Firms from the Bottom Up: Data, Theories, and Models*, Cambridge, Mass.: MIT Press, forthcoming.

Axtell, Robert L., Alan Kirman, Iain D. Couzin, Daniel Fricke, Thorsten Hens, Michael E. Hochberg, John E. Mayfield, Peter Schuster, and Rajiv Sethi, "Challenges of Integrating Complexity and Evolution into Economics," in David S. Wilson and Alan Kirman, eds., *Complexity and Evolution: Toward a New Synthesis for Economics*, Cambridge, Mass.: MIT Press, 2016, pp. 65–81.

Bailey, Richard M., Ernesto Carrella, Robert Axtell, Matthew G. Burgess, Reniel B. Cabral, Michael Drexler, Chris Dorsett, Jens Koed Madsen, Andreas Merkl, and Steven Saul, "A Computational Approach to Managing Coupled Human–Environmental Systems: The POSEIDON Model of Ocean Fisheries," *Sustainability Science*, Vol. 14, No. 2, 2019, pp. 259–275.

Barrett, Chris, and Richard Beckman, *TRANSIMS: Portland Case Study Report*, Vol. I, *Introduction and Overview*, Los Alamos, N.M.: Los Alamos National Laboratory, 1995.

Bhatta, Basudeb, *Analysis of Urban Growth and Sprawl from Remote Sensing Data*, Berlin: Springer-Verlag, 2010.

Bordini, Rafael H., Jomi F. Hübner, and Renata Vieira, "Jason and the Golden Fleece of Agent-Oriented Programming," in Rafael H. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah Seghrouchni, eds., *Multi-Agent Programming: Languages, Platforms and Applications*, New York: Springer Science+Business Media, Inc., 2005, pp. 3–37.

Borges, Jorge Luis, *Labyrinths*, New York: New Directions, 1962.

Brock, William A., and Cars H. Hommes, "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model," *Journal of Economic Dynamics and Control*, Vol. 22, Nos. 8–9, July 1998, pp. 1235–1274.

Brooks, Rodney A., and Pattie Maes, eds., *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, Mass.: MIT Press, 1994.

Cámara, Javier, Kirstie L. Bellman, Jeffrey O. Kephart, Marco Autili, Nelly Bencomo, Ada Diaconescu, Holger Giese, Sebastian Götz, Paola Inverardi, Samuel Kounev, and Massimo Tivoli, "Self-Aware Computing Systems: Related Concepts and Research Areas," in Samuel Kounev, Jeffrey O. Kephart, Aleksandar Milenkoski, and Xiaoyun Zhu, eds., *Self-Aware Computing Systems*, Cham, Switzerland: Springer Nature, 2017.

Cámara, Javier Palanca, Marti Navarro, Estefania Argente, Ana Garcia-Fornes, and Vicente Julián, "Modeling an Operating System Based on Agents," in Emilio Corchado, Václav Snášel, Ajith Abraham, Michal Wozniak, Manuel Graña, and Sung-Bae Cho, eds., *Hybrid Artificial Intelligent Systems: 7th International Conference, HAIS 2012 Proceedings*, Pt. 1, Salamanca, Spain, March 28–30, 2012, pp. 588–599.

Carrella, Ernesto, Steven Saul, Kristin Marshall, Matthew G. Burgess, Reniel B. Cabral, Richard M. Bailey, Chris Dorsett, Michael Drexler, Jens Koed Madsen, and Andreas Merkl, "Simple Adaptive Rules Describe Fishing Behaviour Better Than Perfect Rationality in the US West Coast Groundfish Fishery," *Ecological Economics*, Vol. 169, March 2020.

Chandra, Rohit, Leonardo Dagum, Dave Kohr, Dror Maydan, Jeff McDonald, and Ramesh Menon, *Parallel Programming in OpenMP*, San Francisco, Calif.: Morgan Kaufmann Publishers, 2001.

Chapman, Barbara, Gabriele Jost, and Ruud van der Pas, *Using OpenMP: Portable Shared Memory Parallel Programming*, Cambridge, Mass.: MIT Press, 2007.

Ciancarini, Paolo, and Michael J. Wooldridge, eds., *Agent-Oriented Software Engineering*, Vol. 1957, New York: Springer, 2000.

Clark, Colin W., *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*, 2nd ed., New York: Wiley Inter-Science, 2005.

Clark, Colin W., *The Worldwide Crisis in Fisheries: Economic Models and Human Behavior*, New York: Cambridge University Press, 2006.

Clark, Robin, and Steven O. Kimbrough, "The Spontaneous Emergence of Language Variation from a Homogeneous Population," presented at the Computational Social Science Society of America, Santa Fe, N.M., 2015.

Cliff, Dave, "Methodological Mess-Ups in Modelling Markets with Minimal-Intelligence Agents," presentation at the First Conference on Zero/Minimal Intelligence Agents, New Haven, Conn.: Yale University, October 23, 2020.

Colander, David, Peter Howitt, Alan Kirman, Axel Leijonhufvud, and Perry Mehrling, "Beyond DSGE Models: Toward an Empirically Based Macroeconomics," *American Economic Review*, Vol. 98, No. 2, May 2008, pp. 236–240.

Comer, Kenneth W., *Who Goes First? An Examination of the Impact of Activation on Outcome Behavior in Agent-Based Models*, dissertation, Fairfax, Va.: George Mason University, 2014.

Connable, Ben, "Authentically Describing and Forecasting Human Behavior for Policy Analysis: A Review and a Path Forward," in Aaron B. Frank and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches,* Santa Monica, Calif.: RAND Corporation, RR-A1275-1, 2022. As of July 2022: https://www.rand.org/pubs/research_reports/RRA1275-1.html

Cooper, David J., and John H. Kagel, "Other-Regarding Preferences: A Selective Survey of Experimental Results," in John H. Kagel and Alvin E. Roth, *The Handbook of Experimental Economics*, Vol. 2, Princeton, N.J.: Princeton University Press, 2015, pp. 217–289.

Crooks, Andrew, and Ates Hailegiorgis, "Disease Modeling Within Refugee Camps: A Multi-Agent Systems Approach," in R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds., *2013 Winter Simulations Conference (WSC)*, Washington, D.C.: IEEE, 2013, pp. 1697–1706.

Czupryna, Marcin, Christian L. E. Franzke, Sascha Hokamp, and Jürgen Scheffran, "An Agent-Based Approach to Integrated Assessment Modelling of Climate Change," *Journal of Artificial Societies and Social Simulation*, Vol. 23, No. 3, 2020.

Davies, Martin, and Tony Stone, *Mental Simulation*, Cambridge, Mass.: Blackwell Publishers, 1995.

Dawid, Herbert, and Domenico Delli Gatti, "Agent-Based Macroeconomics," in Cars Hommes and Blake LeBaron, eds., *Handbook of Computational Economics*, Vol. IV, *Heterogeneous Agent Modeling*, Amsterdam: Elsevier, North-Holland, 2018, pp. 63–156.

D'Souza, R. M., M. Lysenko, and K. Rahmani, "SugarScape on Steroids: Simulation over a Million Agents at Interactive Rates," presented at AGENT 2007 Conference on Complex Interaction and Social Emergence, Evanston, Ill., 2007.

Duffy, John, "Agent-Based Models and Human Subject Experiments," in Leigh Tesfatsion and Kenneth L. Judd, eds., *Handbook of Computational Economics*, Vol. 2, *Agent-Based Methods*, New York: North-Holland, 2006, pp. 949–1011.

Edwards, Paul N., *The Closed World: Computers and the Politics of Discourse in Cold War America*, Cambridge, Mass.: MIT Press, 1996.

Einav, Liran, and Jonathan Levin, "Economics in the Age of Big Data," *Science*, Vol. 346, No. 6210, November 7, 2014.

Epstein, Joshua M., and Robert Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, Cambridge, Mass.: MIT Press, 1996.

Eubank, Stephen, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltán Toroczkai, and Nan Wang, "Modelling Disease Outbreaks in Realistic Urban Social Networks," *Nature*, Vol. 429, May 13, 2004, pp. 180–184.

Fagiolo, Giorgio, and Giovanni Dosi, "Exploitation, Exploration and Innovation in a Model of Endogenous Growth with Locally Interacting Agents," *Structural Change and Economic Dynamics*, Vol. 14, No. 3, 2003, pp. 237–273.

Fagiolo, Giorgio, and Andrea Roventini, "Macroeconomic Policy in DSGE and Agent-Based Models," working paper, Paris, France: EconomiX, 2012.

Farmer, J. Doyne, and Duncan Foley, "The Economy Needs Agent-Based Modeling," *Nature*, Vol. 460, No. 6, August 2009, pp. 685–686.

Farmer, J. Doyne, Paolo Patelli, and Ilija I. Zovko, "The Predictive Power of Zero Intelligence in Financial Markets," *Proceedings of the National Academy of Sciences*, Vol. 102, No. 6, February 8, 2005, pp. 2254–2259.

Feynman, Richard P., Robert B. Leighton, and Matthew Sands, *The Feynman Lectures on Physics*, Vol. II, *Mainly Electromagnetism and Matter*, Reading, Mass.: Addison-Wesley Publishing Company, 1964.

Fujiwara, Yoshi, Corrado Di Guilmi, Hideaki Aoyama, Mauro Gallegati, and Wataru Souma, "Do Pareto-Zipf and Gibrat Laws Hold True? An Analysis with European Firms," *Physica A: Statistical Mechanics and Its Applications*, Vol. 335, No. 1, April 1, 2004, pp. 197–216.

Gaffeo, Eduardo, Mauro Gallegati, and Antonio Palestrini, "On the Size Distribution of Firms: Additional Evidence from the G7 Countries," *Physica A: Statistical Mechanics and Its Applications*, Vol. 324, No. 1, June 1, 2003, pp. 117–123.

Geanakoplos, John, Robert Axtell, J. Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang, "Getting at Systemic Risk via an Agent-Based Model of the Housing Market," *American Economic Review*, Vol. 102, No. 3, May 2012, pp. 53–58.

Gelernter, David, *Mirror Worlds, or: The Day Software Puts the Universe in a Shoebox . . . How It Will Happen and What It Will Mean*, New York: Oxford University Press, 1992.

Georges, Christophre, and Javier Pereira, "Market Stability with Machine Learning Agents," *Journal of Economic Dynamics & Control*, Vol. 122, January 2021.

Gerst, M. D., P. Wang, A. Roventini, G. Fagiolo, G. Dosi, R. B. Howarth, and M. E. Borsuk, "Agent-Based Modeling of Climate Policy: An Introduction to the ENGAGE Multi-Level Model Framework," *Environmental Modelling & Software*, Vol. 44, June 2013, pp. 62–75.

Gilbert, Nigel, "Varieties of Emergence," in Charles Macal and David Sallach, eds., *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, Chicago: University of Chicago, October 2002, pp. 41–50.

Gode, Dhananjay K., and Shyam Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, Vol. 101, No. 1, February 1993, pp. 119–137.

Gode, Dhananjay K., and Shyam Sunder, "What Makes Markets Allocationally Efficient?" *Quarterly Journal of Economics*, Vol. 112, No. 2, May 1997, pp. 603–630.

Grieves, Michael, and John Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in Franz-Josef Kahlen, Shannon Flumerfelt, and Anabela Alves, eds., *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*, Switzerland: Springer Nature, 2017, pp. 85–114.

Gunaratne, Chathika, *Evolutionary Model Discovery: Automating Causal Inference for Generative Models of Human Social Behavior*, dissertation, Orlando, Fla.: University of Central Florida, 2019.

Guo, Jinzhong, Qi Xu, Qinghua Chen, and Yougui Wang, "Firm Size Distribution and Mobility of the Top 500 Firms in China, the United States and the World," *Physica A: Statistical Mechanics and Its Applications*, Vol. 392, No. 13, July 1, 2013, pp. 2903–2914.

Higdon, David M., Robert L. Axtell, Venkatramani Balaji, Lawrence E. Buja, Katherine V. Calvin, Kathleen M. Carley, Rebecca Castaño, Ronald R. Coifman, Omar Ghattas, James A. Hansen, et al., *From Maps to Models: Augmenting the Nation's Geospatial Intelligence Capabilities*, Washington, D.C.: The National Academies Press, 2016.

Hogeweg, P., and B. Hesper, "Individual-Oriented Modelling in Ecology," *Mathematical and Computer Modelling*, Vol. 13, No. 6, 1990, pp. 83–90.

Holland, John H., and John H. Miller, "Artificial Adaptive Agents in Economic Theory," *American Economic Review*, Vol. 81, No. 2, May 1991, pp. 365–370.

van der Hoog, Sander, *Deep Learning in (and of) Agent-Based Models: A Prospectus*, arXiv.org, June 20, 2017.

Huberman, Bernardo A., and Natalie S. Glance, "Evolutionary Games and Computer Simulations," *Proceedings of the National Academy of Sciences*, Vol. 90, No. 16, August 1993, pp. 7716–7718.

Jackson, Matthew O., *Social and Economic Networks*, Princeton, N.J.: Princeton University Press, 2008.

Jennings, Nicholas R., Francisco J. Garijo, and Magnus Boman, "Agent-Oriented Software Engineering," presented at the Multi-Agent System Engineering: 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Valencia, Spain, July 1999.

Kang, Sang Hoon, Zhuhua Jiang, Chongcheul Cheong, and Seong-Min Yoon, "Changes of Firm Size Distribution: The Case of Korea," *Physica A: Statistical Mechanics and Its Applications*, Vol. 390, No. 2, January 15, 2011, pp. 319–327.

Kiran, Mariam, Paul Richmond, Mike Holcombe, Lee Shawn Chin, David Worth, and Chris Greenough, "FLAME: Simulating Large Populations of Agents on Parallel Hardware Architectures," in Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, eds., *AAMAS '10: Proceedings of the 9th International Conference on Autonomous Agent and Multiagent Systems*, Toronto, Canada, 2010.

Kreindler, Gabriel E., and H. Peyton Young, "Fast Convergence in Evolutionary Equilibrium Selection," *Games and Economic Behavior*, Vol. 80, July 2013, pp. 39–67.

Krugman, Paul, *Peddling Prosperity: Economic Sense and Nonsense in the Age of Diminished Expectations*, New York: W. W. Norton & Company, 1994.

Kuhn, Thomas S., *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought*, New York: Vintage Books, 1957.

Lakatos, Imre, "Falsification and the Methodology of Scientific Research Programmes," in Imre Lakatos and Alan Musgrave, eds., *Criticism and the Growth of Knowledge*, New York: Cambridge University Press, 1970, pp. 91–196.

Lamperti, Francesco, Antoine Mandel, Mauro Napoletano, Alessandro Sapio, Andrea Roventini, Tomas Balint, and Igor Khorenzhenko, "Towards Agent-Based Integrated Assessment Models: Examples, Challenges, and Future Developments," *Regional Environmental Change*, Vol. 19, No. 3, 2019, pp. 747–762.

Lamperti, Francesco, Andrea Roventini, and Amir Sani, "Agent-Based Model Calibration Using Machine Learning Surrogates," *Journal of Economic Dynamics & Control*, Vol. 90, May 2018, pp. 366–389.

Langton, Christopher G., ed., *Santa Fe Institute Studies in the Sciences of Complexity*, Vol. VI, *Artificial Life*, Reading, Mass.: Addison-Wesley Publishing, 1989.

Langton, Christopher G., ed., *Santa Fe Institute Studies in the Sciences of Complexity*, Vol. XVII, *Artificial Life III*, Reading, Mass.: Addison-Wesley Publishing, 1994.

Langton, Christopher G., Charles Taylor, J. Doyne Farmer, and Steen Rasmussen, *Santa Fe Institute Studies in the Sciences of Complexity*, Vol. X, *Artificial Life II*, Redwood City, Calif.: Addison-Wesley Publishing, 1992.

Laughlin, R. B., David Pines, Joerg Schmalian, Branko P. Stojkovic, and Peter Wolynes, "The Middle Way," *Proceedings of the National Academy of Sciences*, Vol. 97, No. 1, January 4, 2000, pp. 32–37.

Lave, Charles A., and James G. March, *An Introduction to Models in the Social Sciences*, New York: Harper & Row, 1975.

Lee, Edward A., *The Problem with Threads*, Berkeley, Calif.: Electrical Engineering and Computer Sciences, University of California at Berkeley, January 10, 2006.

Luce, R. Duncan, and Howard Raiffa, *Games and Decisions: Introduction and Critical Survey*, New York: John Wiley & Sons, 1957.

Luke, Sean, *Multiagent Simulation and the MASON Library*, Fairfax, Va.: George Mason University, 2015.

Lustick, Ian S., and Philip E. Tetlock, "The Simulation Manifesto: The Limits of Brute-Force Empiricism in Geopolitical Forecasting," *Futures & Foresight*, Vol. 3, No. 2, June 2021.

Mahabir, Ron, Peggy Agouris, Anthony Stefanidis, Arie Croitoru, and Andrew T. Crooks, "Detecting and Mapping Slums Using Open Data: A Case Study in Kenya," *International Journal of Digital Earth*, Vol. 13, No. 6, 2020, pp. 683–707.

Mansoor, Ali, and Bryce Quillin, *Migration and Remittances: Eastern Europe and the Former Soviet Union*, Washington, D.C.: World Bank, 2006.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory*, New York: Oxford University Press, 1995.

McCabe, Stefan, Dale Brearcliffe, Peter Froncek, Marta Hansen, Vince Kane, Davoud Taghawi-Hejad, and Robert L. Axtell, "A Comparison of Languages and Frameworks for the Parallelization of a Simple Agent Model," working paper, Fairfax, Va.: Department of Computational and Data Sciences, George Mason University, forthcoming.

McFadden, Daniel, and Paul A. Ruud, "Estimation by Simulation," *Review of Economics and Statistics*, Vol. LXXVI, No. 4, November 1994.

Menger, Carl, *Investigations into the Method of the Social Sciences*, trans. Francis J. Nock, Grover City, Pa.: Libertarian Press, Inc., [1871] 1985.

Mirowski, Philip, *More Heat than Light: Economics as Social Physics, Physics as Nature's Economics*, New York: Cambridge University Press, 1989.

Morgan, Mary S., and Margaret Morrison, *Models as Mediators: Perspectives on Natural and Social Science*, New York: Cambridge University Press, 1999.

Nagel, Kai, and Christopher L. Barrett, "Using Microsimulation Feedback for Trip Adaptation for Realistic Traffic in Dallas," *International Journal of Modern Physics C*, Vol. 8, No. 3, 1997, pp. 505–525.

Nasar, Sylvia, *A Beautiful Mind: A Biography of John Forbes Nash, Jr., Winner of the Nobel Prize in Economics, 1994*, New York: Simon & Schuster, 1998.

Nash, John, *Parallel Control*, Santa Monica, Calif.: RAND Corporation, RM-1361, 1954. As of November 7, 2021:
https://www.rand.org/pubs/research_memoranda/RM1361.html

North, Michael J., and Charles M. Macal, *Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation*, New York: Oxford University Press, 2007.

Nowak, Martin A., and Robert M. May, "Evolutionary Games and Spatial Chaos," *Nature*, Vol. 359, October 29, 1992, pp. 826–829.

Nyman, Linus, and Mikael Laakso, "Notes on the History of Fork and Join," *IEEE Annals of the History of Computing*, Vol. 38, No. 3, July–September 2016, pp. 84–87.

Ostrom, Elinor, *Governing the Commons: The Evolution of Institutions for Collective Action*, New York: Cambridge University Press, 1990.

Ostrom, Elinor, Joanna Burger, Christopher B. Field, Richard B. Norgaard, and David Policansky, "Revisiting the Commons: Local Lessons, Global Challenges," *Science*, Vol. 284, No. 5412, April 9, 1999, pp. 278–282.

Ostrom, Elinor, Roy Gardner, and James Walker, *Rules, Games, and Common-Pool Resources*, Ann Arbor, Mich.: University of Michigan Press, 1994.

Quinn, Michael J., *Parallel Programming in C with MPI and OpenMP*, New York: McGraw-Hill, 2003.

Rai, Varun, and Adam Douglas Henry, "Agent-Based Modelling of Consumer Energy Choices," *Nature Climate Change*, Vol. 6, No. 6, June 2016, pp. 556–562.

Reinders, James, *Intel Threading Building Blocks: Outfitting C++ for Multi-Core Processor Parallelism*, Sebastopol, Calif.: O'Reilly, 2007.

Rizi, Seyed M. Mussavi, Maciej M. Latek, and Armando Geller, "Merging Remote Sensing Data and Population Surveys in Large, Empirical Multiagent Models: The Case of the Afghan Drug Industry," presented at the Third World Congress on Social Simulation, Kassel, Germany, 2010.

Rodriguez, Sebastian, Nicolas Gaud, and Stéphane Galland, "SARL: A General-Purpose Agent-Oriented Programming Language," presented at the 2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), Warsaw, Poland, August 11–14, 2014.

Rosenzweig, Roy, "Scarcity or Abundance? Preserving the Past in a Digital Era," *American Historical Review*, Vol. 108, No. 3, June 1, 2003, pp. 735–762.

Rupp, Karl, "Microprocessor Trend Data," data repository, last updated July 16, 2020. As of November 10, 2021:
https://github.com/karlrupp/microprocessor-trend-data

Schelling, Thomas C., *The Strategy of Conflict*, Cambridge, Mass.: Harvard University Press, 1960.

Scott, Stephen Lewis, *Computational Modeling for Marine Resource Management*, dissertation, Fairfax, Va.: George Mason University, 2016.

Shoham, Yoav, "Agent-Oriented Programming," *Artificial Intelligence*, Vol. 60, No. 1, March 1993, pp. 51–92.

Shoham, Yoav, "An Overview of Agent-Oriented Programming," in Jeffrey M. Bradshaw, ed., *Software Agents*, Menlo Park, Calif.: American Association for Artificial Intelligence/MIT Press, 1997, pp. 271–290.

Shubik, Martin, "Bibliography on Simulation, Gaming, Artificial Intelligence and Allied Topics," *Journal of the American Statistical Association*, Vol. 55, No. 292, December 1960, pp. 736–751.

Shubik, Martin, ed., *Game Theory and Related Approaches to Social Behavior*, New York: John Wiley & Sons, Inc., 1964.

Simon, Herbert A., "Giving the Soft Sciences a Hard Sell," *Boston Globe*, May 3, 1987.

Simon, Herbert A., and Charles P. Bonini, "The Size Distribution of Business Firms," *American Economic Review*, Vol. 48, No. 4, September 1958, pp. 607–617.

Smets, Frank, and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, Vol. 97, No. 3, June 2007, pp. 586–606.

Smith, Vernon L., "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, Vol. 70, No. 2, April 1962, pp. 111–137.

Stanley, Michael H. R., Luís A. N. Amaral, Sergey V. Buldyrev, Shlomo Havlin, Heiko Leschhorn, Philipp Maass, Michael A. Salinger, and H. Eugene Stanley, "Scaling Behaviour in the Growth of Companies," *Nature*, Vol. 379, No. 29, February 1996, pp. 804–806.

Tenner, Edward, *The Efficiency Paradox: What Big Data Can't Do*, New York: Alfred A. Knopf, 2018.

von Neumann, John, and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton, N.J.: Princeton University Press, 1944.

Vu, Tuong Manh, Charlotte Probst, Joshua M. Epstein, Alan Brennan, Mark Strong, and Robin C. Purshouse, "Toward Inverse Generative Social Science Using Multi-Objective Genetic Programming," presented at the Genetic and Evolutionary Computation Conference (GECCO '19), Prague, Czech Republic, July 13–17, 2019.

Walras, Léon, *Éléments d'Économie Politique Pure, Édition Définitive, Revue et Augmentée*, Paris: Pichon et Durand-Anzias, 1926.

Weidmann, Nils B., and Idean Salehyan, "Violence and Ethnic Segregation: A Computational Model Applied to Baghdad," *International Studies Quarterly*, Vol. 57, No. 1, March 2013, pp. 52–64.

Wellman, Michael P., "Market-Oriented Programming: Some Early Lessons," in Scott H. Clearwater, ed., *Market-Based Control: A Paradigm for Distributed Resource Allocation*, Singapore: World Scientific, 1996.

Wellman, Michael P., Amy Greenwald, and Peter Stone, *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*, Cambridge, Mass.: MIT Press, 2007.

Wilensky, Uri, and William Rand, *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with NetLogo*, Cambridge, Mass.: MIT Press, 2015.

Winkelman, Zev, "Using Technology to Improve the Agility of Force Generation Processes," in Aaron B. Frank and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, Santa Monica, Calif.: RAND Corporation, RR-A1275-1, 2022. As of July 2022:
https://www.rand.org/pubs/research_reports/RRA1275-1.html

Zhang, Jianhua, Qinghua Chen, and Yougui Wang, "Zipf Distribution in Top Chinese Firms and an Economic Explanation," *Physica A: Statistical Mechanics and Its Applications*, Vol. 388, No. 10, May 15, 2009, pp. 2020–2024.