

Aggregation in agent-based models of economies

SCOTT E. PAGE

Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI 48106, USA;
e-mail: spage@umich.edu

Abstract

Agent-based models are often described as bottom-up because macro-level phenomena emerge from the micro-level interactions of agents. These macro-level phenomena include fixed points, cycles, dynamic patterns, and long transients. In this paper, I explore the link between micro-level characteristics—learning rules, diversity, network structure, and externalities—and the macro-level patterns they produce. I focus on why we need agent-level modeling, on how these models produce emergent phenomenon, and on how agent-based models help understand outcomes of social systems in a way that differs from the analytic, equilibrium approach.

1 Introduction

Economics resides within the social sciences. That means that its focus lies not on the individual *per se* but on the interactions between individuals and on how those interactions aggregate. Economic interactions comprise a range of activities from exchanges of work between friends to impersonal purchases over the Internet. To make sense of how economic behavior aggregates, a common conceit in economics involves a representative agent. If this agent is truly representative, then its behavior will be reflected in the economy writ large. If a tax increase causes the representative agent to work less, then we can conclude the same holds economy wide.

To many, the idea that such a representative agent exists seems a dubious proposition (Kirman, 1992). That skepticism can be broadened to apply to science writ large. In the words of the physicist Phil Anderson ‘more is different’ (Anderson, 1972). What Anderson meant was that in physical, biological, and social systems, aggregate phenomena can exhibit properties quite distinct from the properties of their parts. Single cell organisms exhibit behaviors far different from those of electrons and atoms, people perform a range of functions that far outstrip the capabilities of the cells that comprise them, and organizations and societies produce products, movements, and cultures that could not be the property of a single individual.

Aggregate phenomena that exhibit these unanticipated properties are commonly referred to as emergent. The wonder that arises when introduced to the concept of emergence should not discourage less romantic, scientific explorations into how and why the whole differs from its parts. In social systems, as in physical and chemical systems, we cannot nor should not assume that what happens at the macro level can be determined by the properties of the human agents alone. We must also take into account their interactions (Howitt, 2006).

In this article, I provide a brief introduction into how systems of interacting human agents aggregate. I focus on agent-based models (Tesfatsion, 1997; Bankes, 2002). These models consist of rule-following agents interacting in space and time. Introductory accounts describe four core characteristics of agent-based models: learning adaptation, geography networks, externalities, and diversity (Epstein & Axtell, 1996; de Marchi, 2005; Miller & Page, 2007). These same characteristics are evident in everyday life. People differ from one another. So do firms. People and firms

also adapt to their environments, which are often local and network based. And finally, firms and people interact locally and through networks, and the actions they take influence the payoffs and actions of others.

When considering the aggregation of agent-based models, combinations of these characteristics matter as much or more than the characteristics themselves. For example, adding diversity often has slight implications for aggregation as does adding networks, but adding both diversity and networks to a system tends to have substantial implications.

Aggregate behavior in an agent-based model can take many forms. It can result in equilibria, it can produce cycles and patterns, it can produce bubbles and crashes, and it can sometimes produce chaos. Most agent-based models produce multiple types of aggregate phenomena. Locally, a model might produce patterns, but globally the same model might produce an equilibrium. In some cases, we can determine whether an agent-based model will result in an equilibrium, a pattern, or chaos, but often the only way to determine the outcome is to construct the model in a computational platform, run it, and see what arises. Even computational implementation can be problematic. A model implemented on one platform may produce different results when implemented on another (Axelrod *et al.*, 1996).

In this manuscript I explore the generative nature of agent-based models by focusing on three themes. First, I show that aggregation models that do not take an agent-based approach often lose information that can result in errors, especially when a model is complex. Second, I discuss emergent aggregate phenomena. Third, I discuss the potential for agent-based models to reveal basins of attraction when systems possess a variety of outcomes.

2 Aggregate models vs. agent-based models

Agent-based models include substantial micro-level detail. It is natural to ask: ‘How much micro-level detail is necessary?’ To begin to answer that question, I want to consider a simple discrete time model of aggregation. At time step t , each agent can be thought of as being in a *state*. The state of agent i at time t , denoted by x_i^t , consists of all relevant information, including the agent’s location and direction, preferences, information, wealth, and abilities. Therefore, if the model contains n agents, all relevant information can be written as a vector of the agents’ states. This vector of states changes as a function of the model. Let F denote the *agent-level mapping* from the agents’ states at time t to their states at time $t + 1$ ¹. This agent-level mapping can be deterministic or stochastic.

Rather than consider the state of each agent, we could construct an *aggregation operator* A that maps the vector of the agents’ states at time t into an aggregate variable Y^t or vector of variables \bar{Y}^t . In the latter case, the dimension of the vector must be strictly lower than the sum of the dimensions of the agents’ state vectors. We could then construct an *aggregate-level mapping* H that maps the aggregate variable at time t to its new state at time $t + 1$. What I have described so far can be characterized in the following diagram adapted from Axtell and McRae (personal communication) and Iwasa *et al.* (1987).

This diagram shows a logic for when we need the agent-based particulars and when we do not. First, we must care about Y^t and not about the vector $(\bar{x}_1^t, \bar{x}_2^t, \dots, \bar{x}_n^t)$. Second, the diagram must *commute*: $H[A(x_1^t, x_2^t, \dots, x_n^t)] = A[F(x_1^t, x_2^t, \dots, x_n^t)]$. In a dynamical systems framework, conditions under which this diagram commutes can be found in Iwasa *et al.* (1987).

Many agent-based models do aggregate. Consider a model of consumer demand for a product in which each agent’s state can be written as a one-dimensional real number that denotes the agent’s income. Assume that each agent’s income increases by one unit in each time step and that agent i ’s demand for the product equals $a_i + bx_i^t$. The function that describes the change in the agents’ states is given by $F(x_1^t, x_2^t, \dots, x_n^t) = (x_1^t + 1, x_2^t + 1, \dots, x_n^t + 1)$. Let Y^t denote aggregate income, that is, the sum of the agents’ incomes. A straightforward calculation shows

¹ In a more general model, the mapping would depend on time and be denoted by F_t .




Figure 1 The aggregation diagram




Figure 2 An initial state of the Voter model

that the diagram commutes. Let $A(\vec{x}_1^t, \vec{x}_2^t, \dots, \vec{x}_n^t) = \sum_{i=1}^n \vec{x}_i^t$, and $H(Y^t) = Y^t + n$. It follows that $H(Y^t) = A[F((\vec{x}_1^t, \vec{x}_2^t, \dots, \vec{x}_n^t))]$. To calculate the sum of the agents' demands, we can use either the agent-based model or the aggregate model. In the former case, we just sum the individual demands of the agents. In the latter case, aggregate demand in time t equals $A + bY^t$, where $A = \sum_{i=1}^n a_i$.

In this example, agents spend a fixed proportion of their income on the product. Therefore, the entire population of agents also spends a fixed proportion of income on the product. And, as a result, the aggregate model works perfectly. This condition—identical proportional behavior across agents—allows aggregation to work.

In agent-based models, agents often interact over networks and their actions depend on the states of their neighbors. These feedbacks stifle attempts to capture model dynamics with an aggregate model. Consider the following canonical complex systems model known as the Voter model. This model includes both networks and interactions. In this simple version, five agents are arranged in a circle. Each agent has two *neighbors*: one on the left and one on the right. In the first period, an agent is assigned either action 1 (denoted by a solid disc) or action 0 (denoted by a circle). Its action in subsequent periods remains the same unless both of its neighbors take the opposite action, in which case it switches its action to match that of its neighbors. Figure 2 below shows a possible initial state.

The state of the five agents at time t can be represented by a binary vector of length five. Suppose that we employ an aggregation function A so that Y^t equals the sum of the agents' states so that $A(0,0,1,1,0) = 2$ and $A(1,1,1,0,1) = 4$. For $Y^t \leq 1$ and $Y^t \geq 4$ an aggregate-level mapping, H , exists which makes the diagram in Figure 1 commute. We need only set $H(Y^t) = 0$ if $Y^t \leq 1$ and $H(Y^t) = 5$ if $Y^t \geq 2$. The problem arises for Y^t in the set $\{2,3\}$. Up to symmetry there are two possible configurations for Y^t equals 2. One is shown in Figure 2. The other is shown in Figure 3.

In the Voter model depicted in Figure 2, in the next period, exactly one agent will be in state 1 (the agent surrounded by two agents in state 1). All other agents will either remain in state 0 or move to state 0. To match this instance of the agent-based model, the aggregate model would have




Figure 3 A second possible initial state of the Voter model




Figure 4 A 'blinking' state in a six-agent Voter model

to be such that $H(2) = 1$. But the initial states depicted in Figure 3 are an equilibrium. According to the rules of the model, no agent will change its action. To match this instance of the agent-based model, the aggregate model would have to be such that $H(2) = 2$. Therefore, the aggregate model cannot possibly be correct in all circumstances, which is why the agent-based model is needed.

Although ostensibly a stylized example of how aggregation can fail, the Voter model has empirical relevance. Let an agent in state 1 denote a criminal and an agent in state 0 denote a law-abiding citizen. Suppose that we have data on age, income, years of schooling, etc. that enable us to predict with some success the likelihood that someone will initially turn to crime. But suppose that the agents' decisions in subsequent periods depend on the actions of their friends (represented by neighbors in our graph). If so, to predict the aggregate level of crime it may not be sufficient to just look at averages. The network configuration would matter. Glaeser *et al.* (1996) construct just such a model of criminal activity and subject it to empirical testing. They find that crime data reveals these clusters which suggest the empirical relevance of the kind of spillovers assumed in the Voter model.

Expanding the Voter model to six agents produces the possibility of an equilibrium in the aggregate model but a cycle in the agent-based model. Consider the following initial actions by the agents (Figure 4).

This initial configuration of states will create a 'blinker'—every agent who was in state 1 will move to state 0 and every agent who was in state 0 will move to state 1. But, at the aggregate level, the total number of agents in each state remains unchanged. Thus, aggregate-level equilibrium masks agent-level dynamics (Epstein and Axtell, 1996).

As agent-based models become more complex—as they include greater diversity, more sophisticated behavioral rules, more complicated feedback structures, and more subtle network effects—the ability of aggregate models to trace the trajectory of states becomes limited. That is not to say that aggregate models and systems dynamics models have no use. An analysis of their

value lies outside the scope of this article. Nor should one infer that agent-based models offer a panacea. Agent-based models require substantial amounts of data as well as knowledge of how actors respond to their information, both of which may be in short supply. The take-away from this analysis is that aggregation implies *information loss*. The extent of that information loss increases as models become more complex.

3 Aggregation failure and emergence

Agent-based models are well suited for modeling complex systems. One of the most interesting, perhaps *the* most interesting feature of complex systems is that they can produce emergent phenomena. Bedau (1991) describes emergent phenomena as constructed of and generated by micro-level processes but as simultaneously autonomous from those processes. In other words, the aggregate output from a complex system can categorically differ from the parts that comprise it. This can be in form—such as the flocking of birds or two cycles that was produced earlier in this paper in the Voter model—or it can be in function—the ability of collections of neurons to make calculations. This potential for getting some leverage on the concept of is one of the most captivating features of agent-based models.

A natural starting point for thinking about emergence is a philosophical construct known as the principle of the *fallacy of composition*, which states that it is a mistake to infer that something is true of the whole because it is true of some part, or perhaps even every part of the whole. The fallacy refers to ‘truth’ perhaps owing to its philosophical foundations. The social sciences focus less on truth than on specific properties such as efficiency or fairness. Therefore, it will be helpful to decompose the fallacy into two separate phenomena. In the first case, a property can exist in the parts but not in the whole. In the latter case, a property can fail to exist in the parts but emerge in the aggregate.

*A system exhibits **aggregation failure** if every agent exhibits a property P but the aggregate does not exhibit that property.*

The canonical example of aggregation failure in social systems is the Condorcet triple. Each of three voters has transitive preferences. Voter 1 prefers A to B to C, and given transitivity, A to C. Voter 2 prefers B to C to A. And, voter 3 prefers C to A to B. If we define the aggregate’s preferences by majority rule, then the aggregate prefers A to B, B to C, and C to A. Therefore, although each individual satisfies the transitivity property, the aggregate does not.

Alternatively, a system can produce *emergent* properties.

*A system produces an **emergent property** if no agent exhibits a property P but the aggregate does (Axtell, 2003).*

In some cases, emergent properties *cannot* exist at the micro level. It is not possible for a single neuron to be conscious. Contrast this with the aggregation failure example of transitivity in preferences. In that case, transitivity can hold, at least in theory, at both levels. Often, although emergent phenomena could exist at the micro level, they do not. Epstein and Axtell (1996) construct a model in which agents who can only move horizontally and vertically form an aggregate that moves diagonally.

In markets, we can think of efficiency as an emergent phenomena. Suppose that we create an economy with equal numbers of two types of agents: one type has 10 apples, the other type has 10 bananas. Assume that every agent has the same preferences given by the utility function $U(A,B) = A^{0.5}B^{0.5}$. Given these preferences, an agent would be happiest with equal numbers of apples and bananas, but an agent would prefer any combination of apples and bananas to having all of one type of fruit. Consider an economy in which agents form random pairs, one of each pair makes random trade offers, and the trade is executed if both prefer the trade. Agents using this trading rule are called *zero-intelligent agents* (Gode and Sunder, 1993). The result of such a model will be that *on average* apples will be traded for bananas in equal quantity. Thus, the market price will be *efficient*. This efficiency arises even though none of the agents optimize, nor do they act




Figure 5 A two-dimensional Voter model

‘as if’ they are optimizing. They are following relatively simple rules. A similar phenomena exists in what is called the *wisdom of the multitudes*. Collections of individuals none of whom have the correct predictive model can combine to make accurate assessments².

In markets and aggregate forecasts, the collective can prove more capable than the individuals that comprise it. That said, the juxtaposition of micro-level random behavior and the emergence of an efficient average price is surprising. As is the combination of multiple, diverse models into an accurate collective prediction. It is this surprise factor that underpins informal characterizations of emergence.

The Condorcet triple and the market with zero-intelligent agents provide a best- and worst-case scenario from an institutional design perspective (Reiter, 1977). On the one hand, we want to avoid creating institutions in which good properties of individuals, for example, rationality, get lost in the aggregation such as was the case in the Condorcet triple. On the other hand, we want to find a way to construct institutions in which desirable aggregate properties, for example, efficiency, emerge even though none of the agents possesses that property.

3.1 Emergent spatial and dynamic patterns

Agent-based models are particularly adept at helping us to understand spatial patterns and distributions. As a first example of a spatial pattern, consider a two-dimensional version of the Voter model laid out on a checkerboard as shown in Figure 5.

In this model, assume that the set of *neighbors* of an agent are those agents that lie to the North, South, East, and West of the Agent. To create symmetry, assume that the North neighbor of an agent in the top row is the corresponding agent on the bottom row (and vice versa) and that the East neighbor of an agent in the right-most column is the corresponding agent in the left-most column (and vice versa). Each agent then has four neighbors. Assume that agents choose the action of a majority of their four neighbors, but stick with their own action in case of a tie. Applying this rule to the configuration shown in Figure 5 gives the following equilibrium configuration (Figure 6).

This pattern was not built into the model but emerges from the agents’ rules. In this example, the spatial patterning has greater influence on the equilibrium configuration than the initial number of agents in state 1. Notice that in the initial configuration (Figure 5) 22 of the 49 agents were in state 1 and that in the equilibrium configuration 28 of the agents were in state 1. Therefore, this instance of the model moves from a configuration in which a minority of the agents choose state 1 to a configuration in which a majority choose that state. Given that agents match the majority action in their neighborhood, this outcome is surprising. But the spatial correlations trump the averages. In fact, in the six periods the model takes to attain the equilibrium state, the number of agents in state 1 takes the following non-monotonic sequence: 22, 22, 21, 22, 24, 26, 28.

² See Page (2008) for a survey.




Figure 6 An equilibrium configuration in a two-dimensional Voter model

The pattern in Figure 6 is static. Agent-based models can also produce dynamic patterns. Wolfram (2001) showed how two-state cellular automata can produce random patterns as well as patterns that perform computations over time. The capability of agent-based models to produce such patterns hints at their potential. Two of the most important phenomena in economics—business cycles and stock market crashes—can be thought of as patterns. Sadly, neither is fully understood. This may be partly because equilibrium models, which assume that all forces in the economy balance, are ill-suited to explaining either phenomena.

Agent-based models have been constructed that produce business cycles (Nirei, 2006) as well as stock market crashes (LeBaron, 2001, 2006). These models can also produce other phenomena, such as clustered volatility in prices (Arthur *et al.*, 1997), which also exist in the data. The emergence of clustered volatility is often unavoidable. In an agent-based model, agents respond to changes in their local environments. As their environments become volatile, more agents have incentives to change their actions: adaptation begets adaptation. In contrast, in environments which have settled down, few agents may have incentives to change what they are doing. Therefore, a phenomenon that seems puzzling from an equilibrium perspective may be natural in an agent-based context.

Business cycles and crashes are aggregate dynamic phenomena that occur only if certain combinations of assumptions are built into or emerge from the model. Cycles can arise from the accumulation of lags through feedbacks. Models that include hierarchies of agents such as networks models of suppliers appear more capable of producing cycles than are models in which agents just make trades. Crashes, in contrast, can arise either from lack of heterogeneity (LeBaron, 2006) or from cascading failures (Bak, 1996).

That agent-based models produce both static and dynamic patterns is not by itself a singular accomplishment. Mathematical models can produce them. So can graphical models and chemical models. An advantage of agent-based models although is that they enable us to connect micro-level assumption to pattern formations. It is often the case that the emergence of a pattern depends on particulars of the model such as the structure of the network and timing assumptions. Agent-based models provide us a laboratory for exploring how various micro-level assumptions interact to produce macro-level patterns.

Consider timing assumptions for just a moment. In the Voter model described above, the agents update simultaneously in discrete time steps, not unlike a marching band. We might instead construct a model in which agents update asynchronously. In many contexts, this seems the more natural assumption. Huberman and Glance (1993) show that some emergent spatial and dynamic patterns disappear if updating is randomly asynchronous. In models with purposive agents, we might expect the order of updating might be based on incentives. In a real-world Voter model, people who differ entirely from their neighbors might update more quickly than those surrounded by a slim majority (Page, 1997). For the purposes of this survey, the particulars of how updating influences outcomes matters less than the general insight that *it does matter*, as do networks, as do learning rules. Agent-based models help us to see how.

3.2 Emergent distributions

Agent-based models also produce emergent distributions. In quote–unquote normal science we expect distributions to be normal (if idiosyncratic effects are additive) or log-normal (if idiosyncratic effects are multiplicative). Yet, in many agent-based models, the distributions that emerge satisfy *power laws*. In a power-law distribution, the probability that a random variable x equals some value v is proportional to v^{-k} . In the special case of k equal to 1, this means that events of size one hundred are one hundredth as likely as events of size one. If outcomes satisfy a power law, we should expect many many small outcomes and infrequent huge outcomes. The most famous of these models are Per Bak’s sandpile model, which produces avalanches with a power-law distribution (Bak, 1996) and Herb Simon’s preferential attachment model (Simon, 1955), which produces city sizes that satisfy a power law. In Bak’s model, the sand pile reaches a state of *self-organized criticality*, in which small events can trigger huge avalanches.

Many real-world phenomena—such as the sizes of traffic jams, stock market crashes, war deaths, network connections on the World Wide Web, and firms sizes (Axtell, 2001)—satisfy power-law distributions. Even patent production in cities has been shown to be a power law (Bettencourt *et al.*, 2007). Agent-based models have been able to reproduce those distributions (Axtell, 1999; Barabasi, 2002). It is important to note that the power law emerges from the model, that it emerges from the interactions of the agents.

In addition to being of scientific interest, these power laws often have substantial social consequences. Large events—be they market crashes or traffic pile ups—exact massive social costs. Models that produce power-law behavior may not only help us understand how they arise, they may point to interventions that prevent their occurrence.

4 Equilibrium, stability, basins of attraction, and path dependence

As mentioned several times in this article, neoclassical economic models rely on equilibrium as the fundamental solution concept. Set aside for the moment the fact that complexity and not equilibrium may be the more salient feature of most economies. Let us accept equilibrium as something that happens in some economic situations. We can then build from Wolfram (2001) and think of an agent-based model as making a calculation. If we restrict attention to agent-based models that attain equilibrium, we can think of them as calculating an equilibrium of the model. In one of the earliest models, Miller (1996) showed how genetic algorithms can produce equilibria not unlike those seen in experiments with human subjects.

Agent-based modelers highlight the fact that the equilibria located by agent-based models are generated (Epstein, 2003), not derived. To understand why, it is helpful to remind ourselves of the distinction between the *existence* of an equilibrium, the *stability* of an equilibrium, and an equilibrium’s *basin of attraction*. Place a marble in a bowl and the force of gravity will pull it to the bottom. The ball at rest is in a stable equilibrium that has the entire bowl as its basin of attraction. Therefore, we would say that the equilibrium is *globally stable*. Conversely, imagine dropping the same marble above an egg carton. The egg carton would have a dozen stable equilibria. But those equilibria would be only *locally stable* because a moderately firm push on the marble would send it from one well into another. Finally, if we turned our bowl upside down and delicately balanced the ball on what was the bottom of the bowl, we did have an *unstable equilibrium*. Even the slightest perturbation would cause the ball to roll down the side of the bowl.

In a physical system, the stability of equilibria and their basins of attraction are well defined. The same cannot be said of social systems. The reason stems from the potential for diverse learning rules. Nash equilibrium requires that no agent can do better by deviating. Thus, any learning rule that moves an agent to a higher payoff will not deviate from a Nash Equilibrium. Stability properties as well as the size of basins of attraction can depend on learning rules. Two learning rules applied from the same starting point need not give the same equilibrium. This becomes more likely as the number of equilibria increases.

The existence of an equilibrium, even a symmetric, stable, efficient one, need not imply that it will be generated by an agent-based model. Page and Tassier (2004) investigate an agent-based model of the Groves–Ledyard mechanism which has such an equilibrium. They find that under a wide class of parameter settings the symmetric, efficient equilibrium has a small basin of attraction. Therefore, unless the initial conditions are chosen close to symmetric equilibrium, the agents do not go there. In fact, the agents head off to the boundaries of the message space. If the message space is bounded, the resulting equilibrium is inefficient. If unbounded, the agents head off toward infinity forever. Thus, the unique, stable, efficient interior equilibrium is not likely. Instead, the aggregate behavior of agents results in an inefficient race to the boundary.

4.1 Stability and basins of attraction

Agent-based models are often used to investigate contexts that include features such as networks, heterogeneity, and feedbacks. In such models, multiple equilibrium are commonplace. Page (2007) shows a ‘rule of six’ result for models with agents with finite sets of actions. If the number of types (actions) plus the size of the interacting groups equals six, then multiple equilibria are possible. In other words, a model with two types of agents meeting in groups of size three must attain a unique equilibrium (provided an ergodicity assumption is met). But if three types of agents meet in groups of size three, then multiple equilibria become possible. The rule of six does not guarantee multiple equilibria, but it does demonstrate the mild conditions sufficient for them to exist.

When confronted with a model that has multiple stable equilibria, a social scientist can follow either of two approaches. First, she can apply a refinement criterion to select from among them. She might, for example, assume individuals prefer efficient, symmetric, or risk-dominant equilibria (Young, 2001). The game theory literature contains a plethora of such refinements. Second, she can apply a learning model. This can be accomplished with mathematics or with an agent-based model. Mathematical derivations of basins of attraction can be difficult, which is why so many people use agent-based models for this type of question. In an agent-based model, she would assume a behavioral rule on the part of agents (more on this in a moment) and watch how the model aggregates. This second approach is generative and can help discern the likelihood of different equilibria (Young, 2001)³.

Recall the five-person Voter model from above. That model has four classes of stable equilibria: all 1’s, all 0’s, three consecutive 1’s, and three consecutive 0’s. The model has 32 possible initial configurations. The equilibrium with all 1’s will be reached if the initial configuration begins with all 1’s, with four 1’s, or with three non-consecutive 1’s. The numbers of each of these three types of configurations equal one, five, and five, respectively. Each of these initial configurations can be said to belong to the *basin of attraction* for the all-1’s equilibrium. If we believe that we have a good approximation of behavior, we should expect that equilibria with larger basins of attraction should be more likely to arise. Young (2001) shows that learning models that include noise tend to select *risk dominant*, that is, safer equilibria. This contrasts with the refinement approach, which would argue in favor of efficient equilibria.

At this point, it is worth returning to the question of whether how agents learn matters for stability and basins of attraction. The economics and psychology literatures contain an abundance of learning rules. Vriend (2000) distinguishes between population-based and individual-based learning strategies. In the former, agents learn from the population. For example, they may replicate the actions of more successful agents. In the later, agents alter their actions based only on their own beliefs and experiences. Best-response learning is an example of individual learning. It turns out that for a wide class of learning rules, local stability properties of equilibria do not vary much with choice of learning rules (Hofbauer, 2000). This is because learning rules move in the

³ This approach can be used even if a game has a unique equilibrium. Andreoni and Miller (1995) show how even though two auction designs may have the same equilibrium, one may provide a better learning environment. It is equilibrium may be easier to attain given a particular learning rule.

direction of improving actions and given a perturbation of an equilibrium, going back to that equilibrium will typically be an improvement.

Basins of attraction are another matter altogether. Golman and Page (2009) show that it is possible to construct a three-by-three game in which the basins of attraction under replicator dynamics, a population-based learning rule, have almost no overlap with the basins of attraction for best-response dynamics, an individual-based learning rule. They further show that a necessary condition for this to be the case some non-evolutionarily stable strategy be an initial best response almost everywhere. Less formally, this means that there must be an initial dynamic that moves away from the eventual equilibria of the model. This mathematical result aligns with the intuition that the more complex the dynamics, the more the learning rule should matter in determining basins of attraction. It is also further evidence of how micro-level details matter for aggregate phenomena.

4.2 Path dependence

The analysis so far has assumed deterministic dynamics. In many agent-based models, actions and outcome can be stochastic. If so, basins of attraction need not be well defined in the sense that the same initial point might map to different long-run equilibria or patterns. This potential for *path dependence* has been shown in many agent-based models. The most famous of these models is the Polya urn model (Arthur, 1994). In the Polya urn model, an urn contains one blue ball and one red ball. In each period, a ball is selected. If the ball is red (blue), a second red (blue) ball is added to the urn. The same rule applies in each period. This model produces an infinite number of equilibria. In fact, any proportion of red balls is equally likely to arise. Thus, the equilibrium depends on the path of events. The urn model can be used to explain choice over technologies, locations, or fashions.

The extreme path dependence in the urn model depends partly on *increasing returns*: when a red ball is selected, more red balls are likely to be selected in the future. While that is true, the general notion that path dependence and increasing returns are logically linked is false (Page, 2006). In fact, spatial and resource constraints may be far more likely the cause of path dependence in most agent-based models than increasing returns or positive feedbacks (Hong and Page, 2001).

5 Summary

In this brief survey, I have attempted to highlight three of the key issues pertaining to aggregation in agent-based models. First, it may be useful to model at the level of agent in order to get accurate results, given that aggregate models fail to capture micro-level dynamics. I do not in any way mean to imply that representative-agent models and systems dynamics models have not been useful. To the contrary, they provide the foundation on which agent-based model can be developed. My point is that agent-based models complement these other scientific approaches.

Second, I have provided a description of several types of emergent behavior: macro-level properties that were not explicitly built into the micro-level assumptions of the model. Agent-based models can produce patterns such as crashes and cycles. They can produce spatial clusters of behaviors such as crime. And they can produce emergent distributions of firms sizes and city sizes.

Finally, I have shown how the generative nature of agent-based models provides a method of equilibrium selection in more complex models. These results can depend on the choice of learning rule. Some might see this as a weakness of agent-based models—that their aggregative properties can depend on subtle changes in behavioral rules. Others take a contrary view. If such changes matter in these models, then they might also matter in the real world that these models intentionally approximate. If so, empirical and theoretical understanding of aggregation in real economic systems may require richer modeling of behavior than is the current practice—at least in more complex environments.

In sum, the idea that the whole can differ from its parts is neither new nor novel. Our familiarity with the idea should not blind us to its depth or relevance. In economies, as well as political and social systems, unanticipated aggregate phenomena can produce prosperity or wreak havoc.

Agent-based models provide a methodology for exploring aggregation of systems of diverse, spatialized, interacting agents to help us identify conditions that better ensure the former while reducing the probability of the latter.

Acknowledgement

The author acknowledges the National Science Foundation for an IGERT grant and the Army Research Office Grant No. W911NF1010379.

References

- Anderson, P. 1972. More is different. *Science* **177**, 393–396.
- Andreoni, J. & Miller, J. H. 1995. Auctions with adaptive artificial agents. *Games and Economic Behavior* **10**(1), 39–64.
- Arthur, B. 1994. *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press.
- Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R. G. & Tayler, P. J. 1997. Asset pricing under endogenous expectations in an artificial stock market. In *The Economy As An Evolving Complex System II*, Arthur, W. B., Durlauf, S. & Lane, D. (eds). Addison-Wesley.
- Axelrod, R., Axtell, R., Epstein, J. & Cohen, M. D. 1996. Aligning simulation models: a case study and results. *Computational and Mathematical Organization Theory* **1**, 123–141.
- Axtell, R. 2001. Zipf distribution of U.S. firm sizes. *Science* **293**(5536), 1818–1820.
- Axtell, R. 2003. *A Positive Theory of Emergence for Multi-Agent Systems*. Working paper, Center on Social and Economic Dynamics. The Brookings Institution.
- Axtell, R. & McRae, C. *Model Aggregation for Dynamical Systems*. Personal communication.
- Axtell, R. L. 1999. *The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions*. Working paper. Santa Fe Institute.
- Bak, P. 1996. *How Nature Works: The Science of Self Organized Criticality*. Springer Verlag.
- Bankes, S. 2002. Agent-based modeling: a revolution? *Proceedings of the National Academy of Sciences* **99**, 7199–7200.
- Barabasi, L. 2002. *Linked: How Everything is Connected to Everything Else and What it Means*. Basic Books.
- Bedau, M. A. 1991. Weak emergence. *Nos* **31**(Supplement: Philosophical Perspectives, 11, Mind, Causation, and World), 375–399.
- Bettencourt, L. M. A., Lobo, J., Helbing, D., Kuehnert, C. & West, G. B. 2007. Growth, innovation, scaling, and the pace of life in cities. *Proceedings National Academy of Sciences* **104**, 7301–7306.
- de Marchi, S. 2005. *Computational and Mathematical Modeling in the Social Sciences*. Cambridge University Press.
- Epstein, J. 2003. *Generative Social Science*. Princeton University Press.
- Epstein, J. & Axtell, R. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press.
- Gode, D. K. & Sunder, S. 1993. Allocative efficiency of markets with zero intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy* **101**, 119–137.
- Golman, R. & Page, S. 2009. Basins of attraction and equilibrium selection under different learning rules. *Evolutionary Economics*, **20**, 49–72.
- Glaeser, E., Sacerdote, B. & Scheinkman, J. 1996. Crime and social interactions. *Quarterly Journal of Economics* **111**(2), 507–548.
- Hofbauer, J. 2000. From Nash and Brown to Maynard Smith: equilibria, dynamics and ESS. *Selection* **1**, 81–88.
- Hong, L. & Page, S. 2001. Problem solving by teams of heterogeneous agents. *Journal of Economic Theory* **97**, 123–163.
- Huberman, B. & Glance, N. 1993. Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences of the United States of America* **90**(16), 7716–7718.
- Howitt, P. 2006. Coordination issues in long-run growth. In *Handbook of Computational Economics, II: Agent-Based Computational Economics*, Tesfatsion, L. & Judd, K. (eds). North Holland.
- Iwasa, Y., Andreasen, V. & Levin, S. 1987. Aggregation in model ecosystems. I. Perfect aggregation. *Ecological Modeling* **37**, 287–302.
- Kirman, A. 1992. Whom or what does the representative individual represent? *The Journal of Economic Perspectives* **6**, 117–136.
- LeBaron, B. 2001. Stochastic volatility as a simple generator of apparent financial power laws and long memory. *Quantitative Finance* **1**(86), 621–631, 929–937.
- LeBaron, B. 2006. Agent-based financial markets: matching stylized facts with style. In *Post Walrasian Macroeconomics: Beyond the DSGE Model*, Colander, D. (ed.). Cambridge University Press, 221–235.

- Miller, J. H. 1996. The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior and Organization* **29**(1), 87–112.
- Miller, J. H. & Page, S. E. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press.
- Nirei, M. 2006. Threshold behavior and aggregate fluctuation. *Journal of Economic Theory* **127**, 309–322.
- Page, S. E. 1997. On incentives and updating in agent based models. *Computational Economics* **10**, 67–87.
- Page, S. E. 2006. Essay: path dependence. *Quarterly Journal of Political Science* **1**(1), 87–115.
- Page, S. E. 2007. Type interactions and the rule of six. *Economic Theory* **30**(2), 223–241.
- Page, S. E. 2008. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press.
- Reiter, S. 1977. Information, incentives, and performance in the (new)² economid. *American Economic Review* **67**(1), 226–234.
- Simon, H. A. 1955. On a class of skew distribution functions. *Biometrika* **42**(3–4), 425–440.
- Tesfatsion, L. 1997. How economists can get a-life. In *The Economy as a Complex Evolving System II*. Arthur, W.B., Durlauf, S. & Lane, D. (eds). Addison Wesley, 533–565.
- Vriend, N. 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control* **24**, 1–19.
- Wolfram, S. 2001. *A New Kind of Science*. Wolfram Media.
- Young, P. 2001. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press.