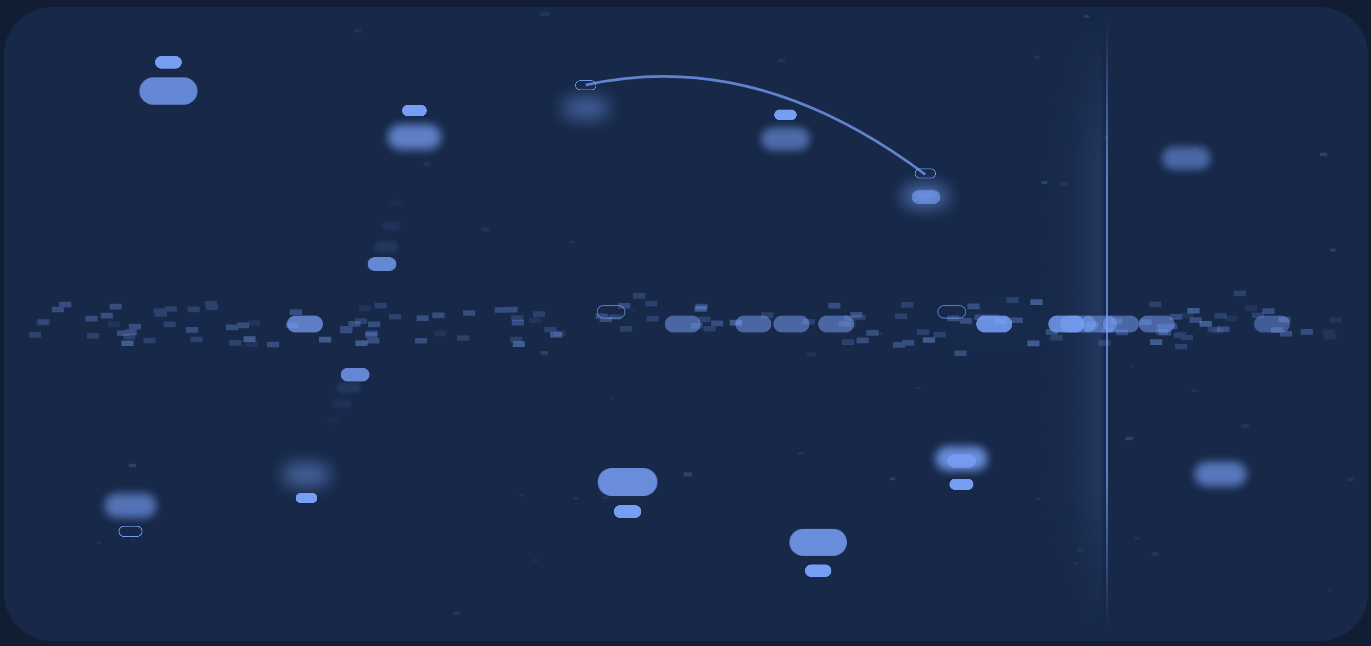
Rillor

EXPLAINER

How agent-based market simulation works

Agent-based market simulation builds a market from simulated traders and its rules, then checks the result against real data or a known answer. LLM agents now join as traders.



Rillor

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

An agent-based market simulation builds a market from simulated traders and the rules they trade under, then checks what emerges against real market data or against an outcome fixed before the run [1, 2, 3].

- An agent-based model is a computer simulation of decision-makers, called agents, and institutions that interact through prescribed rules [1]. In a market model the building blocks are the agents' design, the market mechanism that fixes prices and the environment that links the agents [4].
- Early models showed that market patterns can come from the traders' interaction. In the 1996 Santa Fe artificial stock market, adaptive traders produced bubbles, crashes and market-like price and volume behavior once they explored new forecasting rules faster [5]. Programs that bid at random within a budget, called zero-intelligence traders, reached allocative efficiency, the profit earned as a share of the maximum possible, close to 100 percent in double auctions [6, 7].
- Stylized facts are statistical regularities found in price changes across many markets, such as fat tails and volatility clustering, and reproducing them is the usual test of a market model [2, 4, 8]. In most studies, especially of large models, that test is qualitative and seldom goes beyond showing that a model reproduces a set of them [9].
- A 2019 test of two setups of the ABIDES market simulator found heavy tails in both [2]. Neither reproduced the absence of autocorrelation in returns or matched the skew and spread of volume-volatility correlations in real data, and the authors reported "significant discrepancies" between simulated and real markets [2].
- Calibration fits a model's parameters to known facts [10]. A 2019 comparison on five models found Bayesian estimation consistently ahead of frequentist methods and concluded that calibration techniques "require further development" for large-scale models [9]. A grid of 10 values for each of 5 parameters, at 5 seconds a run, takes about 6 CPU days [11].
- LLM agents enter agent-based market models as one more kind of trader. In one 2025 preprint, an LLM made each buy or sell decision while rules set price and volume, and adding such agents moved a measure of price behavior near all-time highs from about zero mostly into the range seen in real stock data [12].
- An induced-value market gives each agent a private value or cost set before the run, a standard tool of experimental economics, so the efficient allocation is known before the first trade and every outcome can be scored against it [3, 7]. A calibrated replica of a real market is judged instead by how close its statistics come to real data [2, 9].

Background

Why economists build markets from agents

In 2009, J. Doyne Farmer and Duncan Foley argued in Nature that economics needed agent-based models [1]. They defined one as a computerized simulation of a number of decision-makers, the agents, and institutions that interact through prescribed rules [1]. Agents can range from consumers to policy-makers and professional traders, and institutions can include banks and the government [1]. They set this against the equilibrium models then in use, such as those used by the Federal Reserve, which in their words "stripped away most of the structure of a real economy", with no banks or derivatives [1].

The Bank of England's 2016 explainer describes agent-based models as explaining a system's behavior by simulating each individual agent within it [10]. In economics, it says, they have shown how business cycles occur and how statistics observed in financial markets, such as fat tails, arise. By December 2016 the Bank had developed agent-based models of two markets, corporate bonds and housing [10].

A 2025 review in the Journal of Economic Literature by Robert Axtell and Farmer lists the field's accomplishments in financial markets as understanding clustered volatility, market impact, systemic risk and housing markets, and it sets out the hurdles the field still has to overcome [13]. For financial models, Simon Cramer and Torsten Trimborn put the main purpose plainly: to replicate stylized facts and to identify sufficient conditions for them to arise [4].

A simulated market can be judged against real market data, usually through stylized facts, or against an answer fixed before the run, as laboratory economics does with induced values. This explainer is for quant researchers, risk teams and model teams who need to know what such a model contains, how it is checked and whether LLM agents belong in one.

Method

Question. What does an agent-based market model contain, how is it checked against real markets, what limits calibration, and how do LLM agents change the design?

Sources. 33 sources, each opened on 8 November 2026: papers and reviews on agent-based models and stylized facts, a handbook chapter on agent models and human experiments, central bank and research office papers, calibration studies, simulator papers and code, studies with LLM agents and sources on experimental markets.

How. Figures come from each paper's abstract, tables or text, as the paper reports them. For six sources only the abstract or bibliographic record was read [6, 8, 13, 14, 15, 16], and Gode and Sunder's efficiency figures come from Duffy's handbook chapter [7]. LLM studies are described by model size and kind. Numbers computed for this report are labeled.

Scope. Models of financial and experimental markets. Macroeconomic agent-based models appear only as background.

Findings

1. What an agent-based market model contains

Cramer and Trimborn review an abstract agent-based economic market model whose building blocks are agent design, market mechanism and environment [4]. The market mechanism is a rule that fixes a price. The environment is any extra coupling or spatial correlation among agents [4].

Agents. In the order-driven market of Carl Chiarella, Giulia Iori and Josep Perelló, each agent's expected return has fundamentalist, chartist and noise-trader components [17]. The authors find the chartist strategy mainly responsible for the fat tails and clustering in simulated prices, and trace large price changes to large gaps in the book [17]. A 2025 preprint calls this fundamentalist-chartist-noise agent, first introduced by Chiarella and Iori, "a widely adopted agent model" for limit order book simulations [12].

Market mechanism. In the Santa Fe market, agents pass their demand parameters to a specialist who declares the price that clears the market for one risky stock, with a risk-free bond alongside [5]. ABIDES runs a continuous double auction at the same nanosecond time resolution as real markets [18].

Building block	What it sets	Example	Sources
Agent design	How each trader forms expectations and orders	Fundamentalist, chartist and noise-trader components	[4, 17]
Market mechanism	How a price is fixed	A specialist who clears the market, or a continuous double auction	[4, 5, 18]
Environment	Coupling among agents beyond the market	Herding, or a spatial structure among agents	[4]

Table 1. Parts of an agent-based market model, from the sources cited.

2. What early models showed: the Santa Fe market and zero-intelligence traders

The Santa Fe artificial stock market. W. Brian Arthur, John Holland, Blake LeBaron, Richard Palmer and Paul Tayler built a market of heterogeneous agents who continually adapt their expectations [5]. Their working paper, dated 12 December 1996, used 25 agents, each holding 100 predictors conditioned on 12 market descriptors, with dividends following an AR(1) process with rho of 0.95 [5]. Each experiment ran for 250,000 periods and was repeated 25 times under different random seeds [5]. A genetic algorithm revised the predictors every 1,000 periods on average in the slow-learning regime and every 250 in the medium-exploration regime [5].

The outcome depended on that rate. When agents explored new forecasting models slowly, the market settled into the rational-expectations equilibrium. At a higher rate it self-organized into a complex pattern: technical trading emerged, temporary bubbles and crashes occurred, and prices and volume showed the GARCH behavior characteristic of actual market data [5].

In a 1999 Nature paper, Thomas Lux and Michele Marchesi fed a multi-agent model news with neither power-law scaling nor time dependence in volatility, and interactions among agents generated both [15].

Zero-intelligence traders. Dhananjay Gode and Shyam Sunder took the opposite route and removed intelligence. Their 1993 programs submitted random bids and offers [6]. Adding a budget constraint, so that no trader sold below cost or bought above value, was enough to raise the allocative efficiency of their double auctions close to 100 percent [6]. They concluded that a double auction's efficiency derives largely from its structure [6].

John Duffy's handbook chapter reports their figures: average efficiency of 78.3 percent with unconstrained random traders, 98.7 percent with budget-constrained ones and 97.6 percent with human subjects [7]. Allocative efficiency here is the total profit earned divided by the maximum possible profit, the sum of consumer and producer surplus [7].

The result has limits. Duffy reports that in double auctions with large avoidable costs, studied by Van Boening and Wilcox in 1996, zero-intelligence traders did significantly worse than human traders [7]. In a market whose supply and demand were refreshed continuously, studied by Brewer and colleagues in 2002, the prices of budget-constrained random traders showed no tendency toward convergence, while human traders converged [7]. Duffy judges agent models by how well they explain human experimental data [7].

3. Stylized facts as the usual test

The catalogue. Rama Cont's 2001 paper in Quantitative Finance sets out stylized empirical facts drawn from price variations across many markets: distributional properties, tail properties and extreme fluctuations, pathwise regularity, and linear and nonlinear dependence of returns [8]. The paper finds them common to a wide variety of markets and instruments [8]. Cramer and Trimborn count at least thirty documented stylized facts besides fat tails, volatility clustering and the absence of autocorrelation [4]. The Get Real study, which tests simulators against them, lists the following among others [2].

Stylized fact	What it says	Source
Absence of autocorrelation	Linear autocorrelation of returns is insignificant beyond about 20 minutes	[2]
Heavy tails and aggregational normality	Returns have fat tails at short horizons and look closer to normal as the return period grows	[2]

Stylized fact	What it says	Source
Volatility clustering	The autocorrelation of squared returns stays positive over several days	[2]
Order flow facts	Facts about order flow, such as U-shaped intraday volume and the times between orders	[2]

Table 2. Stylized facts as listed in the Get Real study, which also lists long-range dependence, market impact and correlation across assets [2].

How the test is applied. Donovan Platt observed in 2019 that in most studies, especially of large models, validation procedures are qualitative and "seldom venture beyond" showing that a model reproduces a set of stylized facts [9]. Calibration is often manual adjustment of parameters, and a very large number of models recover a similar number of stylized facts [9]. The Bank of England's explainer adds that agent-based models typically produce stylized facts rather than quantitative forecasts, that agreement with them can be assessed in various ways, and that a model with many construction options can spuriously reproduce data that look similar to empirical data [10].

The measures move too. Cramer and Trimborn call excess kurtosis "a very volatile measure" that changes heavily between indices, and find that the time horizon changes several measures substantially [4].

A quantitative test: Get Real. In 2019, Svitlana Vyetrenko and colleagues compared two ABIDES configurations with real data [2]. The first, sparse_zi_100, held 100 zero-intelligence agents [2]. The second, rmsc01, held 1 market maker, 50 zero-intelligence agents, 25 heuristic belief learning agents and 24 momentum agents [2]. The real benchmarks were minute returns of 50 randomly sampled U.S. listed equities on each trading day of 2011, and order book data for one large U.S. bank stock on NASDAQ in June 2019 [2].

Check	sparse_zi_100	rmsc01
Heavy tails and aggregational normality	Shown	Shown, closer to real data
Absence of autocorrelation	Not captured	Not captured
Correlation of volume and volatility	Less skewed and less spread than real data	Less skewed and less spread than real data

Table 3. Results of the Get Real comparison for the checks run on both configurations [2].

For rmsc01 the study also found intraday volume shaped like the reverse of the U seen in real markets, and a poor fit for the times between orders [2]. The authors concluded that "there are still significant discrepancies between simulated markets and real ones", that the time correlation of order flow and intraday volume patterns are not well reproduced, and that herding and clustering may need agents that learn online [2]. They also note that if a stylized fact arises in a market of zero-intelligence agents only, it must come from the mechanism governing the market rather than from strategic behavior [2].

4. Calibration and its limits

Calibration and validation. The Bank of England's explainer separates the two. Calibration is "adjusting the model to fit with the known facts", and validation checks that the model's output is reasonable given what is known [10]. It calls calibration even more difficult in agent-based models, because they typically produce stylized facts rather than quantitative forecasts, and notes that results are hard to communicate because they must be presented alongside the assumptions behind them [10].

Methods. The method of simulated moments is the traditional simulated minimum distance approach, and Platt notes that it needs an arbitrary choice of moments [9]. Jakob Grazzini and Matteo Richiardi note that the criterion function of an agent-based model has no simple analytical expression, and show conditions under which ergodic models can be consistently estimated by simulated minimum distance [14].

Platt compared calibration methods on five models: an AR(1) process, an ARMA(2,2)-ARCH(2) process, a random walk with structural breaks, the Brock and Hommes heterogeneous-beliefs asset pricing model and the INET Oxford agent-based model of the UK housing market [9]. Bayesian estimation consistently outperformed frequentist objective-function methods and gave reasonable estimates in many contexts. Platt concluded that calibration techniques "require further development" before they can definitively calibrate large-scale models [9].

Compute. Francesco Lamperti, Andrea Roventini and Amir Sani worked out what a brute-force search of parameters costs, for a model with 5 parameters at 5 seconds a run on one CPU core [11].

Values per parameter	Runs for 5 parameters	Time on one core	Source
10	100,000	About 6 CPU days	[11]
15	759,375 (computed)	About 1.5 months	[11]
20	3,200,000 (computed)	About 6 months	[11]

Table 4. Cost of a grid search at 5 seconds a run. Run counts for 15 and 20 values are computed for this report from the grid sizes in [11].

A sixth parameter pushes the search past 10 years [11]. Their answer was a machine-learning surrogate, which on average predicted 500 times faster than running the Brock and Hommes model, tested on 10,000 out-of-sample points, and 3,750 times faster than the "Island" growth model [11].

Calibration on a real market. A Bank of England staff working paper calibrated an agent-based model of corporate bond trading against US corporate bond trading data [19]. Where available, inputs came from market data, and others were calibrated by matching statistical features of market returns such as autocorrelations, volatility and fat tails [19]. The model found that market makers' sensitivity to demand and the share of momentum traders strongly influence how far yields overshoot or undershoot after a shock. It also found that growth in passive investment increases the tail risk of large yield dislocations after shocks, even if average volatility falls [19].

Learning the market instead. Andrea Coletta and colleagues note that poorly calibrated simulators can lead to misleading conclusions [20]. Because labeled data on each trader's strategy is not public, they learn a single "world" agent from historical data instead of assuming individual strategies [20].

5. Simulators in use: ABIDES and JAX-LOB

ABIDES. ABIDES, the Agent-Based Interactive Discrete Event Simulation environment, was described by David Byrd, Maria Hybinette and Tucker Hybinette Balch in April 2019 [18]. It supports tens of thousands of trading agents interacting with an exchange agent, configurable pairwise network latencies and messages modeled after NASDAQ's published ITCH and OUCH protocols [18].

A discrete event kernel drives each run, and one seed generates every agent's random number generator, so a run is identical when the same seed is used in the same configuration [18]. The authors name questions that historical data alone cannot answer, such as the market impact of large orders and the value of co-location, because each day can be re-simulated with known, limited changes [18]. The code is open source on GitHub under a BSD-style license [2, 21]. Its main branch was last changed on 19 November 2020 [21]. ABIDES-Gym, from 2021, wraps the simulator in the Gym interface for reinforcement learning, with two benchmark market environments for training daily investor and execution agents [22].

JAX-LOB. JAX-LOB, released in August 2023 by Sascha Frey and colleagues, is a limit order book simulator that runs on GPUs to process thousands of books in parallel [23]. Its authors note that many applications need multiple books, either to calibrate agent-based models or to train reinforcement learning agents [23]. It reported at least a sevenfold speedup in reinforcement learning training against equivalent CPU-based simulators, 550 against 74 steps per second on the same hardware, an NVIDIA A40 GPU with a 32-core AMD EPYC CPU [23].

Its related-work section divides order book simulators into agent-based models and replay of historical messages such as LOBSTER data, and its abstract also names generative models [23]. It states that "the primary limitation of ABMs" is the difficulty of creating sufficiently realistic agent behavior [23].

Generative simulators. MarS, presented at ICLR 2025, generates orders with a "Large Market Model", an order-level generative foundation model built for realistic, interactive and controllable order generation [24].

Stress studies. In an Office of Financial Research working paper, Richard Bookstaber and Mark Paddrik used a limit-order-book framework with liquidity demanders, liquidity suppliers and market makers to study how liquidity recedes in a crisis [25].

6. LLM agents as a new kind of trader

A 2023 survey of agent-based modeling with LLMs names the challenges as environment perception, human alignment, action generation and evaluation, and calls the field "new and quickly evolving" [26].

Hybrid agents. FCLAgent, a 2025 preprint by Ryuji Hashimoto, Takehiro Takayanagi, Masahiro Suzuki and Kiyoshi Izumi, lets an LLM make each buy or sell decision based on the agent's situation, while order price and volume follow standard rule-based methods [12]. The baseline market held 1,000 agents in the style of Chiarella and Iori, and up to 5 of them were replaced by LLM agents [12]. Each LLM agent traded a larger size per order, which kept the number of LLM agents and the computation time down [12]. Runs lasted 500 simulated trading days, with five trials per setting, on an 8-billion-parameter open-weight model and the PAMS simulator [12]. The real benchmark was tick data for 18 liquid Japanese stocks from 5 January 2015 to 20 August 2021 [12].

The paper reports that its simulations reproduce path-dependent patterns that conventional agents fail to capture, and that the reference points behind loss aversion vary with market trajectories [12]. In its Table 2, the coefficient on nearness to the all-time high at 10 days is -0.73 in real data, scaled by 10^-1 [12]. With no LLM agents the simulation gives -0.01, plus or minus 0.76, and with one to five LLM agents -0.67, -0.42, -0.32, -1.20 and -1.31, with trial spreads up to 1.90, against a one-standard-deviation band of -1.27 to -0.20 across real stocks [12].

The paper's other stylized-fact checks require only the sign: positive kurtosis, positive autocorrelation of absolute returns over a range of lags and a positive correlation of absolute returns with volume [12].

Measure	Real data	Simulated	Pass rule
Kurtosis	7.85	5.11 to 8.01	Positive
Autocorrelation of absolute returns at lag 5	0.14	0.02 to 0.04	Positive
Correlation of absolute return and volume	0.46	0.07 to 0.08	Positive

Table 5. Figures from the FCLAgent paper's Table 3, where every simulated row carries a check mark [12].

Every row passes the sign test. The simulated autocorrelation is at most 0.04 against 0.14, under a third of the real value, and the return-volume correlation at most 0.08 against 0.46, under a fifth (computed for this report from [12]).

Markets made only of LLM agents. Alejandro Lopez-Lira's 2025 preprint puts LLMs into a simulated stock market as heterogeneous competing traders, with a persistent order book, partial fills and dividends [27]. The models kept to their assigned strategies as value investors, momentum traders or market makers. The markets showed price discovery, bubbles, underreaction and strategic liquidity provision, and the author warns that prompts can generate correlated behaviors that affect market stability [27].

TwinMarket built a simulated stock market of LLM agents with social interaction and scaled it to 1,000 agents, of which only a subset trade, over 15 June to 15 July 2023 [28]. Its authors report that individual actions can trigger group behavior, such as bubbles and recessions, and checked the market against four stylized facts: fat tails, the leverage effect, the volume-return relationship and volatility clustering [28]. ASFM paired LLM traders with a real order matching system and reports reactions consistent with the real stock market in two controllable scenarios [29].

Leakage and fragility. StockAgent's authors say their design avoids the test set leakage of earlier LLM trading simulations by preventing the model from using prior knowledge about the test data [30]. In a bond market model augmented with an LLM for client decisions, Alicia Vidler and Toby Walsh found that "even slight trade aversion" encoded in the LLM led to a complete stop in trading, and that human-like variability often led to systemic agent collapse across simulations [31].

7. Two designs: calibrated replicas and induced-value markets

Induced values. The Royal Swedish Academy of Sciences' background to the 2002 prize names control of subjects' preferences as a major difficulty, and describes the induced-value method, from Chamberlin in 1948 and developed by Vernon Smith in 1976, as the way to give subjects the right monetary incentives [3]. The method "has now become a standard tool in experimental economics" [3].

In the double auction experiments of Smith's 1962 paper in the Journal of Political Economy [32], each seller had a private reservation price and each buyer a private maximum, and no subject could compute the equilibrium [3]. In one of them, trading prices came close to the theoretical equilibrium price, and their standard deviation fell over five periods as prices converged [3].

Agent models against human markets. Tomaso Poggio, Andrew Lo, Blake LeBaron and Nicholas Chan built a computer simulation of a repeated double auction market designed to match experimental markets with human subjects [16]. Across six designs they replicated several findings of human experimental markets and found "intriguing differences" [16]. In 2026, Pawel Struski and colleagues replicated seminal double auction experiments with LLM agents in place of humans [33]. Their markets showed slower or no convergence toward market equilibrium and less efficient allocations than human markets, with differences across model families and roles [33].

The two designs side by side. Table 6 is a synthesis made for this report from the sources cited.

	Calibrated replica	Induced-value market
Question it answers	Does the model reproduce the statistics of a real market?	What do agents do when the right outcome is known?
Known before the run	The real data to match, while the parameters are estimated	Every value and cost, so the efficient allocation and the equilibrium
How it is judged	Distance between simulated and real statistics, often stylized facts	Efficiency, prices against equilibrium, convergence
Main limit	Calibration techniques need further development for large models, and simulated markets still differ from real ones	The rules alone can produce high efficiency
Sources	[2, 9, 11]	[3, 6, 7]

Table 6. Two designs for a simulated market, a synthesis of the sources cited.

The zero-intelligence results bear on both. If random traders within a budget reach efficiency close to 100 percent in a double auction [6, 7], a high efficiency score says much about the rules and less about the agents. Efficiency therefore needs reading beside measures of behavior such as convergence toward equilibrium and differences from human markets in the same format [16, 33].

Implications

For teams building an agent-based market model. State which stylized facts the model must reproduce, with which measures, horizons and pass marks, and report simulated values next to real ones. A sign test passes a simulated autocorrelation under a third of the real value [12], and kurtosis changes between indices and horizons [4]. Run the market with zero-intelligence agents first, since any fact that appears there comes from the mechanism [2].

For calibration. Size the compute before choosing a parameter grid [11]. Compare estimation methods on the model at hand, since Bayesian estimation did better in Platt's comparison while large models stayed hard to calibrate definitively [9].

For LLM agents in a market model. In FCLAgent the LLM makes one decision, rules set price and volume, and a larger trade size per LLM agent keeps their number and the computation time down [12]. Check that the test period cannot leak into what the model already knows [30], and read each result as one study with its own model, setup and period [12, 27, 28, 31].

For testing an agent before it trades or bids. An induced-value market gives each run a known answer, so efficiency, prices and convergence can be scored against it [3], and results can be compared with human markets in the same format [16, 33].

Where Rillor works. Rillor runs LLM agents in simulated markets and auctions and measures how they bid, and every run can be replayed. Its testbed is built on the designs of experimental economics. Agents receive induced values, so the efficient outcome is known before the first bid.

Limits of this analysis

- For six sources only the abstract or bibliographic record was read [6, 8, 13, 14, 15, 16], and Gode and Sunder's figures come from Duffy's chapter [7].
- The Get Real comparison covers two ABIDES configurations against data from 2011 and 2019 [2]. This report found no comparable published test of the newer simulators among its sources.
- Each LLM study has its own model, setup and period, and FCLAgent has one version [12, 27, 28, 29, 31]. This report found no replication of them among its sources.
- Platt's comparison dates from 2019, and the cost figures of Lamperti and colleagues assume 5 seconds a run on one core [9, 11].
- The Bank of England and Office of Financial Research papers are research in progress by their own description [19, 25]. This report makes no claim about how central banks or regulators use agent-based models today, and gives no counts of simulator users.
- Simulated results are research, and nothing here is investment advice.

Sources

1. The Economy Needs Agent-Based Modelling. J. Doyne Farmer and Duncan Foley, Nature 460, 685-686, 6 August 2009, Institute for New Economic Thinking at the Oxford Martin School, viewed 8 November 2026. https://www.inet.ox.ac.uk/publications/the-economy-needs-agent-based-modelling

2. Get Real: Realism Metrics for Robust Limit Order Book Market Simulations. Svitlana Vyetrenko, David Byrd, Nick Petosa, Mahmoud Mahfouz, Danial Dervovic, Manuela Veloso and Tucker Hybinette Balch, NeurIPS 2019 Workshop on Robust AI in Financial Services, arXiv:1912.04941, 10 December 2019, viewed 8 November 2026. https://arxiv.org/abs/1912.04941

3. Foundations of Behavioral and Experimental Economics: Daniel Kahneman and Vernon Smith. The Royal Swedish Academy of Sciences, Advanced information on the Prize in Economic Sciences 2002, 17 December 2002, viewed 8 November 2026. https://www.nobelprize.org/uploads/2018/06/advanced-economicsciences2002.pdf

4. Stylized Facts and Agent-Based Modeling. Simon Cramer and Torsten Trimborn, arXiv:1912.02684, 2 December 2019, viewed 8 November 2026. https://arxiv.org/abs/1912.02684

5. Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. W. Brian Arthur, John H. Holland, Blake LeBaron, Richard Palmer and Paul Tayler, Santa Fe Institute working paper, 12 December 1996, viewed 8 November 2026. https://sites.santafe.edu/~wbarthur/Papers/Arthur-HollandStockMarket.pdf

6. Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. Dhananjay K. Gode and Shyam Sunder, Journal of Political Economy 101(1), 119-137, February 1993, RePEc/IDEAS record, viewed 8 November 2026. https://ideas.repec.org/a/ucp/jpolec/v101y1993i1p119-37.html

7. Agent-Based Models and Human Subject Experiments. John Duffy, in Leigh Tesfatsion and Kenneth L. Judd (eds.), Handbook of Computational Economics, Volume 2, chapter 19, 949-1011, Elsevier, 2006, author's copy, viewed 8 November 2026. https://sites.socsci.uci.edu/~duffy/papers/duffy2006.pdf

8. Empirical properties of asset returns: stylized facts and statistical issues. Rama Cont, Quantitative Finance 1(2), 223-236, 2001, EconPapers (RePEc) record, viewed 8 November 2026. https://econpapers.repec.org/RePEc:taf:quantf:v:1:y:2001:i:2:p:223-236

9. A Comparison of Economic Agent-Based Model Calibration Methods. Donovan Platt, arXiv:1902.05938, 15 February 2019, viewed 8 November 2026. https://arxiv.org/abs/1902.05938

10. Agent-based models: understanding the economy from the bottom up. Arthur Turrell, Bank of England Quarterly Bulletin 2016 Q4, 16 December 2016, viewed 8 November 2026. https://www.bankofengland.co.uk/quarterly-bulletin/2016/q4/agent-based-models-understanding-the-economy-from-the-bottom-up

11. Agent-Based Model Calibration using Machine Learning Surrogates. Francesco Lamperti, Andrea Roventini and Amir Sani, arXiv:1703.10639, version 2, 6 April 2017 (first version 30 March 2017), viewed 8 November 2026. https://arxiv.org/abs/1703.10639

12. Agent-Based Simulation of a Financial Market with Large Language Models. Ryuji Hashimoto, Takehiro Takayanagi, Masahiro Suzuki and Kiyoshi Izumi, arXiv:2510.12189, 14 October 2025, viewed 8 November 2026.

https://arxiv.org/abs/2510.12189

13. Agent-Based Modeling in Economics and Finance: Past, Present, and Future. Robert L. Axtell and J. Doyne Farmer, Journal of Economic Literature 63(1), 197-287, March 2025, American Economic Association, viewed 8 November 2026. https://www.aeaweb.org/articles?id=10.1257/jel.20221319

14. Estimation of ergodic agent-based models by simulated minimum distance. Jakob Grazzini and Matteo Richiardi, Journal of Economic Dynamics and Control 51, 148-165, 2015, RePEc/IDEAS record, viewed 8 November 2026. https://ideas.repec.org/a/eee/dyncon/v51y2015icp148-165.html

15. Scaling and criticality in a stochastic multi-agent model of a financial market. Thomas Lux and Michele Marchesi, Nature 397, 498-500, February 1999, viewed 8 November 2026. https://www.nature.com/articles/17290

16. Agent-Based Models of Financial Markets: A Comparison with Experimental Markets. Tomaso Poggio, Andrew W. Lo, Blake LeBaron and Nicholas T. Chan, working paper, 2001, listed with abstract on Andrew W. Lo's publications page, MIT, viewed 8 November 2026. https://alo.mit.edu/publications/page/18/

17. The Impact of Heterogeneous Trading Rules on the Limit Order Book and Order Flows. Carl Chiarella, Giulia Iori and Josep Perelló, Journal of Economic Dynamics and Control 33, 525 (2009), arXiv:0711.3581, 22 November 2007, viewed 8 November 2026. https://arxiv.org/abs/0711.3581

18. ABIDES: Towards High-Fidelity Market Simulation for AI Research. David Byrd, Maria Hybinette and Tucker Hybinette Balch, arXiv:1904.12066, 26 April 2019, viewed 8 November 2026. https://arxiv.org/abs/1904.12066

19. An agent-based model of dynamics in corporate bond trading. Karen Braun-Munzinger, Zijun Liu and Arthur Turrell, Bank of England Staff Working Paper No. 592, 15 April 2016, viewed 8 November 2026. https://www.bankofengland.co.uk/working-paper/2016/an-agent-based-model-of-dynamics-in-corporate-bond-trading

20. Learning to simulate realistic limit order book markets from data as a World Agent. Andrea Coletta, Aymeric Moulin, Svitlana Vyetrenko and Tucker Balch, arXiv:2210.09897, 26 September 2022, viewed 8 November 2026. https://arxiv.org/abs/2210.09897

21. ABIDES: Agent-Based Interactive Discrete Event Simulation (source code repository and README). abides-sim, GitHub, created 6 March 2019, last push 6 July 2023, viewed 8 November 2026. https://github.com/abides-sim/abides

22. ABIDES-Gym: Gym Environments for Multi-Agent Discrete Event Simulation and Application to Financial Markets. Selim Amrouni, Aymeric Moulin, Jared Vann, Svitlana Vyetrenko, Tucker Balch and Manuela Veloso, arXiv:2110.14771, 27 October 2021, viewed 8 November 2026. https://arxiv.org/abs/2110.14771

23. JAX-LOB: A GPU-Accelerated limit order book simulator to unlock large scale reinforcement learning for trading. Sascha Frey, Kang Li, Peer Nagy, Silvia Sapora, Chris Lu, Stefan Zohren, Jakob Foerster and Anisoara Calinescu, arXiv:2308.13289, 25 August 2023, viewed 8 November 2026. https://arxiv.org/abs/2308.13289

24. MarS: a Financial Market Simulation Engine Powered by Generative Foundation Model. Junjie Li, Yang Liu, Weiqing Liu, Shikai Fang, Lewen Wang, Chang Xu and Jiang Bian, ICLR 2025, arXiv:2409.07486, version 2, 13 March 2025 (first version 4 September 2024), viewed 8 November 2026. https://arxiv.org/abs/2409.07486

25. An Agent-based Model for Crisis Liquidity Dynamics. Richard Bookstaber and Mark Paddrik, Office of Financial Research Working Paper 15-18, 22 September 2015, viewed 8 November 2026. https://www.financialresearch.gov/working-papers/files/OFRwp-2015-18_Agent-based-Model-for-Crisis-Liquidity-Dynamics.pdf

26. Large Language Models Empowered Agent-based Modeling and Simulation: A Survey and Perspectives. Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu and Yong Li, arXiv:2312.11970, 19 December 2023, viewed 8 November 2026. https://arxiv.org/abs/2312.11970

27. Can Large Language Models Trade? Testing Financial Theories with LLM Agents in Market Simulations. Alejandro Lopez-Lira, arXiv:2504.10789, 15 April 2025, viewed 8 November 2026. https://arxiv.org/abs/2504.10789

28. TwinMarket: A Scalable Behavioral and Social Simulation for Financial Markets. Yuzhe Yang, Yifei Zhang, Minghao Wu, Kaidi Zhang, Yunmiao Zhang, Honghai Yu, Yan Hu and Benyou Wang, arXiv:2502.01506, version 5, 18 October 2025 (first version 3 February 2025), viewed 8 November 2026. https://arxiv.org/abs/2502.01506

29. Simulating Financial Market via Large Language Model based Agents. Shen Gao, Yuntao Wen, Minghang Zhu, Jianing Wei, Yuhan Cheng, Qunzi Zhang and Shuo Shang, arXiv:2406.19966, 28 June 2024, viewed 8 November 2026. https://arxiv.org/abs/2406.19966

30. When AI Meets Finance (StockAgent): Large Language Model-based Stock Trading in Simulated Real-world Environments. Chong Zhang, Xinyi Liu, Zhongmou Zhang and others, arXiv:2407.18957, version 5, 23 June 2026 (first version 15 July 2024), viewed 8 November 2026. https://arxiv.org/abs/2407.18957

31. Shifting Power: Leveraging LLMs to Simulate Human Aversion in ABMs of Bilateral Financial Exchanges, A bond market study. Alicia Vidler and Toby Walsh, arXiv:2503.00320, version 2, 4 March 2025 (first version 1 March 2025), viewed 8 November 2026. https://arxiv.org/abs/2503.00320

32. An Experimental Study of Competitive Market Behavior. Vernon L. Smith, Journal of Political Economy 70(2), 111-137, 1962, Chapman University Digital Commons record, viewed 8 November 2026. https://digitalcommons.chapman.edu/economics_articles/16/

33. Competitive Market Behavior of LLMs. Pawel Struski, Jakub Swistak, Inez Okulska and Przemyslaw Biecek, arXiv:2609.02580, 2 September 2026, viewed 8 November 2026. https://arxiv.org/abs/2609.02580

About this report

Published by Rillor on 12 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. How agent-based market simulation works. 12 November 2026. https://rillor.com/insights/agent-based-market-simulation

Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

NVIDIA, AMD, Ceph, Kubernetes and related product names are trademarks of their owners and are used only to identify products. No affiliation or endorsement is implied.

Superintelligence (SI) outputs can be wrong. Important decisions should include human review.

Work with Rillor

Send us the agent you want tested, the market or auction format and the behavior you need measured. Contact us at rillor.com/contact?type=simulation or contact@rillor.com. More on market simulation: rillor.com/markets/simulation