Oxford Review of Economic Policy, 2025, 41, 571–590
https://doi.org/10.1093/oxrep/graf027
Article

Quantitative agent-based models: a promising alternative for macroeconomics

J. Doyne Farmer

Institute for New Economic Thinking at the Oxford Martin School, Smith School for Enterprise and the Environment, UK, Santa Fe Institute, and Macrocosm Inc, E-mail: dorothy.nicholas@inet.ox.ac.uk

I would like to thank Baillie Gifford for supporting my professorship at Oxford and David Vines and Alison Gomm for their editorial assistance.

Abstract

Agent-based models (ABMs) are dynamic computer simulations that abandon utility maximization and instead assume that agents are boundedly rational and make decisions using heuristics, myopic reasoning, and/or learning algorithms. Because ABMs do not need to compute optima they are more tractable, allowing a higher level of realism. Recent research has developed quantitative agent-based models that make time series predictions, modelling a specific economy at a specific point in time; some of these address questions that mainstream models cannot even ask, and some make predictions that are superior or equal to their mainstream equivalents. After explaining what ABMs are and how they are built in more detail, I review four examples of models from my own work for leverage cycles, the 2008 housing bubble, Covid, and a general-purpose micro-macro model. I conclude by discussing the advantages and disadvantages of agent-based models in comparison to standard models.

Keywords: agent-based modelling, complexity economics, macroeconomics, financial stability.

JEL codes: C18, D58, D83, E14, E17

I. Introduction

Since rational expectations began to be adopted in economics in the 1960s, there has always been an alternative strand of thought that we should also be exploring approaches based on Herb Simon's notion of bounded rationality. This naturally leads to what is called 'agent-based modelling', which is based on computer simulation of boundedly rational agents rather than utility optimization within a purely mathematical framework. Over the last 60 years quite a lot of progress has been made, but until recently there were no quantifiable successes based on empirical benchmarks. On a recent panel, a Nobel prize-winning economist whom I shall not name said (more or less) 'agent-based modelling has been around for a long time, but there is still very little success to show for the effort'. My goal here is to convince you that this economist was speaking out of ignorance. Computers are now a billion times more powerful than they were in Herb Simon's day, there is vastly more data, and we know a great deal more about human behaviour. Agent-based models are the natural way to take advantage of this new power. Among other things, they provide a flexible modelling framework to take advantage of advances in behavioural economics, which is only beginning to be explored. One of the new developments is *quantitative* agent-based models (ABMs) that make time series predictions about a specific economy at a specific point in time. These efforts are beginning to be successful, even though the person-hours invested are tiny compared to the effort invested in creating the corresponding mainstream tools. These models provide proofs of principle for the promise of agent-based modelling. In this paper I explain the advantages of agent-based models and discuss how their disadvantages could be overcome by pouring more effort into this area. In doing so I hope to convince you that agent-based models deserve more attention from the mainstream.

This paper is organized as follows: I first explain what agent-based models are and how they differ from standard models in economics. I then present four examples, including a very simple, stylized toy model for leverage cycles in the financial system, a highly realistic model for the Washington DC housing market, a model that was successfully used to predict the economic impacts of the Covid pandemic and a detailed quantitative micro-macro model whose predictions are roughly as accurate as those of mainstream macro models. I then discuss the pros and cons of agent-based models and conclude with some reflections about why they have not been accepted by the mainstream in academic economics, even though they are gaining some traction in industry and with avant-garde central banks and other policy-makers.

I want to emphasize that this is a very incomplete review, largely focused on my own work. For a more comprehensive survey see Axtell and Farmer (2025), or for something easy to digest see Farmer (2024).

II. What is an agent-based model?

Agent-based modelling provides a way to do economics that is radically different from the prevailing mainstream approaches. An agent-based model is a simulation of the economy, which is typically constructed like this:

1. Identify the essential agents to be modelled. Assign each agent a method for making decisions. This can involve heuristics, simple reasoning, search, learning, AI, or a mixture of any of these. When learning algorithms are used, goals need to be assigned, but these need not satisfy the requirements of utility functions.
2. Mimic the institutions, market mechanisms, and other economic factors of interest in software, with as much verisimilitude as necessary (remembering that it can be counterproductive to introduce details that don't matter).
3. Simulate the collective interactions of the agents as time unfolds.

Agent-based models deviate from standard equilibrium-based economic models in several fundamental ways. Perhaps the biggest point of departure is in step (1), the model for agent decision-making. There is a great deal of flexibility in how this is done, but the key point of departure is that there is no assumption of optimality. The agents may *attempt* to optimize something, which might even be a standard utility function, but there is typically no requirement that they make decisions that achieve the optimum. In very simple situations, the agents may find optimal decisions, but as in the real world, this is very unlikely when things get more complicated. Agent-based models are powerful precisely because they make it possible to study situations where it is impossible to compute optimal decisions (as is typically the case in reality). Thus, they may incorporate goals along with learning algorithms that *attempt* to achieve those goals, but there is no assumption that the goals are achieved.

Here are some of the ways agent-based modellers incorporate decision-making:

- *State the obvious*. There are many circumstances for which there is an obvious decision rule. For example, a class of stock traders called *value investors* (think Warren Buffett) follow the rule 'buy undervalued assets'.
- *Ask experts,* like psychologists and sociologists, who have done lab experiments or field studies or performed surveys to understand how people made decisions in a similar context.
- *Study data where similar decisions have been made*. For example, in our housing model for Washington DC, discussed later, we used data from the US census on people's housing choices, conditioned on income and wealth. In many cases ABMs use econometrics from micro-data (but see the later discussion about the Lucas critique).
- *Do experiments in the context of the model*. One can potentially slow down an agent-based model to human speed, unplugging the computer's decision-making algorithms and replacing them with the decisions of real people, and then build a model mimicking the

choices of the human decision-makers. This is a promising approach that remains largely unexplored.[1]

- *Use learning algorithms* to train agents to make better decisions through time. This is the closest thing to a one-size-fits-all approach and is the most powerful way for agent-based models to answer the Lucas critique. Although there are many learning algorithms, they often lead to similar behaviours.

Many current agent-based models endow agents with simple heuristics, such as 'buy undervalued assets', 'imitate the best', or 'trial and error'. Using simple heuristics is an example of a more general approach that I call 'modelling behaviour directly'. As illustrated in the examples that are given later, sometimes such simple decision-making rules can be very effective in mimicking real-world behaviour, but in other cases this is inadequate. Note also that all three of these heuristics automatically adapt to new regimes, suggesting they may pass the Lucas critique in many different circumstances.

The other radical departure from mainstream modelling is step (3), simulation. The interactions of agents are explicitly mimicked by the computer. This can be done via equations, e.g. that might specify how much an agent will buy or sell or using algorithms that might be difficult to render with equations. Although algorithms may be written as mathematical expressions, at least in principle, they are not necessarily equations, i.e. in general they are not statements that something equals something else. Consider, for example, a heuristic call *adaptation level adaptation*, which is widely used for selling goods such as used cars or houses (let's assume it's a house). When a house is put on the market, the seller looks for recent sales of comparable houses and marks her house up from this price by a small amount, typically about 5%. Then if the house doesn't sell after a month or so, the seller marks it down by about 10%, and continues marking it down until either the house sells, or the seller decides to take it off the market. This is easily written in software as an iterative expression involving 'if–else' statements and inequalities, and while it can in principle be written as a mathematical expression involving a combination of logical statements and real numbers, it is much more easily written as computer algorithm, and is very different from the equations that are easily incorporated into DSGE models. (See the later discussion of the housing market model.)

Algorithms are also more flexible for describing market institutions. Consider, for example, the dynamics of the continuous double auction, which in common parlance is called a 'limit order book', which is widely used in financial markets. Agents can submit different types of orders, such as market orders, which generate immediate transactions, or limit orders, which accumulate, and they can also cancel limit orders. It is difficult to implement the highly state-dependent dynamics of the limit order book with equations, but very easy to implement them with a computer algorithm.

This differs from standard practice in mainstream economics. While it is now common to use computers to solve mathematical equations, e.g. to find the decisions that maximize utility, this is very different from directly simulating the processes in the economy using algorithms, as agent-based models do. The use of simulation is much more flexible and makes it possible to capture economic institutions with more verisimilitude than is possible using mathematical equations and solving first-order conditions.

By *verisimilitude*, I mean that an attempt is made to model the economy as it is, rather than as if it were something else. In other words, an attempt is made to mimic economic processes realistically. This doesn't necessarily mean literally: good modellers focus on the effects that matter and omit those that are irrelevant. Models should be as simple as possible, though no simpler. The goal is to capture the essence of how things really work. Verisimilitude just means that it is usually much better to model the way the economy actually works, rather than to pretend it works another way. This contrasts with the 'as if' arguments that are frequently used to justify assumptions in economic models that are manifestly wrong but might nonetheless yield reasonable predictions. In this paper I provide several examples of how agent-based models use assumptions with more verisimilitude than those of standard models.

[1] Cars Hommes and his collaborators run economic experiments and explain the results of these experiments using simple agent-based models. See Hommes (2013).

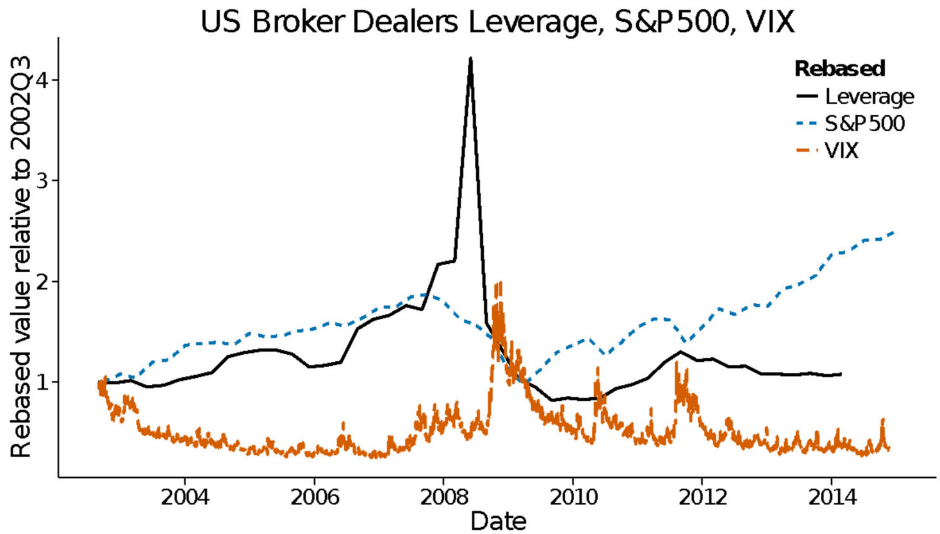

Figure 1. Comparison of the US stock market (dashed), broker-dealer leverage (solid), and volatility (jagged).
Note: The scales are adjusted so they are all equal to one in 2002.

Agent-based models are explicitly dynamic in a stronger sense than dynamic equilibrium models, such as DSGE models. At each timestep the agents in the model gather information and make decisions. The computer then calculates the economic consequences of these decisions, which may generate new information. This is then added to any exogenous information processes (shocks), and the agents once again make decisions. This simulation goes around and around its loop until the simulation stops. There is no assumption of equilibrium (at least in the game theoretic sense), though in some cases the simulation may converge to an equilibrium. There is a great deal of flexibility in how simulations are done; for example, the agents might synchronously make their decisions at each timestep, without knowing the decisions of the other agents, or the simulation may proceed asynchronously, with agents making decisions one at a time, updating the state of the world after each agent's decision.

I now give a few examples of agent-based models, drawing on my own work. Agent-based models can be simple toy models, designed to illustrate a conceptual point, or they can be complicated, striving to achieve realism, e.g. to make quantitative predictions or evaluate a policy.

Let's begin with a very simple, conceptual agent-based model designed to understand how leverage can cause systemic risk, leading to booms and busts in financial markets.

III. The Basel leverage cycle

I now describe a model that was designed to understand the role of leverage and risk control in the Great Financial Crisis of 2008 (Aymanns *et al*., 2016). Figure 1 compares US broker-dealer leverage to the stock market beginning in 2002. Over the next 6 years the price of the S&P index and the leverage steadily climb while volatility falls. In 2008 the leverage soars, the market crashes, and volatility rises dramatically. The model illustrates that this can be explained by a simple and plausible mechanism.

The model uses a standard framework, with two representative agents, corresponding to two types of investors: an investment bank that uses leverage and a value investing fund that does not use leverage. They trade a single asset, whose supply is fixed—their alternative is to hold cash. They are both given initial endowments of the asset and cash, and then they begin to trade. At each timestep, both the fund and the bank evaluate their demand for the asset, the price is set so that supply equals demand and the process is repeated. There's also a random shock that alters the fund's estimate of the fundamental value, so that it is always changing in time according to

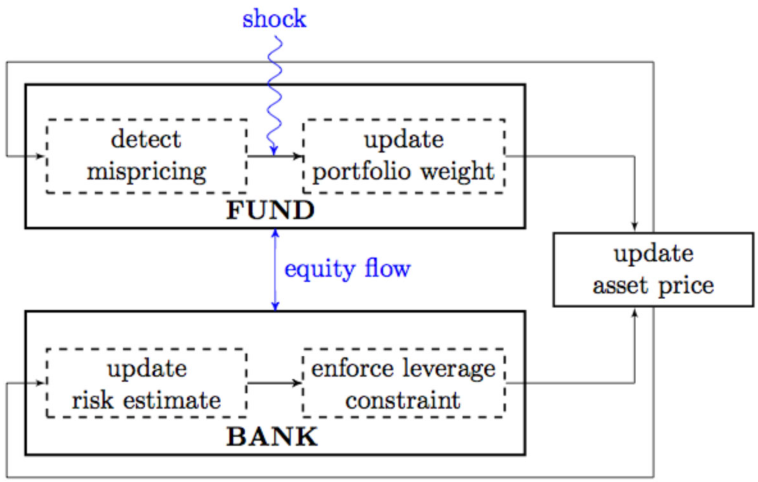

Figure 2. A simple schematic diagram of the Basel leverage-cycle model.
Note: There are two agents: the *fund*, which represents all unleveraged value investors, and the *bank*, which represents all investors with leverage targets. The equity flow is from their buying and selling of the asset, and the shock represents outside noise that randomly influences the valuation perceived by the fund.

a random walk—the resulting shocks can be turned up or down or even set to zero. The model can be written as either a discrete or continuous dynamical system, with seven equations. When the solution converges to a fixed point it can be solved algebraically, but in many situations, as discussed below, the solution does not converge to a fixed point, even in the absence of noise. A schematic of the model is shown in Figure 2.

The strategies of the two agents consist of simple but widely used heuristics. The fund is a value investor who estimates the value of the asset and buys and holds it when it's undervalued and sells it and holds a short position when it's overvalued—the size of the positions is proportional to the mispricing. The bank is an investment bank that has a leverage target and adjusts its leverage according to the recommendations of Basel II, which are based on estimating value-at-risk. When expected volatility is high, leverage is lowered, and when expected volatility is low, it is increased. Following the widespread practice in financial markets, the estimate of future volatility is backward looking, based on an exponential moving average of past volatility.

The model has a few parameters that control its behaviour. Two key parameters, whose effects on the dynamics are similar, are the relative wealth of the investment bank and the value fund and the risk appetite of the investment bank. If the fund has most of the wealth and/or if the risk appetite of the investment bank is low, then the dynamics are stable. In the zero-noise limit, this means that the price of the asset converges to its fundamental value and the leverage and volatility converge to constant values.

In contrast, if the bank has more wealth, and/or if the risk appetite of the investment bank is sufficiently high, the system begins to oscillate. During the cycle the stock price and leverage slowly build, while the volatility slowly drops. Then the market suddenly crashes and the volatility spikes sharply upward. Just before the crash the leverage spikes upward, then its drops sharply just after the crash, as shown in Figure 3.

One of the nice features of the model is that sensible a priori parameters yield quantitatively plausible results. The scale of the horizontal axis in Figure 3 is measured in years. The frequency of the oscillations depends on the time-window used for the exponential moving average of volatility; with a 2-year window, which is the industry standard, the cycle takes about 10 years, roughly equivalent to the period of smooth growth in the build-up to the great financial crisis of 2008.

Surprisingly, *the oscillations in the unstable phase occur even if the shocks are turned off completely*. This is because the dynamics are chaotic, meaning that the system shows sensitive dependence on initial conditions (corresponding to the exponential separation of nearby states) and

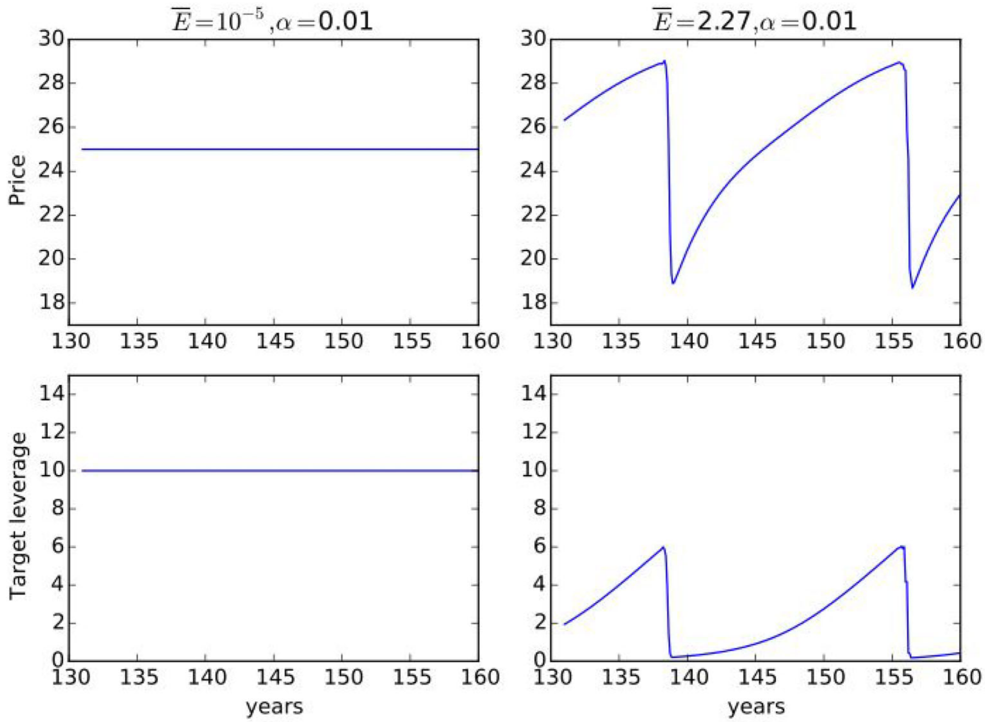

Figure 3. An illustration of the Basel leverage cycle in the deterministic limit, where the shocks are set to zero. *Note*: The parameter $\bar{E}$ is the target equity of the bank, and the parameter α is the risk aversion of the bank—higher α means more risk, which results in higher leverage. 'Target' refers to the fact that the bank bases its decisions on its volatility forecast and the current price, but these may change (in practice the true and target leverage and price are usually very close to each other). On the left, parameters are set so that the banking sector is relatively small, whereas, on the right, the banking sector is relatively large; we could have increased the bank's risk parameter α and achieved the same effect. The scale of the horizontal axis is measured in years. When the dynamics are unstable, as they are on the right, there is a steady climb in the price of the stock, followed by a sharp collapse when the leverage gets too large. This cycle repeats itself, but does so chaotically, with peaks whose amplitudes vary irregularly from cycle to cycle.

there is an endogenous oscillation, even with no external noise, as shown in Figure 4. Chaotic attractors have the property that the dynamics are locally unstable but globally stable. The state is confined to the attractor, but moves around it in an irregular manner. The magnitude of the oscillations varies irregularly, as does the time interval between them, as one would expect from a system with chaotic dynamics.

To understand the cause of this behaviour more intuitively it helps to trace the behaviour as it moves around the attractor. Shortly after the system crashes, the price and leverage are low and the volatility is high. Because the bank uses a moving average of historical volatility to estimate future risk (a smoothing operator), after the crash prices change slowly and volatility slowly drops. This causes the bank to raise its leverage, driving the price up. Increasing leverage makes the system less and less stable, until it becomes unstable and crashes.

The dependence on the size of the banking sector can be understood in terms of two competing effects. The fund buys when prices drop, stabilizing the market. The bank, in contrast, has an inverted demand function: when prices drop its leverage increases (by definition). This forces it to sell the asset. It is thus selling into a falling market, which is destablizing. When the combination of the size of the banking sector and the bank's leverage is too large, the market is unstable and its behaviour is chaotic. The nonlinear behaviour of the system is such that, though the system is locally unstable, it is globally stable, which keeps it on the chaotic attractor seen in Figure 4.

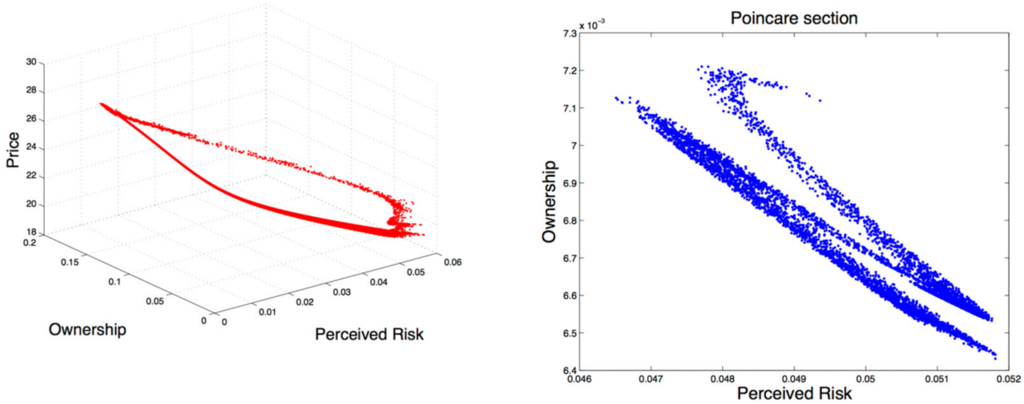

Figure 4. Underlying chaotic dynamics of the Basel leverage cycle.
Note: The panel on the left shows how the system moves clockwise (from the top) through a three-dimensional projection of its phase space, consisting of the stock ownership by the bank, the price of the risky asset, and the risk (volatility) perceived by the bank, which is inversely proportional to the leverage. (When the perceived risk is high the bank lowers it leverage, and vice versa.) The crash occurs when the bank ownership, leverage, and the price are at their maximum, in the upper left in the left figure. A dot is plotted each time step, so one can see speed of motion. As the crash happens there are distinct dots, corresponding to very rapid motion, vs the region on the bottom, where the system moves so slowly that the dots are blurred together. The panel on the right is what is called a Poincaré section, which is made by plotting the ownership vs the perceived risk whenever the price rises above a given threshold; the result is as if one took a slice through the attractor in the figure on the left. This shows the characteristic fractal structure of stretching and folding that always underlies chaotic behaviour.

The transition between the stable and unstable modes of behaviour is sharp: the system jumps discontinuously from a fixed point to a finite amplitude chaotic oscillation. This creates a danger for the financial system—if circumstances change, e.g. because the bank gains more assets, even a very small change can cause the chaotic oscillation to suddenly appear at full size.

The Basel leverage cycle model discussed above illustrates the complementarity of standard economics and agent-based modelling. The model was inspired in part by John Geanakoplos's theory of leverage cycles (Geanakoplos, 2003). His model is an equilibrium-based three-period model with utility-maximizing heterogeneous investors, ranging from optimistic to pessimistic. The Basel leverage cycle model introduces dynamics and shows that systemic financial risk can be caused by prudent investors—it does not necessarily require heterogeneity. It differs from the leverage cycles envisioned by Minsky in that greed enters only indirectly, through the use of leverage by the bank. In fact, it is (boundedly rational) prudence that causes the problem: in the Basel leverage cycle model investors follow institutionally sanctioned risk management. The crash occurs because this leads them to collectively sell into falling markets, behaving as if they had inverted demand functions. This is a good example of what is called an emergent phenomenon, in that it is not obvious a priori that this can lead to a chaotic endogenous oscillation.

To summarize, this model uses a standard framework in finance of two investors who choose between a risky and a non-risky asset. There are four key assumptions:

1. The investment bank uses value-at-risk (Basel II) to regulate its leverage.
2. The investment bank uses an exponential average of historical volatility to estimate future volatility.
3. The value investor holds the asset when it is undervalued and shorts it when it is overvalued.
4. The price is set by letting supply equal demand.

The first three assumptions provide good examples of heuristics, i.e. simple rules that agents follow, either because they are useful or because they are mandated. When I present this model to an economics audience, if I want to rib them a little, I point out that assumption (4), supply equals demand, has the weakest empirical support. My point is that it is possible to construct models using heuristics that are empirically well-grounded in known behaviour. The only exception is that there is some arbitrariness in the precise way the strategy of the value investor is constructed in assumption (3), but numerical experiments make it clear that the behaviour of the model is robust with respect to this.

To close the discussion of this model, I want to emphasize that *bounded rationality* is key. The non-optimality of the agents' rules of thumb are an important feature of the model. When Christoph Aymanns and I presented this model at the OECD, a discussant from the BIS objected, correctly pointing out that the chaotic oscillations would never occur if the bank followed rational expectations. But our response was to ask whether this is a defect or a virtue? Does anyone really think that typical investors understood what was about to happen in the run-up to 2008? At a high level, this model illustrates that the bounded rationality of real agents is important, and is necessary to achieve verisimilitude. As we discuss later, bounded rationality often leads to endogenous oscillations in systems where it is not present with rational agents. (See Farmer (2024) for several examples and an extended discussion; Asano *et al*. (2021) is particularly relevant.)

IV. Housing markets

I now describe an agent-based model which sits at the other end of the spectrum, in that the model is much more complicated, and is calibrated carefully against empirical data to be as realistic as possible. The original model for Washington DC was built in a project led by Rob Axtell, John Geanakoplos, Peter Howitt, and me, supervising a long list of graduate students and postdocs (Geanakoplos *et al*., 2012; Axtell *et al*., 2014). This model mimics the institutional structures for buying and selling houses in the real world and employs rules of thumb that people use to make decisions in housing markets, employing a mixture of forward- and backward-looking expectations. Using detailed data on real-estate transactions and mortgages, and information from the IRS and the US census, our model reproduces all the steps of buying or selling a house. It simulates the decisions of households to rent or buy and their interactions with banks and real-estate agents in taking out loans and buying and selling houses with as much realism as possible. The result is a model that was able to quantitatively reproduce the housing bubble in Washington DC reasonably well.

Agent-based modelling in economics has a long history of qualitative agent-based macro models, which study the behaviour of a hypothetical generic economy at a generic point in time. These models have made valuable contributions to understanding the likely effect of policies on economic outcomes. For a more extensive discussion see Axtell and Farmer (2025). Our housing market model is perhaps the first quantitative agent-based model in economics. Constructing it required considerable computer programming and data ingestion, but at the same time, it is conceptually simple—the way it works will be familiar to anyone who has ever bought a house. Calibrating our model with detailed data on real-estate transactions and mortgages, and information from the IRS and the US census, we simulated the decisions of households to rent or buy and their interactions with banks and real-estate agents in taking out loans and buying and selling houses. To make the most realistic possible simulation of how houses are bought and sold, we modelled details such as how people get loans from banks and how real-estate agents help them find the right house. The agents in our model were households, who could be homeowners or renters. We simulated the behaviour of 10 000 households in small runs and up to 100 000 in larger runs—fewer than the actual number of households in Washington but enough to provide a good sample. We created artificial households with different levels of income and wealth, so that they were representative of the actual households in the Washington DC area in roughly the right proportions.

The central question that each of our households faced was, 'Should I rent or buy?' This decision is of course strongly influenced by whether or not households expect prices to rise. In our model of the decision-making of prospective buyers, each household estimated future house prices based on the recent trend; if house prices have been going up, the households assumed they would keep going up, at least for a while. The fact that people make house-buying decisions this way is well supported by empirical studies. At the same time, our households thought about how future prices would affect them. Thus, our model for prospective buyers used a mixture of forward- and backward-looking expectations, based on research supported by data and common sense.

We also simulated the process of taking out a housing loan. When our simulated households decided to buy a house, unless they were rich enough to buy the house outright, they went to a simulated bank to ask for a loan. The bank would consider each loan application, consider the household's income and wealth, and accept or reject the application accordingly. The simulated bank offered several kinds of loans, ranging from simple, old-fashioned fixed-rate 30-year mortgages to complicated, newfangled loans with variable rates and balloon payments. This was important: in the build-up to the actual crisis, lending policy became more liberal, so that it was easier to get a loan and the terms were more generous, often with higher leverage and a higher proportion of complicated loans. We were able to acquire a detailed data set with a record of almost all mortgages in Washington, so we could adjust our bank algorithm on a year-by-year basis to match the types of mortgages that were actually issued and the lending criteria that were in place at the time. We could also run counterfactual scenarios, for example holding the types of loans constant. Similarly, the interest rates at which loans were given could be adjusted to match historical rates, or we could make up hypothetical interest rates to assess their importance in driving the bubble.

We were able to acquire data based on tax returns that told us how many people in a given age and income range entered or left Washington each year. In the course of our simulation, for each year, we could introduce new households (who needed to buy or rent new houses) or put new houses on the market when they moved away from Washington. We were able to check our predictions against detailed real-estate records.

At the beginning of the simulation, we assigned households diverse incomes and wealth to match the real data in 1997 and created a range of houses for them to live in, corresponding to the range of values of existing houses. Then, for each household, we simulated subsequent savings and consumption behaviour. Most households' incomes tended to slowly increase, but unlucky households could become unemployed and default on their mortgages.

Agent-based models have a big advantage for studying housing markets because of the way the market works, which is not at all close to market clearing. Housing markets operate through bilateral trading. Sellers set the offering price via aspiration-level adaptation, as described earlier. They choose an initial price and gradually lower it until either the house sells or they remove it from the market. Meanwhile buyers search for houses in their preferred price range, which depends on their income and wealth, and houses are removed from the market when there is a transaction. We carefully analysed a comprehensive record provided by real estate agents that makes it clear that houses are sold using aspiration-level adaptation and implemented this in the model.

The housing market in our model was constructed to be as realistic as possible. Each month a given number of new houses were chosen at random to be put on the market (using a constant rate in time) and additional houses might be put on the market because the owners defaulted on their mortgage payments. These were added to the unsold houses already on the market. The rate at which new buyers entered depended on several factors, including household income and savings, perceptions about future housing prices, interest rates, and the availability of bank credit. Then we cleared the market, matching as many houses as possible based on the offered selling prices and the buyers' reservation prices. This typically leaves behind many unsold houses and unsatisfied buyers, meaning the market did not clear. The operation of housing markets is hard to write in mathematical form but easy to simulate—one of many examples of how agent-based models allow greater verisimilitude.

The fact that houses are sold via aspiration-level adaptation means that there can be large imbalances in supply and demand—housing markets often operate very far from market clearing. If there are more sellers than buyers, then some of the sellers are left without making a transaction, and if there are more sellers than buyers, some of the buyers are left without making a transaction. During the 2008 housing bubble, there were times when there was more than an order of magnitude more buyers than sellers, and after the bubble burst, the reverse was true. (This was true both empirically and in our model.) One consequence of this, observed in real markets, is that prices respond sluggishly, which means that assuming that supply equals demand is a poor approximation.

We also modelled the rental market. This meant we had to model landlords, who buy houses to rent to others. This part of our model was less accurate, because we couldn't get good data on rental prices.

We simulated the model on a monthly timescale. We began by initializing the model to match conditions at the beginning of 1997. Each month, the model would run through a series of steps like this:

1. Demography is updated, so that new households enter and leave the area to match the historical record.
2. The banks survey all their existing loans and make foreclosure decisions.
3. Households are updated. This requires a long list of possible actions and decisions for each household, including: receive income; consume; make a rental or mortgage payment; decide whether or not to buy a new house; decide whether or not to sell or delist a house that is already on the market; apply for a loan or refinance an existing loan.
4. The banks consider loan applications and offer terms for mortgages.
5. Houses and buyers enter the market, and houses are bought and sold.

If we allowed ourselves to adjust parameters in-sample, it was easy to fit the data well. To address the problem of overfitting, we set the free parameters for each part of the model separately, using microdata, and without looking at the time series we wanted to fit. For example, we adjusted the parameters for how the households set their housing budgets by looking at census data (which recorded how much individual households actually spent), and the parameters for their strategies for selling houses by looking at real-estate quotes, and so on. This was a lot of work—even more than constructing the original model. Then we put all the modules together and pushed Run. We started the model in 1997 and let it run without any intervention until 2010. The last value in the simulation was, in a sense, a 13-year-ahead prediction, though of course it was strongly influenced by the exogenous inputs. To test the validity of our model, we then compared our simulation's prediction to what actually happened in the Washington housing market during those 13 years. The results are shown in Figure 5.

The plots shown in Figure 5 consider nine different features of the housing market. The match is by no means perfect; for example, the simulated bubble peaks before the real bubble and prices go a little higher. The model doesn't fit every feature perfectly, but most of the time it's in the ballpark. It is important to look at the scales on the plots—for home ownership rates, for example, we predicted rates around 60%, whereas in reality they varied between 65 and 70%. So we weren't doing badly given that we were missing the data needed to calibrate some of the features, such as the rental market.

In addition to trying to match reality, we also did some counterfactual experiments. When we held interest rates constant at the initial (high) value, for example, the bubble in our simulation was damped. This was expected; low interest rates reduce monthly payments and so create more demand for housing. We also experimented with holding lending policy constant, with banks continuing to issue the same type of loans they were making in 1997, which were mostly fixed-rate with down payments of 20% or more. The result was that the housing bubble almost disappeared, which confirmed our hypothesis that shifts in lending policy played a bigger role in causing the housing bubble than other factors, such as interest rates. The important role of lending policy is understood now, but it was not well established at the time we developed our model.

This model was subsequently simplified and further developed in collaboration with the Bank of England and used to provide advice about setting lending policy in the UK (Baptista *et al.*, 2016; Carro *et al.*, 2022). Models of this type are now used by many European central banks for analysing lending policy in the housing market.

The main conclusion of our model was that lending policy was the main driver of the housing bubble. If we kept lending policy fixed as it was in the 1990s, there was almost no bubble, but when lending policy was allowed to loosen, as it did leading up to the crisis, we reproduced the bubble.

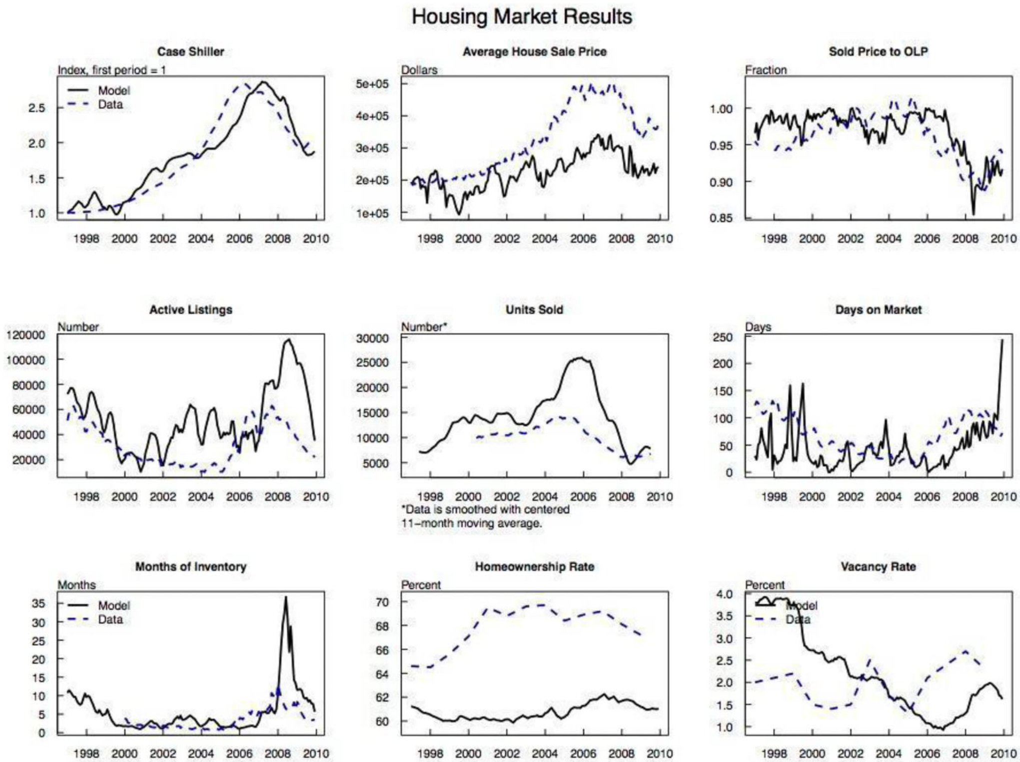

Figure 5. Housing market simulation for Washington DC.
Note: We compare predictions to data for nine different features of the housing market. Solid lines are model outputs, dashed lines are data. The Case–Shiller index is an indicator of how much housing prices have changed; ‘1’ indicates not at all, ‘2’ indicates that prices have doubled, etc.[2] ‘OLP’ is ‘original listing price’. Note the scale of the plots; for example, it looks like homeownership rates are substantially wrong, but the difference is only about 5%.

V. Economic impact of the Covid-19 pandemic

The economic impact of the Covid pandemic combined three exceptional features that made standard models inadequate. First, the shocks were highly industry-specific, making it essential to model the economy with a high degree of resolution; second, the shocks affected both supply and/or demand, depending on the industry; and third, the shocks were huge and sudden.

Our Oxford-based team modelled the impact of Covid on the UK and used this to advise the UK government. There were two main tasks: (i) predict the magnitude of the primary supply and demand shocks on each industry (del Rio Chanona *et al*., 2020) and (ii) predict how these shocks would propagate through the economy, and what the overall effect would be (Pichler *et al*., 2020, 2022).

To calculate the size of the primary supply shocks we made use of the O∗NET database, compiled by the US Bureau of Labor Statistics, which has data on factors such as the distance that workers in a given occupation are typically separated while on the job. This allowed us to determine which occupations were likely to not be allowed to go to work, and other data allowed us to predict who would be able to work remotely. We also used an Italian list of critical industries to determine who would likely be forced to go to work. By mapping occupations into industries, we were then able to estimate the size of the labour shock. To estimate the demand shock, we used a study compiled by the US Congressional Budget Office in 2006, which predicted what would happen to demand for the products of different industries during a flu epidemic.

Our shock propagation model was a dynamical disequilibrium model, which is conceptually very simple. We assumed that an industry cannot produce its product if it lacks either labour, demand, or essential inputs. The essential inputs of each industry were mapped out for us by IHS Markit, a commercial data provider who has since been acquired by S&P, who asked its experts in each industry to say which inputs for each of 52 industries were essential, and which could be omitted for 2 months while still producing the good. For example, the essential inputs for steel include iron, energy, and coke, but do not include restaurants (who manage their cafeterias) or management consultants, even though they both have non-zero entries in the input–output Table. This allowed us to construct a partially-binding Leontief production function that ignored non-essential inputs.

We began by initializing the model with the inventories that were typical for each industry and ran the model dynamically day-by-day. The model was a simple discrete dynamical system. If an industry had no restrictions on labour, demand, or inputs, it would continue to produce its good or service just as it did before the pandemic. However, if there were restrictions on any of these, production was reduced accordingly. We followed this process day-by-day. Although the lockdown was constant over a 2-month period, the state of the system could change over time as inventories were reduced. Because most industries provide inputs for other industries, if an upstream industry's output was reduced, its customer's output could also be reduced on the next day, and so supply shocks propagated downstream. (To clarify terminology, goods flow downstream and money flows upstream.) Similarly, if a retail industry experienced a drop in demand, its suppliers would also experience a reduction in demand, and so demand shocks propagated upstream. The shock waves could also interact, e.g. a supply shock might cause an industry to reduce its output, causing it to need less inputs, creating a demand shock. As inventories were depleted, and as industries reduced their production, supply shortages accumulated, and the primary shocks were amplified accordingly.

We had no prior pandemics that we could use to estimate parameters, so we minimized the number of free parameters and kept the model as simple as possible. Examples of free parameters included the rate at which industries adjust their demand for their inputs as their inventories are depleted, parameters relating to consumption behaviour, and the rate at which industries fire or re-hire in response to changes in production. Whenever possible we took the parameters from the literature, and when this was not possible, we made educated guesses. With hindsight the output was not very sensitive to the free parameters—most of the work was done by the initial conditions, such as the input–output Tables and industry-specific inventories.

As illustrated in Figure 6, our model produced surprisingly accurate predictions. We predicted a GDP shock of 21.5% in the second quarter of 2020; when the dust settled, the correct figure was 22.1%. The industry-by-industry predictions were reasonably good, and the time profile of the predictions of all the major economic indicators over the subsequent year matched what actually happened quite well. A post-mortem showed that while there were a few places where we got lucky, by and large the predictions succeeded for the right reasons.

In a sense this model is not even an agent-based model—there are no agents making decisions. All the work is done by the production function and inventory dynamics. Our postmortem showed that the use of the partially-binding Leontief production function was essential to the accuracy of the model. During the pandemic, industries didn't substitute so much as they just got along without inessential inputs. The model succeeded because it had verisimilitude, i.e. it directly mimicked the processes that were causing economic output to drop. That said, it is not a complete model; for example, there are no prices or markets. (At the time we clearly stated that the model did not predict inflation, which we correctly conjectured would not be a factor during the first year of the pandemic.)

The predictions of the model were made in real time and communicated to the UK government. Based on the analysis of several scenarios, we concluded that the least-bad policy was to keep upstream industries open but close customer-facing downstream industries. By 'least bad', we meant that this scenario caused only a small increase in deaths while doing relatively little damage to the economy. (To analyse deaths we built our own simple epidemiological model.) This was the policy that was enacted by the government. (For a state-of-the-art ABM, coupling an economic model to an epidemiological model, see Pangallo *et al.* (2023).)

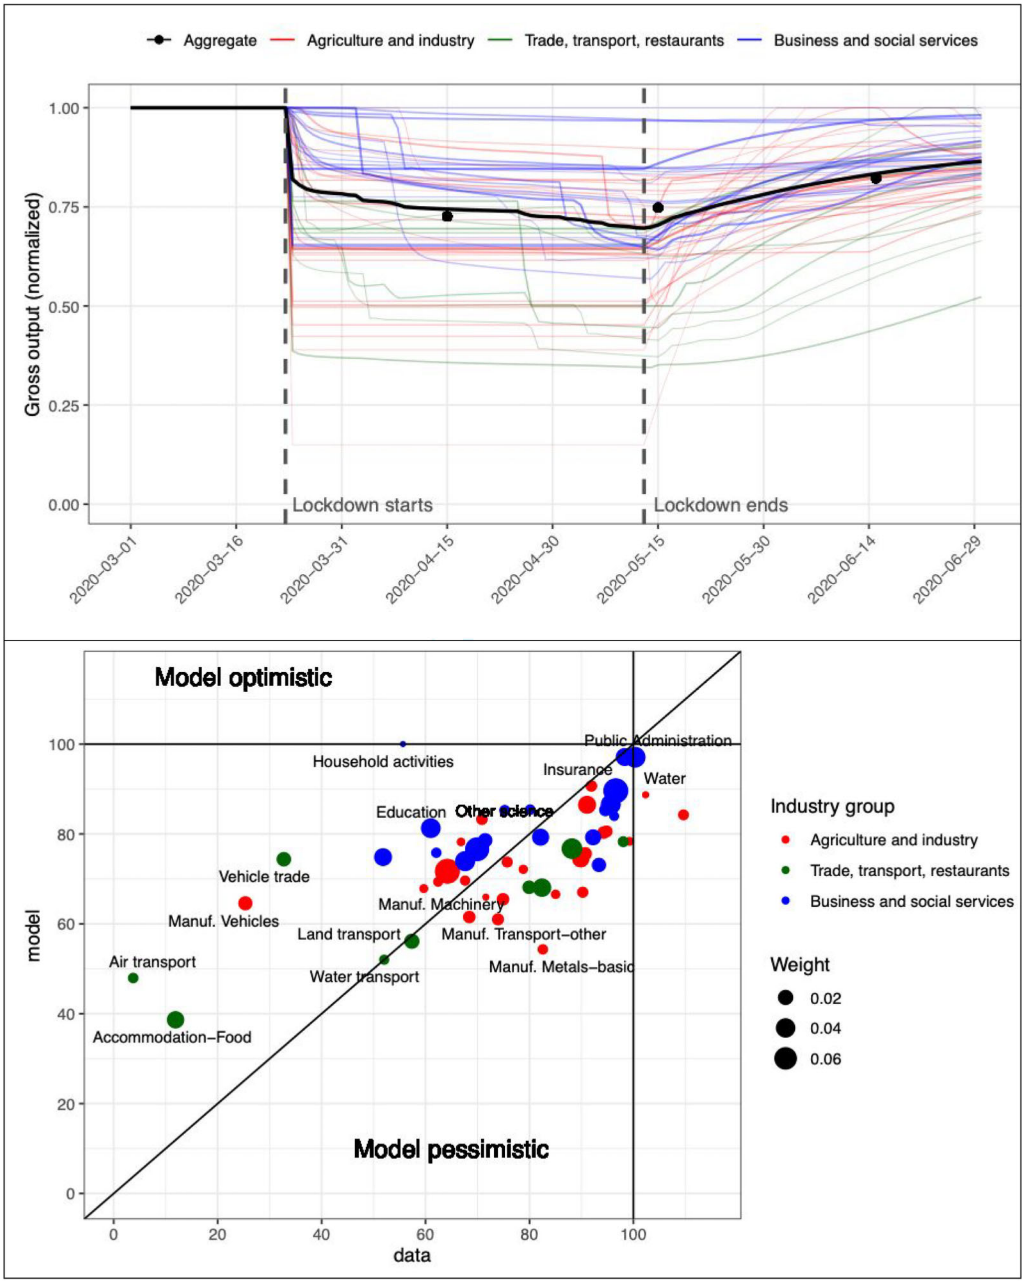

Figure 6. Real time predictions made by the Covid model of Pichler *et al.* on 8 May 2020, when the true values were unknown.

Notes: Upper panel: The solid black line is the model's prediction of aggregate gross output for the UK, while coloured lines are predictions for individual industries. The black dots are measurements of aggregate output (which appeared a year later). Lower panel: Industry-by-industry comparison of values predicted by the model against those measured later in the data. Disk size is proportional to industry gross output. Perfect predictions would lie on the diagonal line.

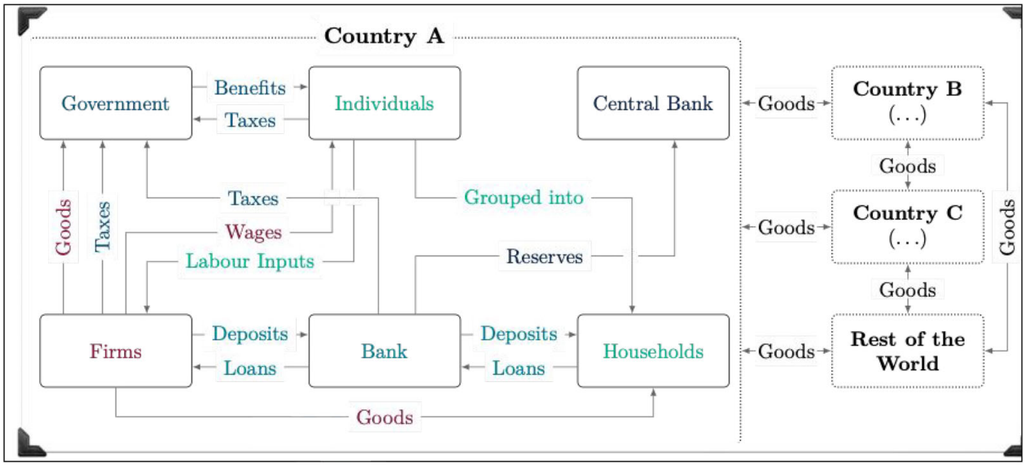

Figure 7. A schematic of the Oxford micro-macro model.
Source: Wiese *et al.* (2024).

VI. Quantitative agent-based micro-macro models

There have recently appeared quantitative micro-macro models that make time series predictions for a given economy at a given point in time (Poledna *et al.*, 2023; Hommes *et al.*, 2024). In the Oxford model (Wiese *et al.*, 2024), for example, initial conditions are constructed with synthetic populations of simulated individuals that closely match demographics, including age, education, gender, and geography (see Figure 7). These individuals, who are organized into households, consume diverse goods and services and either work for a simulated firm or are unemployed. Each household lives in a house that they own or rent. Simulated firms combine labour, capital, and raw materials to produce goods and services, competing with other firms in the same sector. Banks and other financial institutions provide capital to firms, lend to households, and employ labour. Governments formulate fiscal policy, buy goods and services, and employ labour. Central banks lend to commercial banks and set interest rates. The model is (almost) stock-flow consistent.[3]

All the existing models are currently calibrated for only two countries, the country of interest and the rest of the world, but it should soon be possible to run such models for all the major countries of the world simultaneously. I call these 'micro-macro' models because the economy is built from the bottom up, and they can answer questions at the micro level as well as the macro level. Thus, they are capable of answering many questions that conventional macro models cannot even ask.

The remarkable early result from these models is that their predictions are comparable to those of the best DSGE models. I say 'remarkable' because this is even though this work is still in early stages, and there are many known problems with all the existing models—problems that are fixable but require a combination of more careful thought and data gathering to be resolved. We are currently in the process of carefully benchmarking the performance of the Oxford model against state-of-the-art Bayesian Value at Risk models, doing this for many different countries in order to get statistically significant comparisons.

VII. Discussion

(i) Advantages of agent-based models

Agent-based models have several important advantages over standard mainstream models. Perhaps the most important and least controversial is tractability. To quote Xavier Gabaix and David Laibson (2008):

[3] We later realized that because we do not have a long-term bond we are not handling government debt properly, i.e. no one pays for the government's borrowing. This will be corrected soon.

> Tractable models are easy to analyze. Models with maximal tractability can be solved with analytic methods—i.e. paper and pencil calculations. At the other extreme, minimally tractable models cannot be solved even with a computer, since the necessary computations/simulations would take too long. For instance, optimization is typically not computationally feasible when there are dozens of continuous state variables—in such cases, numerical solution times are measured on the scale of years or centuries.

A realistic model of a complicated economic phenomenon like the 2008 financial crisis requires understanding the behaviour of many agents and many institutions, which require at least 'dozens of continuous-state variables'. Furthermore, any realistic model contains nonlinearities, which means that the search landscape is not convex, but rather rugged and difficult to search. This point should be emphasized—in a nonlinear problem with many agents there can be literally millions, billions, or trillions of local maxima that need to be searched. The reason that standard models fail is because they require calculating the *optimal* decisions for each agent. It can be difficult for real people to do this, which is an important clue that such models lack verisimilitude. If computers can't solve the models, how can we expect human beings to solve them? The tractability problems of the standard modelling approach are unavoidable. Utility-maximizing agents can be very useful in simple settings where the necessary calculations are easy to make, but they will inevitably fail in complicated settings. Using models based on utility-optimizing agents, it will never be possible to make successful models of complicated real-world situations that take the many interacting moving pieces of the economy into account.

In contrast, agent-based models don't try to find optima, so they don't suffer from this problem. In our housing model, for example, we considered as many as 100 000 households, all of whom differed from one another. In current state-of-the-art quantitative micro-models, there are millions of agents, requiring many millions of state variables. It would be hopelessly impossible to maximize each household's utility while taking all the others into account. Instead, agent-based models attempt to mimic the way that real households behave. At each step, each agent processes the information available to her and makes decisions using plausible and simple methods. Each agent's decisions consume very little computer time, and as a result it is possible to model millions of heterogeneous agents on a laptop. Finding solutions in such complicated settings will never be possible for a standard model based on utility maximization.

Agent-based simulation can easily incorporate realistic assumptions about the complicated structure of the economy. We can come closer to modelling the complexities of the real world, gaining verisimilitude in two different ways: the decisions of the agents can be more realistic, and the structure of the economy can be more realistic. But the most important consequence is that we can make *tractable* models, ones that produce solutions even when things get complicated. Such models can be built in modular fashion, making it easy to extend them and to compare alternate theories about behaviour. Although many properties of agent-based models are likely to remain controversial, the advantages of agent-based models for tractability should be clear cut.

Endogenous dynamics emerges naturally in many agent-based models. This contrasts with mainstream DSGE models, which under reasonable assumptions, absent shocks, converge to a fixed-point.[4] As a result, mainstream models need to postulate implausible outside shocks to explain events such as the 2008 financial crisis, even when common sense suggests they were endogenously generated by the economy itself. As discussed in some detail in my popular book (Farmer, 2024), and in several papers (Pangallo *et al*. 2019; Asano *et al*., 2021), one should expect endogenous oscillations to occur under circumstances that combine bounded rationality, nonlinear feedback, complex forecasting problems, delayed feedback, and multiple timescales.

To get some intuition as to why this is true, consider a man balancing a pole, as shown in Figure 8. The man's goal is to maintain the pole in a vertical position, but that goal is never perfectly achieved; the pole oscillates away from being vertical in an irregular manner. Why? To make an economics analogy, think of the angle that the pole makes from vertical as the output gap of the economy. A DSGE modeller would explain the movement of the pole in terms of shocks, making the approximation that the man is a perfect pole balancer. Every time a shock strikes the pole,

[4] According to the turnpike theorem, absent shocks, DSGE models will reach a fixed point unless the discount rate is unrealistically high (Scheinkman, 1976).

Figure 8. A man balancing a pole, illustrating instabilities caused by imperfect control.

the man begins to guide it in an optimal manner toward its vertical position, but then he receives another shock, and starts over. However, in a quiet room we know that this explanation must be wrong for the simple reason that there are no such shocks. The right explanation is that a man cannot do this task perfectly—as the pole begins to tip, he moves his hand to compensate, but he tends to over or undershoot, and the pole wobbles, continuously and endogenously. The problem is complicated by the fact that there is a lag between applying the force to move his hand and the response of the pole. In fact, the dynamics are chaotic,[5] as one expects from control systems with lags between the control and the response (Farmer, 1982). Of course, if a man is standing on the deck of a ship in a hurricane, shocks play an important role, but the point is that the oscillations occur even when the ground is stationary and there is no wind. The lesson for macroeconomics is that it suggests that imperfect control due to bounded rational decision-making plays an important role in causing the endogenous oscillations of the economy. This is also suggested by several agent-based models (Pangallo *et al*. 2019, Asano *et al*., 2021), which show how chaotic oscillations emerge in these situations.

Ability to incorporate insights from behavioural economics

Macroeconomics is in a bind. On one hand, the now well-established field of behavioural economics has shown that people are not rational in many important ways. On the other hand, all the workhorse models of macroeconomics, such as those used by treasury departments and central banks, assume rational expectations, albeit with frictions, and sometimes with learning. There is so far no systematic, well-agreed upon method to incorporate insights from behavioural economics into macroeconomic models.

Agent-based modelling potentially solves this. By using computer simulations, it is possible to incorporate the idiosyncrasies of real human behaviour and the role of institutions with as much verisimilitude as needed. This provides a natural route to bring in insights from behavioural economics to economic modelling.

[5] As in the Basel leverage cycle model, the dynamics are locally unstable. They are also globally stable as long as the man does a good job of balancing the pole, so that it doesn't fall over entirely.

That said, I want to emphasize that I am not saying that rational expectations is always a bad way to model the world. Rational expectations is a good model in simple situations where the agents involved can find the appropriate Nash equilibria. When I was 9 years old, my friends and I went through a fad of playing tic-tac-toe (noughts and crosses). But when we discovered the strategy whereby the second player can always get a draw, the game got boring, and we quit playing. For chess in contrast, even with powerful modern computers, no one has yet deduced even a single Nash equilibrium, and there might be billions of them. Rational expectations and Nash equilibria are a great behavioural model for tic-tac-toe, but a useless model for chess. The point is that rational expectations makes lots of sense when the decision problem of the agents is easy to solve and makes no sense when it is extremely hard to solve.

Over the last 20 years or so my collaborators and I have studied the question of when one should expect players with plausible bounded rationality to converge to equilibria in normal form games (for the culmination of this work see Pangallo *et al*. 2019). The upshot is this: whether or not players will converge to equilibria depends on the complexity of the game, here measured by the size of the payoff matrix, which depends on the number of pure strategies and the number of players. It also depends on whether the game is competitive or cooperative. In the way we use the terms here, a game is competitive if one player tends to lose if the other player wins, and it is cooperative if one player tends to win when the other player wins. We studied all possible normal form games using Monte Carlo methods and theoretical techniques from statistical physics and found that when games are both complex and competitive, the strategies of players with plausible bounded rationality do not converge to equilibria. Instead, the system typically has chaotic dynamics in the strategy space, meaning that the strategies of the players never settle into an equilibrium, but rather they continually evolve in an irregular manner. When the games are sufficiently complex the chaotic dynamics of the strategies can be high dimensional, meaning that they are effectively random. This is because players cannot learn the Nash equilibria unless they possess unrealistically heroic levels of rationality.

This work makes two important points. First, it illustrates the earlier point that rational expectations is a poor approximation in complex settings. Second, it shows that one may not be able to make effective models in complex, competitive settings by assuming that the strategies of agents will settle into fixed points. One instead needs agent-based models that can cope with the endogenous dynamics that are likely to occur in settings that are complicated and competitive.

Calibration

Agent-based models are typically not differentiable and are often complicated, making them difficult to calibrate. Recent work offers solutions to this problem using Bayesian methods that take advantage of developments in machine learning (e.g. Dyer *et al*. (2024) provides a recent example). While this work is still under development, there has been some success in calibrating reasonably large models with several free parameters. This is a key enabling factor for constructing quantitative agent-based micro-macro models.

An intrinsic advantage of agent-based models is that they can be *micro-calibrated*, meaning that they can take advantage of data about individual households and firms, to calibrate the model at the level of individual households and firms, rather than relying on aggregate data alone. This can potentially solve the inherent problem caused by severe data limitations to macro-modelling based on aggregate data. There are only 80 years of data since the Second World War, which means that even at quarterly frequency there are only the order of 300 data points. At quarterly frequency, the standard aggregate time series based on quantities like GDP, unemployment, and inflation, are strongly autocorrelated and cross-correlated. As a result, the amount of information they contain is small. As my graduate student Samuel Wiese was able to show, the standard time series used to calibrate aggregate macro-models can be compressed into about 500 bytes of information. (Note that I mean *bytes*, not kilobytes or megabytes.) Conventional macro-models used by central banks, such as the Smets–Wouters model, have at least 40 or 50 free parameters. This is to be expected—the economy is complicated, and one should expect that good quantitative models need to be correspondingly complicated. However, it is unrealistic to imagine that this many parameters can be usefully estimated with so little data, and even with strong Bayesian priors, this is a big problem.

Micro-calibration can potentially solve this, or at least make it better. Agent-based models are best calibrated with micro-data. Using modern computational power, one can potentially build hyper-realistic models that literally model real firms and households at a one-to-one scale. The challenges to doing this are data and confidentiality. This requires detailed data about households and firms, including supply chain information linking the firms together. Such data are only available for some countries (those with VAT taxes recording the counterparties in transactions). It can be challenging to construct synthetic populations that preserve confidentiality while maintaining the needed economic realism. This approach to modelling the economy remains on the horizon, though some central banks are already working to model their economies this way.

(ii) Disadvantages of agent-based modelling

There are also some disadvantages to agent-based modelling. The first problem that most economists will raise is that the decision-making rules are ad hoc. This is necessarily true, though as discussed earlier, there are many ways of grounding assumptions in reality.

In mainstream economics there are very clear and explicit methods for modelling behaviour based on utility maximization, but as discussed above, in agent-based modelling this is much less cut and dried. So far, agent-based modelling has failed to properly and carefully incorporate what has been learned in behavioural economics, which provides a potentially very fruitful avenue for future research.

Another disadvantage of agent-based modelling is that large-scale quantitative models are labour-intensive. Making a realistic agent-based model typically involves considerable programming and data gathering, which is best done by large teams with an internal division of labour. For example, our model for Washington DC housing markets had thirteen authors, and our more recent micro-macro model has nine authors. In contrast to other fields, like particle physics, where papers can have thousands of authors, the anthropological conventions of the academic economics community do not accommodate this style of work.

(iii) The Lucas critique

Agent-based models using fixed rules potentially suffer from the Lucas critique. I used the word 'potentially' to indicate that this is not always true. For example, there will always be investors who buy undervalued assets, and people will always imitate each other and adjust their behaviour based on trial and error. There are also many circumstances where decision-making is mandated by institutions. For example, the model of leverage cycles discussed earlier assumes that investment banks manage their risk as dictated by Basel II. Of course this can change, but the model can change accordingly. Exploring alternatives then becomes an interesting topic of research. For example, in Aymanns *et al.* (2016) we explored alternatives to Basel II and found some that worked better.

In general, however, there are still many circumstances where one needs to consider how behaviour may change in a new policy regime, and where it is important to take into account the fact that people can think. One way to do this is with learning algorithms. To investigate how agents will respond to a policy change, one can train the learning algorithms first with the old policy regime, and then repeat the training with the new policy regime. Then one can assume that the agents switch their behaviours when policies change. In my view, this is a more fruitful way to address the Lucas critique than rational expectations. This suffers from the disadvantage of non-uniqueness—there are many possible learning algorithms, with varying degrees of effectiveness, raising the challenge of discovering which of them best correspond to how real people behave. Finding the methods that best match real human decision-making presents a wide-open field for new research. This method clearly has two major advantages: it is tractable, and it potentially has more verisimilitude. That is, using techniques from machine learning, such as large language models, one can learn decision-making algorithms for the agents in the model that are not designed to be optimal, but rather are designed to mimic the way real people make decisions.

(iv) Vision of the future

Agent-based modelling has been fiercely opposed by the mainstream in economics. There are essentially no agent-based modellers in US economics departments—most agent-based modellers either sit in departments other than economics or belong to unconventional economics departments in Europe. It is currently virtually impossible to publish agent-based modelling in top economics journals. Nonetheless, there are some cracks appearing that suggest that this might be changing. One is that a review paper on agent-based modelling in economics by Rob Axtell and me recently appeared in the *Journal of Economic Literature*. This paper was solicited by Stephen Durlauf, an open-minded mainstream economist. It remains to be seen whether this indicates a new trend or whether it is an isolated event.

Central banks are another indicator of change. As already mentioned, several central banks, including those of England, Italy, Hungary, and Spain, use agent-based models to analyse policy in housing markets. The Bank of Canada now uses an agent-based micro-macro model along with its other more conventional models for advice about monetary policy, and the central banks of Hungary, Italy, Spain, and England are developing such models. There is openness to change, but it takes a model to beat a model. Agent-based models are unlikely to take hold unless they can demonstrate that they make better predictions. While it seems clear that agent-based micro-macro models can match the performance of mainstream aggregate models, whether they can outperform mainstream models remains an open question.

The development of good quantitative agent-based models poses a chicken-egg problem. Creating large-scale quantitative agent-based models is labour and data intensive. Agent-based models that can beat traditional mainstream models are unlikely to emerge without serious effort by large teams requiring large scale resources. But since agent-based modelling sits outside the mainstream, it has so far not been possible to get such resources.

One possibility for a breakthrough may come from the commercial sector, which seems open to using agent-based models. One of the advantages of agent-based models is that there are many situations where their ability to model the world at the micro level allows them to answer questions that traditional models cannot even ask. This opens a potential window to get the funding needed to make such efforts succeed. To meet this need, my colleagues and I have created a company, called Macrocosm Inc., that is building economic models to guide the green energy transition, but this is still in its early stages.

There is a sense in which mainstream macroeconomics and agent-based macroeconomics are converging. As the list of frictions used in mainstream macro gets longer and longer, the number of ad hoc assumptions increases accordingly. In addition, to match data, mainstream macro has found it necessary to abandon the hypothesis that all households make life-cycle calculations based on rational expectations, and instead incorporate ad hoc assumptions, such as 'hand-to-mouth' households. The similarities are obvious. The real economy is complicated, and so it is likely that complicated models are needed to understand it properly (even if we may be able to get some understanding from simple models). Nonetheless, the differences in the theoretical framework are fundamental: agent-based models abandon the assumption of utility maximization that has been a central component of economics for more than a century. Because the frameworks are so different, I don't think these efforts are going to converge, but they may undergo parallel paths of evolution, and they can certainly learn a great deal from each other.

Of course, it is possible that mainstream economists might change their mind and begin hiring agent-based modellers in economics departments and opening the top journals to publishing this kind of work. At present this possibility seems remote, but if agent-based modellers are able to somehow solve the chicken-egg problem and show that their models can outperform mainstream models, this might change quickly.

References

Asano, Y. M., Kolb, J. J., Heitzig, J., and Farmer, J. D. (2021), 'Emergent Inequality and Business Cycles in a Simple Behavioral Macroeconomic Model', *Proceedings of the National Academy of Sciences*, https://doi.org/10.1073/pnas.2025721118

Axtell, R. L., and Farmer, J. D. (2025), 'Agent-based Modeling in Economics and Finance: Past, Present, and Future', *Journal of Economic Literature*, 63, 197–287. forthcoming, https://doi.org/10.1257/jel.20221319

— — Geanakoplos, J., Howitt, P., Carrella, E., Conlee, B., Goldstein, J., Hendrey, M., Kalikman, P., Masad, D., and Palmer, N. (2014), 'An Agent-based Model of the Housing Market Bubble in Metropolitan Washington DC', https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4710928.

Aymanns, C., Caccioli, F., Farmer, J. D., and Tan, V. (2016), 'Taming the Basel Leverage Cycle', *Journal of Financial Stability*, 27, 263–77, https://doi.org/10.1016/j.jfs.2016.02.004

Baptista, R., Farmer, J. D., Hinterschweiger, M., Low, K., Tang, D., and Uluc, A. (2016), 'Macroprudential Policy in an Agent-based Model of the Housing Market', *Bank of England Report No.* 619, http://www.bankofengland.co.uk/research/Pages/workingpapers/2016/swp619.aspx

Carro, A., Hinterschweiger, M., Uluc, A., and Farmer, J. D. (2022), 'Heterogeneous Effects and Spillovers of Macroprudential Policy in an Agent-based Model of the UK Housing Market', *Industrial and Corporate Change*, 32, dtac030, https://doi.org/10.1093/icc/dtac030.

del Rio-Chanona, R. M., Mealy, P., Pichler, A., Lafond, F., and Farmer, J. D. (2020), 'Supply and Demand Shocks in the Covid-19 Pandemic: An Industry and Occupation Perspective', *Oxford Review of Economic Policy*, 36, graa033, https://doi.org/10.1093/oxrep/graa033.

Dyer, J., Cannon, P., Farmer, J. D., and Schmon, S.M. (2024), 'Black-box Bayesian Inference for Agent-based Models', *Journal of Economic Dynamics and Control*, 161, 104827. https://doi.org/10.1016/j.jedc.2024.104827.

Farmer, J. D. (1982), 'Chaotic Attractors of an Infinite-Dimensional Dynamical System', *Physica D*, 4(3), 366–93. https://doi.org/10.1016/0167-2789(82)90042-2.

— (2024), *Making Sense of Chaos: A Better Economics for a Better World*, London, Allen Lane.

Gabaix, X., and Laibson, D. (2008), 'The Seven Properties of Good Models', in A. Caplan, and A. Schotter (eds), *The Foundations of Positive and Normative Economics: A Handbook*, 292–319, Oxford, Oxford University Press.

Geanakoplos, J. (2003), 'Liquidity, Default, and Crashes: Endogenous Contracts in General Equilibrium', *Advances in Economics and Econometrics: Theory and Applications, Eighth World Conference, August 2000*, Vol. II, Econometric Society Monographs, 170–205.

— Axtell, R., Farmer, J. D., Howitt, P., Conlee, B., Goldstein, J., Hendry, M., Palmer, N. M., and Yang, C.-Y. (2012), 'Getting at Systemic Risk via an Agent-based Model of the Housing Market', *American Economic Review*, 102(3), 53–8, https://doi.org/10.1257/aer.102.3.53.

Hommes, C. (2013), *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*, Cambridge, Cambridge University Press, https://doi.org/10.1017/CBO9781139094276.

— He, M., Poledna, S., Siqueira, M., and Zhang, Y. (2024), 'CANVAS: A Canadian Behavioral Agent-based Model for Monetary Policy', *Journal of Economic Dynamics and Control*, 104986.

Pangallo, M., Heinrich, T., and Farmer, J. D. (2019), 'Best Reply Structure and Equilibrium Convergence in Generic Games', *Science Advances*, *5*(2). https://doi.org/10.1126/sciadv.aat1328.

— Aleta, A., del Rio-Chanona, R. M., Pichler, A., Martin-Corral, D., Chinazzi, M., Lafond, F., Ajelli, M., Moro, E., Moreno, Y., Vespignani, A., and Farmer, J. D. (2023), 'The Unequal Effects of the Health–Economy Trade-off During the COVID-19 Pandemic', *Nature Human Behavior*, 8, 264–75, https://doi.org/10.1038/s41562-023-01747-x.

Pichler, A., Pangallo, M., del Rio-Chanona, R. M., Lafond, F., and Farmer, J. D. (2020), 'Production Networks and Epidemic Spreading: How to Restart the UK Economy?', *arXiv*: 2005.10585 [econ.GN], https://doi.org/10.48550/arXiv.2005.10585.

— (2022), 'Forecasting the Propagation of Pandemic Shocks with a Dynamic Input-Output Model', *Journal of Economic Dynamics and Control*, 144, 104527. November, https://doi.org/10.1016/j.jedc.2022.104527

Poledna, S., Miess, M. G., Hommes, C., and Rabitsch, K. (2023), 'Economic Forecasting with an Agent-based Model', *European Economic Review*, 151, 104306. https://doi.org/10.1016/j.euroecorev.2022.104306.

Scheinkman, José Alexandre. 1976, 'On Optimal Steady States of *n* sector Growth Models when Utility is Discounted'. *Journal of Economic Theory,* 12(1): 11–30. https://doi.org/10.1016/0022-0531(76)90026-0.

Wiese, S., Kaszowska-Mojsa, J., Dyer, J., Moran, J., Pangallo, M., Lafond, F., Muellbauer, J., Calinescu, A., and Farmer, J. D. (2024), 'Forecasting Macroeconomic Dynamics using a Calibrated Data-driven Agent-based Model', INET Oxford Working Paper No. 2024-06, https://www.inet.ox.ac.uk/publications/no-2024-06-forecasting-macroeconomic-dynamics-using-a-calibrated-data-driven-agent-based-model