

## Foundations of complexity economics

W. Brian Arthur 

**Abstract** | Conventional, neoclassical economics assumes perfectly rational agents (firms, consumers, investors) who face well-defined problems and arrive at optimal behaviour consistent with — in equilibrium with — the overall outcome caused by this behaviour. This rational, equilibrium system produces an elegant economics, but is restrictive and often unrealistic. Complexity economics relaxes these assumptions. It assumes that agents differ, that they have imperfect information about other agents and must, therefore, try to make sense of the situation they face. Agents explore, react and constantly change their actions and strategies in response to the outcome they mutually create. The resulting outcome may not be in equilibrium and may display patterns and emergent phenomena not visible to equilibrium analysis. The economy becomes something not given and existing but constantly forming from a developing set of actions, strategies and beliefs — something not mechanistic, static, timeless and perfect but organic, always creating itself, alive and full of messy vitality.

For the past 150 years, economic theory has viewed agents in the economy (firms, consumers, investors) as perfectly rational decision makers facing well-defined problems and arriving at optimal behavior consistent with — in equilibrium with — the outcome caused by this behaviour. This view has brought much insight. But many economists<sup>1–7</sup> have pointed out that it is based partly on assumptions chosen for mathematical convenience and, over the years, have raised doubts about whether it is universally applicable. Since the 1990s, economists have instead begun exploring the economy as an evolving complex system, and out of this exploration has come a different approach — complexity economics.

Complexity economics sees the economy — or the parts of it that interest us — as not necessarily in equilibrium, its decision makers (or agents) as not super-rational, the problems they face as not necessarily well-defined and the economy not as a perfectly humming machine but as an ever-changing ecology of beliefs, organizing principles and behaviours. The approach, which has now spread throughout the economics profession, got its start largely at the Santa Fe Institute (SFI) in the late 1980s. But the basic ideas of complexity economics have an even longer history in economics. Even before Adam

Smith, economists noted that aggregate outcomes in the economy, such as patterns of trade, market prices and quantities of goods produced and consumed, form from individual behaviour, and individual behaviour, in turn, reacts to these aggregate outcomes. There is a recursive loop.

It is this recursive loop that makes the economy a complex system. Complexity, the overall subject<sup>8–11</sup>, as I see it is not a science, rather it is a movement within science, and it has roots in thinking developed in the 1970s in Brussels, Ann Arbor and Stuttgart. It studies how elements interacting in a system create overall patterns, and how these patterns, in turn, cause the elements to change or adapt in response. The elements might be cells in a cellular automaton, or cars in traffic, or biological cells in an immune system, and they may react to neighbouring cells' states, or adjacent cars, or concentrations of B and T cells. Whichever the case, complexity asks how individual elements react to the current pattern they mutually create, and what patterns, in turn, result.

The economics I will describe here drops the assumptions of equilibrium and rationality. But it did not come from an attempt to discard standard assumptions, rather it came from a pathway of thinking about how the economy actually works. So instead of giving a formal description,

I will give a personal account of how this economics was arrived at, based on my own experiences. I will also not attempt to survey the hundreds of studies now in the field. Rather, I will discuss how complexity economics came to be, what logic it is based on, what its major themes are and how it links with complexity and physics. I will talk about ideas rather than technicalities, and build from earlier essays of myself and others<sup>12–21</sup> to illustrate the key points, noting that this approach has variants<sup>22,23</sup> and forerunners<sup>24,25</sup>, and it owes much to earlier work by Thorsten Veblen<sup>1</sup>, Herbert Simon<sup>2</sup> and Friedrich Hayek<sup>26</sup>.

### The logic of the approach

**Standard economics and fundamental uncertainty.** Standard economics, called neoclassical economics, studies how outcomes form in the economy from agents' behaviour, and, to do so, it chooses to make several standard assumptions:

- Perfect rationality. It assumes agents each solve a well-defined problem using perfectly rational logic to optimize their behaviour.
- Representative agents. It assumes, typically, that agents are the same as each other — they are 'representative' — and fall into one or a small number (or distribution) of representative types.
- Common knowledge. It assumes all agents have exact knowledge of these agent types, that other agents are perfectly rational and that they too share this common knowledge.
- Equilibrium. It assumes that the aggregate outcome is consistent with agent behaviour — it gives no incentive for agents to change their actions.

These assumptions are by no means perfectly rigid but they constitute an accepted norm. They are made not because theorists necessarily believe they are true, but because they greatly simplify analysis.

The equilibrium assumption in particular is basic to neoclassical theorizing. General equilibrium theory asks what prices and quantities of goods consumed and produced would be consistent with (in equilibrium with) the overall pattern of prices and quantities in the economy's markets — that is, would pose no incentives for those overall patterns to change. Classical game theory asks what strategies or moves of one player

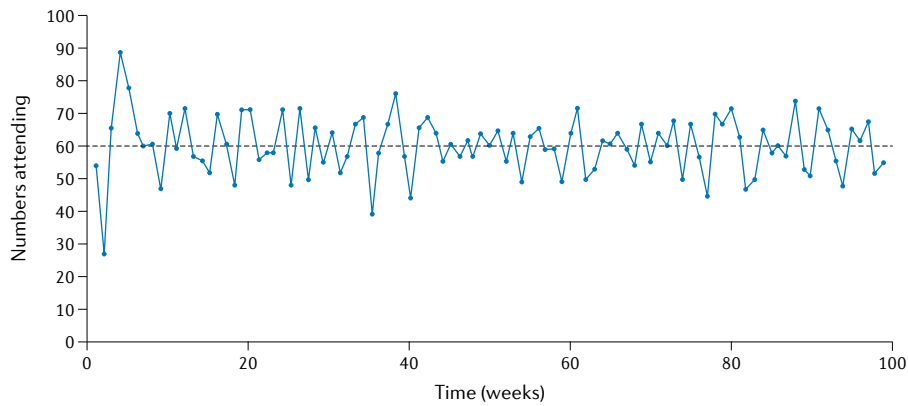


Fig. 1 | **Attendance at the El Farol bar in the first 100 weeks.** Agents attend if they believe the total attendance that week will be no more than 60. Each creates their own set of plausible hypotheses or predictors of attendance, and, every week, acts on their currently most accurate one. Figure reprinted with permission from REF.<sup>12</sup>, AAAS.

would be consistent with the strategies or moves their rivals might choose — that is, would be the best course of action for that player. Rational expectations economics asks what forecasting methods would be consistent with the outcomes these forecasting methods brought about — that is, would statistically, on average, be validated by outcomes.

Overall, this equilibrium approach has worked quite well. It is a natural way to examine questions in the economy and open these up to mathematical analysis, and it illuminates a wide range of issues in economics. I admire its elegance; it has yielded, in Paul Samuelson’s words<sup>27</sup>, an “austere aesthetic grace.” But it severely limits what can be seen. By its definition, equilibrium makes no allowance for the creation of new products or new arrangements, for the formation of new institutions, for exploring new strategies, for events triggering novel events, indeed, for history itself. All these have had to be discarded from the theory. “The steady advance of equilibrium theory throughout the twentieth century,” says David Simpson, “remorselessly obliterated all ideas that did not fit conveniently into its set of assumptions.”<sup>28</sup> Over the past 120 years, economists such as Thorstein Veblen<sup>1</sup>, Joseph Schumpeter<sup>7</sup>, Friedrich Hayek<sup>29</sup>, Joan Robinson<sup>5,30</sup> and others<sup>4,31–35</sup> have objected to the equilibrium framework, each for their own reasons.

All have thought a different economics was needed.

It was with this background in 1987 that the then-new SFI convened a conference to bring together ten economic theorists and ten physical theorists to explore the economy as an evolving complex system. The meeting was a success and, a year later, these initial explorations became SFI’s

first research programme<sup>8,36–38</sup>. I was asked to lead this programme, and, after many discussions, we realized that we kept coming back to the same question: what would economics look like if we went beyond the standard assumptions?

For one thing, agents differ<sup>39</sup>. Companies in a novel market may have different technologies, different motivations and different resources, and they may not know who their competitors will be or, indeed, how they will think. They are subject to what economists call fundamental uncertainty<sup>40</sup>. As John Maynard Keynes described this in 1937, “the prospect of a European war... the rate of interest twenty years hence... About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.”<sup>41</sup> As a result, the decision problem faced by agents is not logically defined and, so, it cannot have a logical solution. It follows that rational behaviour is not well-defined. Therefore, there is no ‘optimal’ set of moves, no optimal behaviour. Faced with this — with fundamental uncertainty, ill-defined problems and undefined rationality — standard economics understandably comes to a halt. It is not obvious how to get further.

**The El Farol problem.** And yet people do act in ill-defined situations, and they do so routinely. As a concrete example, consider the El Farol bar problem<sup>42</sup>. One hundred agents attempt once a week on Thursday nights to forecast attendance at their favourite bar, El Farol in Santa Fe. If they believe the bar will be too crowded — will have more than 60 people, say — they will not go; if they believe fewer than 60 will show up, they go. How will they act?

Deductive logic does not help. Agents’ predictions of how many will attend depend

on their ideas of what others’ predictions will be, which depend, in turn, on their ideas of others’ predictions, and there is an infinite regress. Further, if a shared rational forecasting model did exist, it would quickly negate itself: if it predicted few will attend, all would go; if it predicted many will attend, nobody would go. Agents, therefore, face fundamental uncertainty: they do not know how other agents will decide on their forecasts, and, yet, such knowledge determines attendance. The problem is ill-defined.

One can model this situation by assuming agents act inductively: each creates their own set of plausible hypotheses or predictors, and, every week, acts on their currently most accurate predictor. In other words, a framework for studying the economy should involve agents that form individual beliefs or hypotheses — internal models (possibly several simultaneously) — about how to respond to the situation they are in.

Such agents could be implemented as small, individual computer programs that could differ, explore and learn to get smart. How they could do this — how they could get smart — was inspired by the work of computer scientist John Holland, who had spent much of his career developing methods by which computer algorithms could learn to play checkers/draughts or chess. Holland’s algorithms could ‘recognize’ the current state of the game and learn to associate appropriate moves with it. The moves would be fairly random to start with and not very useful, but, over many games, the program would learn which moves worked in which situations, ‘explore’ new moves and drop ones that did not work — it would get smarter. In economic problems, agents could start with their own arbitrarily chosen or random beliefs, learn which ones worked and explore new ones occasionally, from time to time dropping ones that did not perform well and replacing them with new ones to try out<sup>42–44</sup>. They could, in this way, operate and explore in an ill-defined setting and become more intelligent as they gained experience.

Notice two things about this framework. First, it is dynamic and open to new behaviours, often unthought of ones. The system may converge to an equilibrium in many cases, in others, it may not — it may perpetually discover novel behaviours. So, in general, we have a nonequilibrium economics. Second, the very explorations agents undertake alter their situation, which requires them to explore and adapt afresh, which changes the situation. We are in a world of complexity.

In the case of El Farol, computational experiments show (FIG. 1) that attendance in

the bar (and the collection of forecasts being acted on) self-organizes into an equilibrium pattern that hovers around the comfortable 60 level. The reason is that, if fewer than 60 came in the long term, low forecasts would be valid, so many would come, negating those forecasts; and if more came in the long run, fewer would show up. So an attraction to this level emerges. But, although the population of forecasts on average supports this comfortable level, the actual forecasts in use keep changing. The outcome is a bit like a forest, the shape of which does not change, but the individual trees of which do. Notice that equilibrium in this problem is not assumed, it emerges — self-organizes — because it is a natural attractor.

#### **Agents responding to ill-defined situations.**

The El Farol problem was an early study using our Santa Fe approach, and others followed<sup>45</sup>. Inevitably, we were asked to name this approach, and, in a 1999 Science paper<sup>12</sup>, I labelled it ‘complexity economics’. At the heart of our approach were agents responding to ill-defined situations by ‘making sense’ or recognizing some aspects of them, and choosing their actions, strategies or forecasts accordingly. Ways of modelling this have now widened significantly. Behavioural economics<sup>46</sup> gives insights into how real human agents respond in the context we are looking at. Artificial intelligence or neural nets<sup>47</sup> can be used to model how agents respond to the signals they are getting. Evolutionary programming can create novel unforeseen strategies (as in AlphaGo Zero). Modern psychology shows us how agents use narratives, imagination and calculations to make sense in ill-defined circumstances<sup>48,49</sup>.

Some models in complexity economics use mathematics (such as nonlinear stochastic processes), but, often, the sheer complication of keeping track of the decision processes of multiple agents requires the use of computers. We then build models around agents’ individual behaviour, and, so, agent-based modelling arises naturally<sup>50</sup>. Agent-based models<sup>51–55</sup> are now used all across economics. Some have a few hundred agents; a recent one has 120 million<sup>56</sup>. Some take account of legal and regulatory institutions. Some are designed to simulate reality — the 2008 subprime mortgage meltdown or the economics of the 2020 COVID-19 pandemic. Some investigate theoretical issues — financial asset pricing. But whatever the design of these studies, the idea, as in all of economics, is to explore how outcomes follow from assumed behaviour.

#### **An ecology of behaviours**

In the El Farol problem, agents’ forecasting methods vie to be valid in a situation that is dependent on other agents’ forecasts — they compete in an ‘ecology’ of forecasts. Indeed, a general feature in complexity economics is that agents’ beliefs, strategies or actions are tested for survival within a situation or ecology that these beliefs, strategies or actions together create. They act in a way like species, continually competing or mutually adapting and co-evolving. As a result, a distinct biological evolutionary theme emerges.

Here is an example. In a classic study<sup>57</sup>, a computerized tournament was constructed in which strategies compete in randomly chosen pairs to play a repeated prisoner’s dilemma game. (It is not necessary to understand the details of the prisoner’s dilemma; simply think of the experiment as a repeated game played one-against-one by a current collection of strategies.) Each strategy is a set of fixed instructions for how to act given its and its opponent strategy’s immediate past actions. If strategies perform well over many encounters, they replicate. If they do badly, they die and are removed. Every so often, existing strategies can mutate their instructions, and, occasionally, can deepen by having a lengthier memory of immediate past moves. At the start of the tournament, simple strategies such as tit-for-tat dominate, but, over time, more sophisticated ones show up that exploit them. In time, still more sophisticated strategies emerge to take advantage of these and the simpler ones drop out, and periods of relative stasis alternate with ones of dynamic upheaval (FIG. 2). One can think of each strategy type as a species, well-defined and differing from other species, occasionally mutating to produce a new species. Evolution enters in a natural way that arises from strategies mutually competing for survival and mutating as they go.




Fig. 2 | **Prevalence of strategies in a simulated tournament of the prisoner’s dilemma.** Over time, strategies can evolve based on pressures exerted by other strategies. The lengths of labels indicate the memory depth of strategies, that is, how many previous moves in the game they take into account. Figure reprinted with permission from REF.<sup>139</sup>, Elsevier.

Outcomes for the computerized tournament differ randomly each time it is run. In some runs, an evolutionarily stable strategy appears (one that cannot be invaded by some novel strategy). In other runs, the outcome keeps evolving indefinitely. In some runs, complicated strategies appear early on, in others, they appear only later. But, in spite of these variations, the experiment shows consistent phenomena: the exploitation of strategies by other strategies, emergence of mutual support among strategies, sudden collapses of strategies and takeover by novel ones, periods of stasis followed by ones of turbulent change. The overall scene looks like species competition in palaeozoological times.

Such outcomes are common with complexity in the economy. What constitutes a ‘solution’ — the outcome of the model — is frequently an ecology in which strategies, or actions, or forecasts compete; an ecology that might never settle down, and that shows properties that can be studied qualitatively and statistically.

This vision fits well with Alfred Marshall’s famous dictum in 1890 that “the Mecca of the economist lies in economic biology.”<sup>58</sup>

#### **Simple models, complex phenomena**

A new theoretical framework in a science does not really prove itself unless it explains phenomena that the accepted framework cannot. Can complexity economics make this claim? I believe it can.

Consider the Santa Fe artificial stock market model<sup>59,60</sup>.

The standard, neoclassical theory of financial markets<sup>61</sup> assumes rational expectations: identical investors adopt identical forecasting models that are, on average, statistically validated by the prices they forecast. The theory works convincingly to explain how market prices come about and how they reflect the stream of random

earnings. But it has some key shortfalls: for one, in this theoretical market, no trade at all takes place. The reason is simple. Investors are identical, so if one of them wants to buy, all want to buy and there are no sellers; if one wants to sell, they all want to sell and there are no buyers; the stock price simply adjusts to reflect these realities. Further, the theory cannot account for actual market phenomena such as the emergence of a market psychology, price bubbles and crashes, the heavy use of technical trading (trades based on the recent history of price patterns)<sup>62</sup> and random periods of high and low volatility (price variation).

At SFI, we created a different version of the standard model. We set up an ‘artificial’ stock market inside the computer and our ‘investors’ were small, intelligent programs that could differ from one another. Rather than share a self-fulfilling forecasting method, they were required to somehow learn or discover forecasts that work. We allowed our investors to randomly generate their own individual forecasting methods, try out promising ones, discard methods that did not work and periodically generate new methods to replace them. They made bids or offers for a stock based on their currently most accurate methods and the stock price forms from these — ultimately, from our investors’ collective forecasts. We included an adjustable rate-of-exploration parameter to govern how often our artificial investors could explore new methods.

When we ran this computer experiment, we found two regimes, or phases<sup>60</sup>. At low rates of investors trying out new forecasts, the market behaviour collapsed into the standard neoclassical equilibrium (in which forecasts converge to ones that yield price changes that, on average, validate those forecasts). Investors became alike and trading faded away. In this case, the neoclassical outcome holds, with a cloud of random variation around it. But if our investors try out new forecasting methods at a faster and more realistic rate, the system goes through a phase transition. The market develops a rich psychology of different beliefs that change and do not converge over time; a healthy volume of trade emerges; small price bubbles and temporary crashes appear; technical trading emerges; and random periods of volatile trading and quiescence emerge.

Phenomena we see in real markets emerge.

This last phenomenon of random periods of high and low volatility happens because, if some investors occasionally discover new profitable forecasting methods, they then invest more and this changes the market slightly, causing other investors to also

change their forecasting methods and their bids and offers. Changes in forecasting beliefs thus ripple through the market in avalanches of all sizes, causing periods of high and low volatility.

I want to emphasize something here: such phenomena as random volatility, technical trading or bubbles and crashes are not ‘departures from rationality’. Outside of equilibrium, ‘rational’ behaviour is not well-defined. These phenomena are the result of economic agents discovering behaviour that works temporarily in situations caused by other agents discovering behaviour that works temporarily. This is neither rational nor irrational, it merely emerges.

Other studies<sup>63–66</sup> find similar regime transitions from equilibrium to complex behaviour in nonequilibrium models. It could be objected that the emergent phenomena we find are small in size: price outcomes in our artificial market diverge from the standard equilibrium outcomes by only 2% or 3%. But — and this is important — the interesting things in real markets happen not with equilibrium behaviour but with departures from equilibrium. In real markets, after all, that is where the money is made.

This remark above does not mean that complexity economics always makes small differences. It studies how solutions or structures form, and, often within these, qualitatively new phenomena or major differences emerge.

### A word on agent-based computation

The examples I’ve described contain enough complication with their differing agents’ behaviours that we need to use computation. This is normal. In fact, a closely related approach highlights computation and goes by the label agent-based computational economics<sup>67–70</sup> (Axtell, R. & Farmer, D., manuscript in preparation). It overlaps with the approach I am describing and is the subject of much current interest, so it is worth looking at the relation between the two. I would say this. In the 1980s, computation became available in simple but practical form, and it was computation more than anything else that allowed economic theorists to venture beyond the standard neoclassical assumptions — for instance, to allow complicated inductive reasoning and compute its consequences. If we turn these new possibilities into a theoretical framework, we get complexity economics, or something like it. If we turn them into a solution method, we get agent-based computational economics. So there is no well-marked boundary

between the two approaches. One could, therefore, regard agent-based computational economics as a key method within the framework of complexity economics; or one could regard complexity economics as a conceptual foundation behind agent-based economic modelling. I should note that there are differences: complexity economics uses both mathematics and agent-based computation, and investigates patterns that endogenously form and change in the economy<sup>71</sup>. And agent-based models often concern themselves with computational technicalities, and see themselves as stand-alone and not subject to any particular theoretical foundation. But granted these different emphases, the two approaches blend together. Depending on whether a study emphasizes theory or method, it can fly either flag — or both.

However they are labelled, computational studies are valuable: they offer agent-based behavioural realism and they allow realistic detail; standard economics typically relates average aggregate quantities (outputs produced, say) to average aggregate quantities (inputs used) and, often, the details within such aggregates matter. But, in spite of their advantages, in my experience, computation-based models are still regarded with suspicion in mainstream economic journals — they are held to be ad hoc, open to using arbitrary assumptions or ones chosen for preordained purposes. I agree there is plenty of scope for nefarious modelling, but, as has been pointed out, this is true in equation-based modelling as well<sup>72</sup>. Rigour in a computational setting needs to widen from insistence on correctness of the logic (which, of course, remains imperative) to insistence on strict scientific honesty. It demands careful, verifiable modelling with realistic behaviour and reproducible, analysable results.

A different objection is that equation-based theory uses mathematics with all its majesty and power, and computation-based theory uses, well, computers. But the difference is superficial. Both methods trace a pathway from agent behaviour to its implied outcome. Equation-based models allow one to follow the logical steps of this pathway — how the outcome is implied by the model — and computational models cannot do this. But they compensate in another direction. They are themselves largely collections of equations, and they have the capacity to be expanded to encompass an arbitrary amount of realistic detail. Furthermore, they allow if-then conditions. This means they can allow the changing context of the situation — the ‘if’ clause of where the computation currently is — to direct agents’ behaviour in any way

appropriately called for<sup>73</sup>. This possibility is powerful and, once again, it connects with complexity: agents' behaviour changes the context and the context changes behaviour. On both these counts, computation widens theory's scope.

### Events propagating in networks

**Economic networks.** Very often in a complex system, the actions taken by individual elements are channelled via a network<sup>74</sup> of connections among them. Within the economy, networks arise in many ways, such as trading, information transmission, social influence or lending and borrowing. Several aspects of networks are interesting: how their structure of interaction or topology makes a difference; how markets self-organize within them<sup>75</sup>; how risk is transmitted; how events propagate; how they influence power structures<sup>76</sup>. It is not possible here to cover all themes of interest in network economics<sup>77</sup>. I will simply point out three features.

**Propagation of change.** The topology of a network matters as to whether connectedness enhances its stability or not<sup>78</sup>. Its density of connections matters, too. When a transmissible event happens somewhere in a sparsely connected network, the change will fairly soon die out for lack of onward transmission; if it happens in a densely connected network, the event will spread and continue to spread for long periods. So, if a network were to slowly increase in its degree of connection, the system will go from few, if any, consequences to many<sup>79</sup>, even to consequences that do not die out. It will undergo a phase change. This property is a familiar hallmark of complexity. Notice that the propagation of events brings time inexorably into systems; without such propagation, time disappears.

**Power laws.** Research on networks shows that cascades of events causing further events<sup>79</sup> often follow power laws (the frequency  $p$  of propagation lengths  $x$  follows  $p(x) \sim x^{-a}$  ( $a > 0$ )). And fluctuations related to cascading events often have long-tailed probability distributions (roughly, large deviations have higher probability than they would under Gaussian distributions). Such features occur in all systems — physical, biological, geological — in which events propagate<sup>11</sup>, and they have been familiar in economics at least since the work of Vilfredo Pareto in the early 1900s. But, despite this, standard economics has traditionally assumed that firms, investors and economic events are unconnected and independent, therefore, the changes they cause deviate

from some systemic average in a normal or Gaussian way. Accordingly, finance theory assumed normal fluctuations (as did the famous 1973 Black–Scholes formula for pricing options). This is now changing. Modern network theory shows that power laws and long tails are to be expected in the economy, and empirical studies of price fluctuations bear this out<sup>62,65,80</sup>. Such findings matter in finance. Contemporary financial derivatives markets trade trillions of dollars daily, and traders are forced to take account of such realities<sup>81</sup>.

**Systemic risk.** Networked events have consequences for overall risk in an industry. If firms are unconnected and independent, their ups and downs offset each other, so the possibility that a negative event at the level of one firm could trigger collapse of the industry or economy — called its systemic risk — is relatively low. But when companies are connected in networks of financial dependence, this changes<sup>82</sup>. Banks borrow from or lend to other 'counterparty' banks in their immediate network. If an individual bank discovers it holds distressed assets — counterparty loans that will not be repaid — it comes under pressure to increase its liquidity and call in its loans from its counterparty banks. These, in turn, come under pressure to call in their counterparty loans, and distress can cascade across the network<sup>83</sup>. The overall system can then become threatened or collapse, which is what

happened in 2008. It has been proposed<sup>84</sup> that loans by banks to other individual banks be taxed according to the change in systemic risk they cause, which forces the system to self-organize in a way that minimizes risk.

### Policy

Does complexity economics lead to different policies than the ones neoclassical economics advocates? I believe so. In equilibrium economics, policy typically means adjusting some means of incentive — taxes, regulations, quotas — to gain a desired outcome. Certainly, this approach can be effective, though in cases in which policy is guided by theory based on assumptions adopted for analytical convenience or ideology, it may be dubious. With complexity models<sup>85</sup>, one can bring in much-needed realism<sup>86,87</sup>: agents may differ, in region or class or response; their attitudes can change endogenously<sup>88</sup>; the details of institutions can be built in; and fundamental uncertainty and unseen disturbances can be allowed for. The implications of policy can be explored in ways that go beyond narrow economic 'efficiency'. One can set up policy labs — carefully controlled computer experiments — to test policy designs and game out their consequences. All these are refinements of policy.

But one can go further. Dropping the equilibrium assumption reveals an economy that is open to new behaviour, even to being exploited or gamed by small groups of

#### Box 1 | All systems will be gamed

Standard economics has learned how to stabilize macroeconomic outcomes, avoid depressions, regulate currency systems, manage central banking and carry out antitrust policy. What it has not been able to do is prevent financial and economic crises.

Financial crises happen when small events trigger a cascade of further events that get out of control, or when a small group of players gains control of some part of the system<sup>140,141</sup> to its own private advantage but to the detriment of the system as a whole. Thus, in Russia's 1990s transition from communism to capitalism, a coterie of private players took control of the state's newly freed assets for their own benefit and industrial production plummeted<sup>142,143</sup>. In California's 2000 freeing of its energy market, a small number of suppliers manipulated the market to their own profit and the state's finances suffered<sup>144</sup>. In the USA's mortgage-backed securities market in 2008, financial institutions on Wall Street had obtained looser regulations and created exotic derivative products they greatly benefited from, which caused an unstable structure that spectacularly collapsed<sup>143,145</sup>. Each of these systems was manipulated or 'gamed', and all broke down.

The consequences of economic collapse are serious. So why does equilibrium economics not warn us of these potential failures in advance? The reason is subtle: equilibrium economics is not primed to look for such possibilities. If we assume a system is in stasis or equilibrium, then, by definition, cascades of hazardous events and their consequences cannot happen, and, also by definition, players cannot find ways to manipulate the situation and improve their position. And so, a muted bias precludes the idea of collapse. Complexity, by contrast, sees the economy as a web of incentives open to further actions or to exploitation, so it disposes us to examine economic systems for where they might be open to manipulation or to systemic failure.

Can we program computers to probe for weaknesses? I believe we can. We can model large policy systems and probe them, deliberately or automatically, to see where they might be exploited. We need to adopt such failure-mode practices from structural engineering, or aircraft design<sup>146</sup>, or encryption, and examine where economic systems have weak points and might be manipulated. Doing so would yield more reliable economic and social outcomes.

## Box 2 | The economy creating itself from itself

Where does an economy come from? How does it form itself and change structurally? We are asking how the economy changes in character from canals to railroads or from electronics to algorithms. Economists have long been aware that economies form largely around their means of production (their technologies): industrial processes, machinery, business procedures, transportation methods. And they change structurally as these change. Equilibrium economics acknowledges this but handles it simplistically. Certain technologies exist; certain new ones are somehow invented; production changes, prices change, the equilibrium shifts.

Complexity economics offers a richer story (I will condense it here)<sup>116</sup>. It starts with the observation that technologies are means to human purposes and are constructed, put together — combined — from parts, assemblies, sub-assemblies<sup>116,147,148</sup>. These latter are also means to purposes; thus, new technologies form by combination from existing technologies (albeit with much human ingenuity)<sup>116</sup>. But things do not stop there. A new technology sets in motion a sequence of events — an ‘algorithm’, if you like. When a new technology appears, it replaces existing technologies; calls forth new ones to satisfy its needs; becomes a component available for the creation of further new technologies; and causes the economy, society and their institutions to rearrange themselves<sup>149</sup>. Thus, when the railway locomotive entered the economy, it replaced existing horse-drawn trains; set up needs for the fabrication of iron rails and the organization of railways; caused the canal and horse-drayage industries to wither; became a key component in the transportation of goods; and, in time, caused factories to relocate and towns to grow. The economy transformed itself structurally.

Once set in motion, this sequence of events need not stop. By calling forth new technologies and becoming a component for further technologies, a technology may cause further technologies to be added. These, in turn, bring forth the same sequence and, with this, a cascade of further events. The algorithm may be simple but it ‘calls itself’ within itself, and, in doing so, brings forth rich-patterned change. It does this at all levels and concurrently, causing continuous, unending disruption. The economy, in turn, forms from its technologies, which call for and contribute to the creation of further technologies and, thereby, the economy’s further formation. The economy, thus, continually creates itself from itself.

players (BOX 1), and one can formulate ways to prevent this. And dropping the coarseness of models that implicitly assume average agents makes it possible to look at distributional issues, that is, at different agents being affected differently by policies (discussed below). Complexity widens the policy arena.

**Some frontiers**

It is now more than 30 years since our discussions of complexity economics started at SFI, and many of its ideas are being absorbed into the core of economics. But the new approach is not yet fully central. I believe this is to be expected. For any field to change at a fundamental level, its textbooks, teaching, journal editors and highly trained practitioners must themselves change. Indeed, game theory and behavioural economics each took 40–50 years to be absorbed into the core of economics (Axtell, R. & Farmer, D., manuscript in preparation).

By that measure, complexity economics is still arriving. There are now general texts on the subject<sup>89–96</sup> and research across sub-fields such as macroeconomics<sup>97–99</sup>, labour economics<sup>100,101</sup>, institutional economics<sup>102</sup>, environmental economics<sup>103–106</sup>, finance<sup>107–109</sup>, economics of disease transmission<sup>110,111</sup>, distribution of firms’ sizes<sup>112</sup>, scaling laws<sup>113,114</sup>, ergodicity in economics<sup>115</sup>, technological innovation<sup>116–118</sup> and economic

development<sup>119</sup>. If there are trends, they are towards more behavioural realism, grounding models on large data sets<sup>120</sup>, using computer experiments to study and design systems, and understanding how macro patterns emerge from micro assumptions.

Here are some frontiers I find interesting.

**Formation in the economy.** Neoclassical economics examines equilibrium patterns in the economy: patterns of production, consumption, prices and of quantitative growth in these entities. It cannot readily look at questions of formation — how the arrangements and institutions of the economy come to be in the first place and how the economy changes in character over time. Complexity economics, by contrast, sees the economy as open and subject to novelty, and it can explore formation naturally (BOX 2). It also assumes there may be positive feedbacks (or increasing returns) in the economy; these act to amplify small differences in history and can lead to the lock-in of giant firms, especially in the technology sector (BOX 3). And because complexity economics looks at how structures form or solutions come to be ‘selected’, it connects robustly with the dynamics of evolutionary economics.

Complexity also links with pre-neoclassical approaches in economics — political economy, classical economics and

Austrian economics. These are venerable traditions with different emphases, but, together, they see historical contingency as important, economic structures as perpetually in formation and the economy as rich in process. Because they emphasize process and qualitative formation, they were not easily mathematized, and, in the twentieth century, became sidelined by equilibrium theory. Complexity is connecting with these earlier approaches and giving them new voice<sup>14,28,31,121,122</sup>.

**Econophysics.** Since the 1990s, physicists have been applying physics models and methods to economics, especially within finance<sup>123</sup>. This new field is growing rapidly, and, although it does not quite overlap with complexity economics, it is worth mentioning here because it is physics-based. Studies vary, but the tendency, as in other branches of physics and engineering, is to explore large real data sets and seek simple mechanisms within these. Sometimes, this has had marked success<sup>124,125</sup>.

**Distributional issues.** Neoclassical economics concerns itself greatly with growth and efficiency — the what’s-produced of the economy — and much less with distributional issues — the who-gets-what of the economy. One reason for this is that, for analytical convenience, standard economics often models issues at a coarse-grained level, say at the country level, so that individual regions or groupings of people become unseen or averaged away<sup>126</sup> — the models are mean-field. Then, how these unseen individual agents or groupings will fare under a new policy is unspecified and it’s easy to assume by default that they will benefit equally. In models that allow explicitly diverse agents, as with complexity economics, this ceases to be the case: some may benefit, some may lose. In the early 1990s, standard economic doctrine taught that free trade and globalization were, in most circumstances, beneficial<sup>127</sup>. Offshoring from the USA to locations such as Mexico or China would, therefore, be advantageous: Mexico and China would get new industry and jobs and the USA would get cheaper goods. Such arrangements would, indeed, have been socially optimal if all parts of a given country or territory were the same; they would all benefit equally. But, in practice, regional differences, especially in the USA, mattered. Many economists now believe that offshoring of the US economy to China and Mexico was a major factor in hollowing out jobs in regions such as the US Rust Belt<sup>128</sup>, which has brought grim

consequences to social wellbeing<sup>129</sup> and US politics ever since. Models with agents with realistic, regionally diverse circumstances would have foreseen this outcome, and they open a new capacity to explore distributional issues.

**More realistic modelling.** As discussed above, complexity economics and agent-based computation allow for more realistic modelling across economics and related fields. For example, standard, mean-field, infectious-disease-transmission models assume that the average infected person, on average, infects  $R_0$  further people. With agent-based modelling, one can break out the transmission process, assume diverse agents with diverse circumstances and follow the event-by-event transmission process realistically. More precise detail allows sharper resolution and one sees features that would not be visible otherwise<sup>110</sup>.

**Industry applications.** Industry applications are still at a beginning. Complexity thinking and agent-based computational experiments help where sequences of events and responses to them matter, as occurs in transportation logistics<sup>130</sup> or in citywide traffic management<sup>131</sup>. It also helps where fundamental uncertainty exists, as in planning future operations in the face of unforeseen financial crises, possible wars, epidemics, power outages, abrupt changes in regulation or unexpected actions by competitors. In such cases, optimization may not be appropriate — indeed, it may not be well-defined. A better approach would allow for a multiplicity of candidate responses by computerized ‘agents’ and use complexity methods such as genetic algorithms or evolutionary programming to ‘learn’ and select appropriate responses to given circumstances. In this way, ‘intelligent’ behaviour self-organizes, as with the complexity models I described earlier. What is important in industry is not just efficiency but robustness and resilience — the ability to react to unforeseen circumstances and to recover or transform quickly if something goes wrong. This way of thinking brings a different approach not just to business operations but to management itself. It calls for adaptive, resilient and organic thinking, rather than deterministic, top-down, mechanistic control<sup>132</sup>.

**The autonomous economy.** In the 1960s, the character of the economy in the USA and Europe was heavily determined by large industrial organizations that produced goods and services. In the 1990s, this

changed, and production was sizably offshored. Now, under rapid digitization, the economy’s character is changing again and parts of it are becoming autonomous<sup>133</sup> or self-governing. Financial trading systems, logistical systems and online services are already largely autonomous; they may have overall human supervision, but their moment-to-moment actions are automatic, with no central controller. Similarly, the electricity grid is becoming autonomous (loading in one region can automatically self-adjust in response to loading in neighbouring ones<sup>134</sup>); air-traffic control systems are becoming autonomous and independent of human control<sup>135</sup>; and future driverless-traffic systems, in which driverless-traffic flows respond to other driverless-traffic flows, will likely be autonomous. Such systems have much in common with the operational systems I just described. Besides being autonomous, they are self-organizing, self-configuring, self-healing and self-correcting, so they show a form of artificial intelligence. One can think of these autonomous systems as miniature economies, highly interconnected and highly interactive, in which the agents are software elements ‘in conversation with’ and constantly reacting to the actions of other software elements. A blockchain system (a secure, decentralized, highly autonomous digital ledger) is conversationally interactive

in this way. Indeed, as the economy digitizes, it is increasingly made up of autonomous conversing systems. It becomes ever more an evolving, complex system.

### An overall perspective

In the end, what is my view on this new approach to economics? Here is a brief summary.

Complexity economics relaxes the assumptions of neoclassical economics — the assumptions of representative, hyper-rational agents, each of which faces a well-defined problem and arrives at optimal behaviour given this problem (TABLE 1) — and, thus, gives a different style to economics. It is an economics in which the agents in the economy are realistically human and realistically diverse, in which path-dependence and history matter, in which events trigger events<sup>136</sup> and in which the networks that channel these events matter. It is an economics in which equilibrium is not assumed, if it is present, it emerges; in which rational behaviour is not assumed, in general, it is not well-defined; in which the unexpected crises of the economy can be probed and planned for in advance; in which free markets are not assumed to be optimal for society but can be assessed realistically; and in which distributional issues are not covered up, but can be rigorously scrutinized.

#### Box 3 | Silicon Valley economics

One early theme in complexity economics has been the effects of positive feedbacks. Traditionally, standard economics has assumed negative feedbacks (or diminishing returns). There are only so many good hydroelectric sites in Sweden and, once these are used up, hydro energy runs into diminishing returns — it becomes more costly. Thus, hydro-based and petroleum-based energy share the market in an efficient and predictable way.

But some economic markets — particularly tech ones — show positive feedbacks (increasing returns). If one company or technology gets ahead, it accrues network effects (if more people I deal with are on PayPal’s payment system, it increases my advantage to adopt PayPal), or it can lower costs by spreading its upfront R&D expenses over a wider user base; it, therefore, reaps further advantage. When several such companies compete, one that gets ahead by good fortune or clever strategy may come to dominate or lock in the market. But the winner need not be the best.

Economists have long known about increasing returns. Alfred Marshall in 1890 speculated that, if  $N$  firms competed and each had increasing returns, the market would go to “whatever firm first gets a good start.”<sup>158</sup> But in static equilibrium economics, this causes a problem: if multiple equilibria are possible, we cannot say which one might occur. The outcome is indeterminate.

Complexity economics resolves this indeterminacy<sup>150</sup> by allowing such situations to play out over time. ‘Small random events’ occur — what product launches when, who sat next to whom on an airplane, what design caught the early imagination — and, over time, increasing returns magnifies the cumulation of such events to ‘select’ the outcome randomly. Thus, increasing returns problems in economics are best seen as dynamic processes with random events and natural positive feedbacks — as nonlinear stochastic processes<sup>151</sup>. They may yield different outcomes in different realizations.

Such properties of multiple equilibria, non-predictability, lock-in, inefficiency, historical path dependence and asymmetry in economics are similar to phenomena in physics: multiple metastable states, unpredictability, phase-locking or mode-locking, high-energy ground states, non-ergodicity and symmetry breaking.

Increasing returns have become the basis for our understanding not just of tech markets<sup>152</sup> but of economic geography<sup>153</sup>, international trade<sup>154</sup>, patterns of inequality<sup>155,156</sup> and segregation<sup>157</sup>.

Table 1 | Differences between neoclassical and complexity economics

Feature <sup>a</sup>	Neoclassical economics	Complexity economics
Agents	Representative, with 1, 2, N or a distribution of types	Diverse
Organizing principle	Equilibrium. Agent behaviour consistent with aggregate outcome	Nonequilibrium. Agent behaviour reacts to aggregate outcome
Metaphor	Well-functioning machine	Ecology: of forecasts, actions, strategies
What is faced by agents	Well-defined problem	Ill-defined situation
Behaviour	Agents optimize	Agents face fundamental uncertainty, they try to make sense, explore
Structural change	The equilibrium shifts	Novelty causes endogenous restructuring
Rationality	Perfect and boundless	Rationality usually not defined
Feedbacks	Diminishing returns	Increasing, as well as diminishing, returns
Time	Equilibrium is timeless	History and path taken matter
Dominant theme	Allocation of resources	Formation of structures
System	Closed to new behaviour	Open. System can be exploited
Methods used	Mathematics (quantities, incentives in balance)	Mathematics and computation (algorithmic and event-driven)
Temporary phenomena	Excluded by equilibrium	Possibly emerge
Interaction	Homogeneous	Channelled by networks
Evolution of economy	Outcomes usually seen as in stasis. Not evolving	Economy self-creating, in perpetual novelty

<sup>a</sup>Entries are general; there may be exceptions to them in particular studies.

Because its assumptions are a widening of the neoclassical ones, complexity economics is neither a special case of equilibrium economics nor an addition to it<sup>137</sup>. On the contrary, it is economics done in a more general way. This broadening of principles is not due to a shift in ideology. It is due, I believe, to new tools becoming available to economics: methods to think about decision making under fundamental uncertainty and to deal with nonlinear dynamics and nonlinear stochastic processes. Above all, it is due to computation<sup>138</sup>, which makes it possible to model arbitrarily more complicated and more realistic behaviour.

It would be naive to say that this widening of scope will be a panacea for economics, but it certainly releases economics from the strictures of its neoclassical assumptions. I see this shift in economics as part of a larger shift in science itself. All the sciences are shedding their certainties, embracing openness and process, and asking how structures or phenomena come into being. There is no reason that economics should differ in this regard. Complexity economics sees the economy not as mechanistic, static, timeless and perfect but as organic, always creating itself, alive and full of messy vitality.

W. Brian Arthur <sup>1,2,3</sup>

<sup>1</sup>Santa Fe Institute, Santa Fe, NM, USA.

<sup>2</sup>Intelligent Systems Lab, PARC, Palo Alto, CA, USA.

<sup>3</sup>Stellenbosch Institute for Advanced Study, Stellenbosch, South Africa.

e-mail: [wbarthur@parc.com](mailto:wbarthur@parc.com)

<https://doi.org/10.1038/s42254-020-00273-3>

Published online: 29 January 2021

- Veblen, T. Why is economics not an evolutionary science? *Q. J. Econ.* **12**, 373–397 (1898).
- Simon, H. A. *Models of Man: Social and Rational* (Wiley, 1957).
- Boulding, K. Samuelson's foundations: the role of mathematics in economics. *J. Polit. Econ.* **56**, 187–199 (1948).
- Kirman, A. P. The intrinsic limits of modern economic theory: the emperor has no clothes. *Econ. J.* **99**, 126–139 (1989).
- Robinson, J. Time in economic theory. *Kyklos* **33**, 219–229 (1980).
- McCloskey, D. The trouble with mathematics and statistics in economics. *Hist. Econ. Ideas* **13**, 85–102 (2005).
- Schumpeter, J. *History of Economic Analysis* (Allen & Unwin, 1954).
- Waldrop, M. M. *Complexity* (Simon & Schuster, 1992).
- Mitchell, M. *Complexity: A Guided Tour* (Oxford Univ. Press, 2009).
- Holland, J. H. *Complexity: A Very Short Introduction* (Oxford Univ. Press, 2014).
- Thurner, S., Hanel, R. & Klimek, P. *Introduction to the Theory of Complex Systems* (Oxford Univ. Press, 2018).
- Arthur, W. B. Complexity and the economy. *Science* **284**, 107–109 (1999).
- Arthur, W. B. *Complexity and the Economy* (Oxford Univ. Press, 2015).
- Arthur, W. B. *Complexity and the Economy* 1–29 (Oxford Univ. Press, 2015).

- Arthur, W. B. in *Complexity Economics: Proceedings of the Santa Fe Institute's 2019 Fall Symposium* (eds Arthur, W. B., Beinhocker, E. & Stanger, A.) (SFI Press, 2020).
- Arthur, W. B., Beinhocker, E. & Stanger, A. in *Complexity Economics: Proceedings of the Santa Fe Institute's 2019 Fall Symposium* (eds Arthur, W. B., Beinhocker, E. & Stanger, A.) (SFI Press, 2020).
- Axtell, R. What economic agents do: How cognition and interaction lead to emergence and complexity. *Rev. Austrian Econ.* **20**, 105–122 (2007).
- Colander, D. *The Complexity Vision and the Teaching of Economics* (Edward Elgar, 2000).
- Farmer, J. D. *Economics Needs to Treat the Economy as a Complex System* (INET Conference Paper, 2012).
- Kirman, A. *Complex Economics: Individual and Collective Rationality* (Routledge, 2011).
- Rosser, J. B. On the complexities of complex economic dynamics. *J. Econ. Perspect.* **13**, 169–192 (1999).
- Epstein J. M. *Generative Social Science* (Princeton Univ. Press, 2006).
- Epstein, J. M. in *Handbook of Computational Economics 2. Agent-Based Computational Economics* (eds Tesfatsion, L. & Judd, K. L.) 1585–1604 (Elsevier, 2006).
- Louça, F. Bounded heresies. Early intuitions of complexity in economics. *Hist. Econ. Ideas* **18**, 77–113 (2010).
- Colander, D. in *Handbook of Research on Complexity* (ed. Rosser, J. B.) (Edward Elgar, 2009).
- Hayek, F. A. *The Theory of Complex Phenomena* 332–349 (Taylor & Francis, 1964).
- Samuelson, P. *Economics* (McGraw-Hill, 1967).
- Simpson, D. *The Rediscovery of Classical Economics* (Edward Elgar, 2013).
- Hayek, F. *Individualism and Economic Order* (Univ. Chicago Press, 1948).
- Harris, D. J. *Joan Robinson on "History versus Equilibrium"* (Joan Robinson Centennial Conference, 2003).
- Tabb, W. *Reconstructing Political Economy* (Routledge, 1999).
- Mirowski, P. *Machine Dreams: Economics Becomes a Cyborg Science* (Cambridge Univ. Press, 2002).
- Velupillai, K. V. in *Handbook of Research on Complexity* (ed. Rosser, J. B.) (Edward Elgar, 2009).
- Kirman, A. The economic crisis is a crisis for economic theory. *CEifo Econ. Stud.* **56**, 498–535 (2010).
- Komlos, J. *Foundations of Real-World Economics* (Routledge, 2019).
- Arrow, K., Anderson P. & Pines, D. *The Economy as an Evolving Complex System* (Addison-Wesley, 1988).
- Arthur, W. B. Complexity, the Santa Fe approach, and non-equilibrium economics. *Hist. Econ. Ideas* **18**, 149–166 (2010).
- Fontana, M. The Santa Fe perspective on economics. *Hist. Econ. Ideas* **18**, 167–196 (2010).
- Kirman, A. Whom or what does the representative agent represent? *J. Econ. Perspect.* **6**, 117–136 (1992).
- Knight, F. *Risk, Uncertainty, and Profit* (Houghton Mifflin, 1921).
- Keynes, J. M. The general theory of employment. *Q. J. Econ.* **51**, 209–233 (1937).
- Arthur, W. B. Bounded rationality and inductive behavior (the El Farol problem). *Am. Econ. Rev.* **84**, 406–411 (1994).
- Holland, J., Holyoak, K., Nisbett, R. & Thagard, P. *Induction* (MIT Press, 1986).
- Sargent T. J. *Bounded Rationality in Macroeconomics* (Clarendon, 1993).
- Friedman, D. & Rust J. (eds) *The Double Auction Market: Institutions, Theories, and Evidence* (Addison-Wesley, 1994).
- Thaler, R. H. *Misbehaving: The Making of Behavioral Economics* (Norton, 2015).
- LeBaron, B. Empirical regularities from interacting long and short memory investors in an agent-based financial market. *IEEE Trans. Evolut. Comput.* **5**, 442–455 (2001).
- Bronk, R. *The Romantic Economist: Imagination in Economics* (Cambridge Univ. Press, 2009).
- Beckert J. & Bronk, R. (eds) *Uncertain Futures: Imaginaries, Narratives, and Calculation in the Economy* (Oxford Univ. Press, 2018).
- Holland, J. H. & Miller, J. Artificial adaptive agents in economic theory. *Am. Econ. Rev.* **81**, 363–370 (1991).



51. Arthur, W. B., Durlauf, S. & Lane, D. (eds) *The Economy as an Evolving Complex System II* (Addison-Wesley, 1997).
52. Beinhocker, E. *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics* (Harvard Bus. School Press, 2006).
53. Arthur, W. B. in *Handbook of Computational Economics 2. Agent-Based Computational Economics* (eds Tesfatsion, L. & Judd, K. L.) 1551–1564 (Elsevier, 2006).
54. Blume, L. & Durlauf, S. *The Economy as an Evolving Complex System III* (Oxford Univ. Press, 2006).
55. Farmer, J. D. & Foley, D. The economy needs agent-based modelling. *Nature* **460**, 685–686 (2009).
56. Axtell, R. in *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems* (eds Thangarah, J., Tuyls, K., Jonker, K. & Marsella, S.) (International Foundation for Autonomous Agents and Multiagent Systems, 2016).
57. Lindgren, K. in *Artificial Life II* (Addison-Wesley, 1991).
58. Marshall, A. *Principles of Economics* (1890) 8th edn (Macmillan, 1920).
59. Palmer, R. G., Arthur, W. B., Holland, J., LeBaron, B. & Tayler, P. Artificial economic life: a simple model of a stock market. *Physica D* **75**, 264–274 (1994).
60. Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R. & Tayler, P. in *Economy as an Evolving Complex System II* (CRC Press, 1997).
61. Lucas, R. Asset prices in an exchange economy. *Econometrica* **46**, 1429–1445 (1978).
62. Brock, W. A., Lakonishok, J. & LeBaron, B. Simple technical trading rules and the stochastic properties of stock returns. *J. Finance* **47**, 1731–1764 (1992).
63. Hommes, C. H. in *Handbook of Research on Complexity* (Edward Elgar, 2009).
64. Kopel, M. in *Handbook of Research on Complexity* (Edward Elgar, 2009).
65. LeBaron, B., Arthur, W. B. & Palmer, R. Time series properties of an artificial stock market. *J. Econ. Dyn. Control* **23**, 1487–1516 (1999).
66. Galla, T. & Farmer, J. D. Complex dynamics in learning complicated games. *Proc. Natl Acad. Sci. USA* **10**, 1232–1236 (2013).
67. Tesfatsion, L. & Judd, K. L. (eds) *Handbook of Computational Economics 2. Agent-Based Computational Economics* (Elsevier, 2006).
68. Tesfatsion, L. Agent-based computational economics: a constructive approach to economic theory. in Tesfatsion & Judd, *Handbook of Computational Economics 2* (2006).
69. Tesfatsion, L. Modeling economic systems as locally-constructive sequential games. *J. Econ. Methodol.* **24**, 384–409 (2017).
70. Hommes, C. & LeBaron, B. (eds) *Handbook of Computational Economics, Vol IV: Heterogeneous Agent Modeling* (North-Holland, 2018).
71. Miller, J. & Page, S. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life* (Princeton Univ. Press, 2007).
72. Romer, P. M. Mathiness in the theory of economic growth. *Am. Econ. Rev.* **105**, 89–93 (2015).
73. Arthur, W. B. *Algorithms and the Shift in Modern Science. Beijer Institute Discussion Paper 269* (Swedish Academy Sciences, 2020).
74. Newman, M. *Networks: An Introduction* (Oxford Univ. Press, 2010).
75. Powell, W. & Padgett, J. *The Emergence of Organizations and Markets* (Princeton Univ. Press, 2012).
76. Root, H. *Network Origins of the Global Economy: East vs. West in a Complex Systems Perspective* (Cambridge Univ. Press, 2020).
77. Newman, M., Barabási, A.-L. & Watts, D. (eds) *The Structure and Dynamics of Networks* (Princeton Univ. Press, 2006).
78. Scheffer, M. et al. Anticipating critical transitions. *Science* **338**, 344–348 (2012).
79. Watts, D. A simple model of global cascades on random networks. *Proc. Natl Acad. Sci. USA* **9**, 5766–5771 (2002).
80. Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quant. Finance* **1**, 223–236 (2001).
81. Farmer, J. D. & Geanakoplos, J. The virtues and vices of equilibrium and the future of financial economics. *Complexity* **14**, 11–38 (2008).
82. Geanakoplos, J. et al. Getting at systemic risk via an agent-based model of the housing market. *Am. Econ. Rev.* **102**, 53–58 (2012).
83. Haldane, A. *Rethinking the Financial Network. Speech given at the Financial Student Association, Amsterdam* (Bank of England, 2009).
84. Poledna, S. & Thurner, S. Elimination of systemic risk in financial networks by means of a systemic risk transaction tax. *Quant. Finance* **16**, 1599–1613 (2014).
85. Hoogduin, L. *New Approaches to Economic Complexity: Insights into Complexity and Policy* (OECD, 2016).
86. Durlauf, S. N. Complexity, economics, and public policy. *Polit. Philos. Econ.* **11**, 45–75 (2012).
87. Kirman, A. Complexity and economic policy: a paradigm shift or a change in perspective? *J. Econ. Lit.* **54**, 534–572 (2016).
88. Colander, D. & Kupers, R. *Laissez-Faire Activism in Complexity and the Art of Public Policy: Solving Society's Problems from the Bottom Up* (Princeton Univ. Press, 2014).
89. Pyka, A. (ed.) *Handbook of Complexity Economics* (In Press) (Routledge, 2021).
90. Farmer, D. *The Complexity Economics Revolution* (In press) (Knopf, 2021).
91. Holt, R., Rosser, J. & Colander, D. The complexity era in economics. *Rev. Polit. Econ.* **23**, 357–369 (2010).
92. Rosser, J. B. (ed.) *Handbook of Research on Complexity* (Edward Elgar, 2009).
93. Wilson, D. S. & Kirman, A. (eds) *Complexity and Economics: Towards a New Synthesis for Economics* (MIT Press, 2016).
94. Miller, J. & Page, S. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life* (Princeton Univ. Press, 2007).
95. Kirman, A. in *Rethinking Economics* (eds Fisher, L. et al.) (Routledge, 2018).
96. Castaneda, G. *The Paradigm of Social Complexity* (CEES, 2019).
97. Hommes, C. H. *Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach* (American Economic Association, 2020).
98. Dawid, H. & Delli Gatti, D. in *Handbook of Computational Economics, Vol IV: Heterogeneous Agent Modeling* (eds Hommes, C. & LeBaron, B.) (North-Holland, 2018).
99. LeBaron, B. & Tesfatsion, L. Modeling macroeconomics as open-ended dynamic systems of interacting agents. *Am. Econ. Rev.* **98**, 246–250 (2008).
100. Axtell, R., Guerrero, O. & López, E. Frictional unemployment on labor flow networks. *J. Econ. Behav. Organ.* **160**, 184–201 (2019).
101. Axtell, R. L. in *Handbook of Computational Economics, Vol IV: Heterogeneous Agent Modeling* (North-Holland, 2018).
102. Grabner, C. The complementary relationship between institutional and complexity economics. *J. Econ.* **51**, 392–400 (2017).
103. Burgess, M. G. et al. Opportunities for agent-based modelling in human dimensions of fisheries. *Fish Fish.* **21**, 570–587 (2020).
104. Elsayaw, S. et al. Eight grand challenges in socio-environmental systems modeling. *SESMO 2*, 16226 (2020).
105. Preiser, R., Biggs, R., De Vos, A. & Folke, C. Social-ecological systems as complex adaptive systems: organizing principles for advancing research methods and approaches. *Ecol. Soc.* **23**, 46 (2018).
106. Bailey, R. M. et al. A computational approach to managing coupled human-environmental systems: the POSEIDON model of ocean fisheries. *Sustain. Sci.* **14**, 259–275 (2019).
107. May, R., Levin, S. & Sugihara, G. Complex systems: ecology for bankers. *Nature* **451**, 893–895 (2008).
108. LeBaron, B. Financial price dynamics and agent-based models as inspired by Benoit Mandelbrot. *Eur. Phys. J. Spec. Top.* **225**, 3243–3254 (2016).
109. Battiston, S. et al. Complexity theory and financial regulation. *Science* **351**, 818–819 (2016).
110. Arthur, R. F., Jones, J. H., Bonds, M. H., Ram, Y. & Feldman, M. *PLOS Comp. Biol.* (in the press).
111. Pichler, A., Pangallo, M., del Rio-Chanona, R., Lafond, F. & Farmer, J. D. Production networks and epidemic spreading: How to restart the UK economy? Preprint at [arXiv https://arxiv.org/abs/2005.10585](https://arxiv.org/abs/2005.10585) (2020).
112. Axtell, R. *Dynamics of Firms: Data, Theories and Models: Emergent Scale and Complexity in Modern Economics* (MIT Press, 2020).
113. Barabási, A.-L. & Albert, R. Emergence of scaling in complex networks. *Science* **286**, 509–512 (1999).
114. Malvergne, Y., Saichev, A. & Sornette, D. Zipf's law and maximum sustainable growth. *J. Econ. Dyn. Control* **37**, 1195–1212 (2013).
115. Peters, O. The ergodicity problem in economics. *Nat. Phys.* **15**, 1216–1221 (2019).
116. Arthur, W. B. *The Nature of Technology: What it is and How it Evolves* (Simon & Schuster, 2009).
117. Arthur, W. B. The structure of invention. *Res. Policy* **36**, 274–287 (2007).
118. Arthur, W. B. & Polak, W. The evolution of technology within a simple computer model. *Complexity* **11**, 23–31 (2006).
119. Hildago, C. & Hausmann, R. The building blocks of economic complexity. *Proc. Natl Acad. Sci. USA* **106**, 10570–10575 (2009).
120. Lo Sardo, D. et al. Quantification of the resilience of primary care networks by stress testing the health care system. *Proc. Natl Acad. Sci. USA* **116**, 23930–23935 (2019).
121. Davis, J. The turn in recent economics and return of orthodoxy. *Camb. J. Econ.* **32**, 349–366 (2008).
122. Koppl, R. in *Handbook of Complexity Research* (ed Rosser, J.B.) (Edward Elgar, 2009).
123. Mantegna, R. & Stanley, H. E. *An Introduction to Econophysics* (Cambridge Univ. Press, 2000).
124. Bouchaud, J.-P., Bonart, J., Donier, J. & Gould, M. *Trades, Quotes and Prices: Financial Markets Under the Microscope* (Cambridge Univ. Press, 2018).
125. Buchanan, M. What has econophysics ever done for us? *Nat. Phys.* **9**, 317 (2013).
126. Hanauer, N. & Beinhocker, E. Capitalism redefined. *Democracy* **31**, 30–44 (2014).
127. Krugman, P. Is free trade passé? *J. Econ. Perspect.* **1**, 131–144 (1987).
128. Scott, R. E. *The High Price of "Free" Trade. Economic Policy Institute Briefing Paper 147* (Economic Policy Institute, 2003).
129. Dean, A. & Kimmel, S. Free trade and opioid overdose death in the United States. *SSM Popul. Health* **8**, 100409 (2019).
130. Davidsson, P. et al. in *Applications of Agent Technology in Traffic and Transportation* (eds Klügl F., Bazzan, A. & Ossowski S.) (Birkhäuser, 2005).
131. Chen, B. & Cheng, H. A review of the applications of agent technology in traffic and transportation systems. *IEEE Trans. Intell. Transp. Syst.* **11**, 485–497 (2010).
132. Boulton, J., Allen, P. & Bowman, C. *Embracing Complexity: Strategic Perspectives for an Age of Turbulence* (Oxford Univ. Press, 2015).
133. Arthur, W. B. Where is technology taking the economy? *McKinsey Q. Fall* **697**, 33–43 (2017).
134. Tesfatsion, L. *A New Swing-Contract Design for Wholesale Power Markets* (Wiley/IEEE Press, 2020).
135. Brittain, M. & Wei, P. Autonomous air traffic controller: A deep multi-agent reinforcement learning approach. Preprint at [arXiv https://arxiv.org/abs/1905.01303v1](https://arxiv.org/abs/1905.01303v1) (2019).
136. Wolfram, S. *A New Kind of Science* (Wolfram Media, 2002).
137. Fontana, M. Can neoclassical economics handle complexity? The fallacy of the oil spot dynamic. *J. Econ. Behav. Organ.* **76**, 584–596 (2010).
138. Robertson, D. S. *Phase Change: the Computer Revolution in Science and Mathematics* (Oxford Univ. Press, 2003).
139. Lindgren, K. & Nordahl, M. G. Evolutionary dynamics of spatial games. *Physica D* **75**, 292–309 (1994).
140. Arthur, W. B. *Complexity and the Economy* 103–118 (Oxford Univ. Press, 2015).
141. Boyes, R. *Meltdown Iceland: How the Global Financial Crisis Bankrupted an Entire Country* (Bloomsbury, 2009).
142. Black, B., Kraakman, R. & Tarassova, A. Russian privatization and corporate governance: what went wrong? *Stanford Law Rev.* **52**, 1731–1808 (2000).
143. Cassidy, J. *How Markets Fail: The Logic of Economic Calamities* (Farrar, Straus & Giroux, 2009).
144. Sweeney, J. *The California Electricity Crisis* (Hoover Inst. Press, 2002).
145. Colander, D. et al. The financial crisis and the systemic failure of the economics profession. *Crit. Rev.* **21**, 249–267 (2009).
146. Bibel, G. *Beyond the Black Box: The Forensics of Airplane Crashes* (Johns Hopkins Univ. Press, 2008).
147. Thurston, R. *A History of the Growth of the Steam Engine* (Appleton & Co., 1878).
148. Schumpeter, J. *The Theory of Economic Development* (Oxford Univ. Press, 1912).

149. Perez, C. *Technological Revolutions and Financial Capital* (Edward Elgar, 2002).
150. Arthur, W. B. Competing technologies, increasing returns, and lock-in by historical events. *Econ. J.* **99**, 116–131 (1989).
151. Arthur, W. B. *Increasing Returns and Path Dependence in the Economy* (Univ. Michigan Press, 1994).
152. Tetzeli, R. *A Short History of the Most Important Theory in Technology* (Fast Company, 2016).
153. Krugman, P. Increasing returns and economic geography. *J. Polit. Econ.* **99**, 483–499 (1991).
154. Helpman, E. & Krugman, P. *Market Structure and Foreign Trade* (MIT Press, 1985).
155. Asano, Y., Jakob, J., Kolb, J., Heitzig, J. & Farmer, J. D. Emergent inequality and endogenous dynamics in a

simple behavioral macroeconomic model. Preprint at *arXiv* <https://arxiv.org/abs/1907.02155> (2019).

156. Durlauf, S. A theory of persistent income inequality. *J. Econ. Growth* **1**, 75–93 (1996).
157. Schelling, T. Models of segregation. *Am. Econ. Rev.* **59**, 488–493 (1969).

#### Acknowledgements

The author thanks the Center for Advanced Study in the Behavioral Sciences at Stanford University, where he was a Fellow during the writing of this paper. He also thanks the members of SFI's early Economics Program, particularly John Holland, Stuart Kauffman, David Lane, Blake Le Baron and Richard Palmer. For comments and information, he thanks Ronan Arthur, Robert Axtell, Eric Beinhocker, Richard Bronk, David Colander, Doyne Farmer, Cars Hommes, Alan Kirman,

Blake Le Baron, Kristian Lindgren, Scott Page, Rika Preiser, Leigh Tesfatsion and Stefan Thurner.

#### Competing interests

The author declares no competing interests.

#### Peer review information

*Nature Reviews Physics* thanks Kenneth Judd, Andreas Pyka and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

#### Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© Springer Nature Limited 2021