



A Complete Guide to Agent-Based Modeling For Financial Services

A Bluepaper from Simudyne



1.0

Introduction

Imagine if, as the first signs of stress emerged in the Californian housing market a decade ago, bank executives and regulators had been able to simulate and visualize the far-reaching consequences of mortgage loan defaults in the highly interconnected global market.

Imagine if banks could also simulate different market conditions to demonstrate to regulators that they have tested their algorithms in a broad range of environments and illustrate they have performed necessary due diligence to mitigate against the risk of their algos contributing towards disorderly market conditions.

Finally, imagine if these same banks could monitor the effectiveness of their financial crime detection systems to assess and identify key weaknesses in their processes. If all banks did this, and put into place cost-effective counter-measures, they could help reduce the \$3 trillion global cost of fraud, money laundering and other financial crimes.

In a globalized world, it is no longer enough to consider a bank as though it were isolated from events elsewhere in the financial system. Only by recognizing that the financial market is a complex adaptive system can we begin to tackle areas like risk management, financial crime and disorderly markets. In response to this, banks need simulation models that explicitly capture low-level behavior and the interaction of individual entities within an environment.

Agent-based modeling (ABM) is a technique for modeling complex systems to gain a deeper understanding of system behaviors; they simulate how all kinds of people, regulators, corporations, banks, or investors interact with one other and how that interaction could cause specific things to happen to them and to financial markets more broadly. ABMs are increasingly showing themselves to be the way forward for financial services, but internal inertia and lack of knowledge seems to be a barrier to deployment.

In this guide we explain why banks need to view ABM as a must-have approach for the future—and not just a ‘nice to have’.



2.0

It begins with Enrico Fermi.

The birth of agent-based modeling began with the work of Enrico Fermi in the 1930s as he modeled neutron transport through matter.

This work was instrumental in the development of nuclear power. His ideas led to Monte Carlo simulation, which sought to describe systems of interacting agents and the resulting evolution of the system.

From these hard beginnings, agent-based modeling and computational simulation took off and started transforming decision making in dozens of industries, most recently in financial services.

But, before we dive deeper into that, let's explore a bit more of the beginnings of agent-based modeling and simulation.



3.0

Conway's Game of Life

Agent-based modeling grew out of early work on Cellular Automata, pioneered by John Von Neumann in the 1950s.

Cellular Automata (CA) is a discrete model composed of individual “cell-like” entities that perform a range of functions according to a predetermined set of coded instructions.

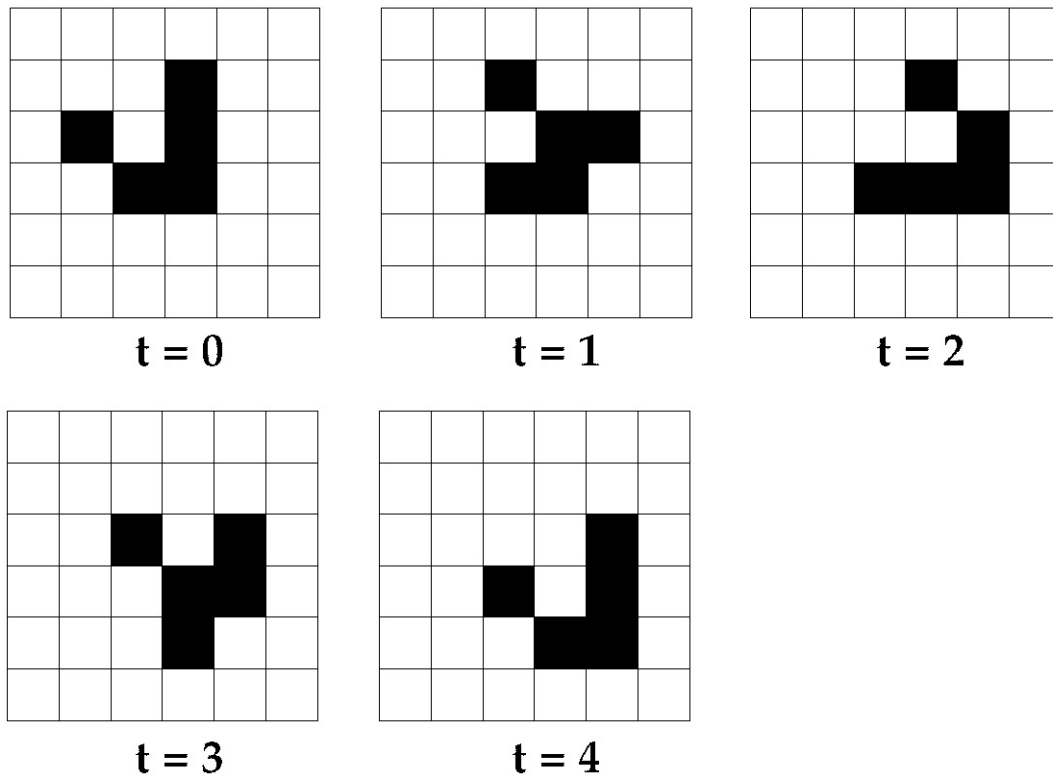
After the pioneering work of Fermi, an important contribution to the field was John Conway's 'Game of Life'. Though not referred to as an agent-based model at the time of its creation, Conway's 'Game of Life' has become the universal go-to for defining and explaining the concept of an ABM.

The game works like so. Start with a grid of squares. Each square is either filled (alive) or empty (dead)—this is the state of each cell. You can start with any pattern you like. Each square has eight directly neighboring cells, and each cell updates its state (alive or dead) based on the state of its eight neighbors. Conway designed four update rules:

1. Any living cell with fewer than two live neighbors dies, as if dying of isolation.
2. Any living cell with two or three live neighbors continues to live.
3. Any living cell with more than three live neighbors dies, as if dying of overpopulation.
4. Any dead cell with exactly three live neighbors becomes a living cell, as if those cells reproduce.

The game proceeds iteratively—like a turn-based game—and at each turn every cell investigates the neighboring cells and updates according to the rules above. Extending the analogy to life, we can think of each turn as a generation of the population and the rules are applied repeatedly at each generation—you can see this in Figure 1.

Figure 1: Conway's Game of Life through time. T = time step.



These simple rules generate complex patterns of movement for the system of cells which seemingly have a life of their own; a life that is far removed from the simplicity of the four rules that govern the cells.

The complex patterns generated by the cells' interactions are an example of "emergent" properties: properties of the whole system which do not belong to any of its parts.

Interestingly, it's impossible to know in advance what a starting configuration of the 'Game of Life' will generate; the patterns produced are not arbitrary or predictable but are in fact *undecidable*. The only way to generate them is to compute all the steps—a computational process called simulation.

The simple ideas introduced by CA—stateful entities which update their states based on local information—underpin the much more advanced approach of agent-based modeling.

4.0

How do ABMs work?



Agent-based modeling and simulation extends the foundational principles of both Fermi's work on neutron transport and Conway's Game of Life.

In agent-based models, the agents are software entities, typically represented computationally as objects. This means that agents themselves hold a state. In Conway's Game of Life, cells have a very simple state—they are either alive or dead. In an agent-based model, the state can be any data that describes the agent; for instance, an agent that is used to represent a person could contain information about their age, height and salary.

In Conway's Game of Life, information from other agents is restricted to just the neighboring eight cells. In an agent-based model, agents can pass information in arbitrarily complex webs of communication. Agents representing people might interact via their social networks which can describe the flow of information between millions of interacting entities. These networks are complex and are formally referred to as graphs. Modeling these graphs is referred to as graph computation.

Components of an ABM

ABMs have three major components:

1. Agents

ABMs contain autonomous models called agents. These agents can be an individual, a group of individuals, or even an organization. Each agent is defined with properties of its own along with relationships with other agents.

2. Topology

The topology of an ABM is the mechanism through which agents can interact with one another in the model. It can be viewed as having two aspects: the rules and the link structure. The topology rules determine the communication forms that the agents can have with one another, whilst the topology structure reflects the links that connect agents together to form a network.

3. Environment

Environmental factors can generally be thought of as shocks that occur to the model. These shocks can come in various forms that affect agent behavior and topology rules. Though the various sets of possible environmental factors can be astronomically large, the number to be tested is normally finite and often predefined by the set of objectives used in the model's initial construction, such as stress tests or policies being evaluated by regulators or risk management groups.

Features

ABMs also have three key features:

1. Heterogeneity: Behavioral differences exist at an agent level.

Agent-based models let you model each individual item of interest; you can model ants, a customer, a household, a small business, a large corporate, a government and so on. You can also inject realistic or irrational behavior into your models. Capturing the heterogeneity of the real world with a 'bottom-up' approach (illustrated in Figure 2) makes the models incredibly interesting.

2. Emergence: Macro insights emerge from agent interactions.

When you then simulate the models, you can see emergent phenomena. This second feature of ABM is profound; birds flocking, locusts swarming and ants building bridges are examples from the natural world. This contrasts with some other modeling approaches in which the actions of many individuals are assumed to lead to a particular outcome, often using a single representative agent.

3. Complexity: Model non-linear behavior of complex adaptive systems.

Agent-based models allow us to recognize that financial markets are interconnected and exhibit non-linear behavior. This means that you're explicitly modeling the bank and the market in which it operates as a complex adaptive system. Rather than ignoring complexity, you're actively embracing it.

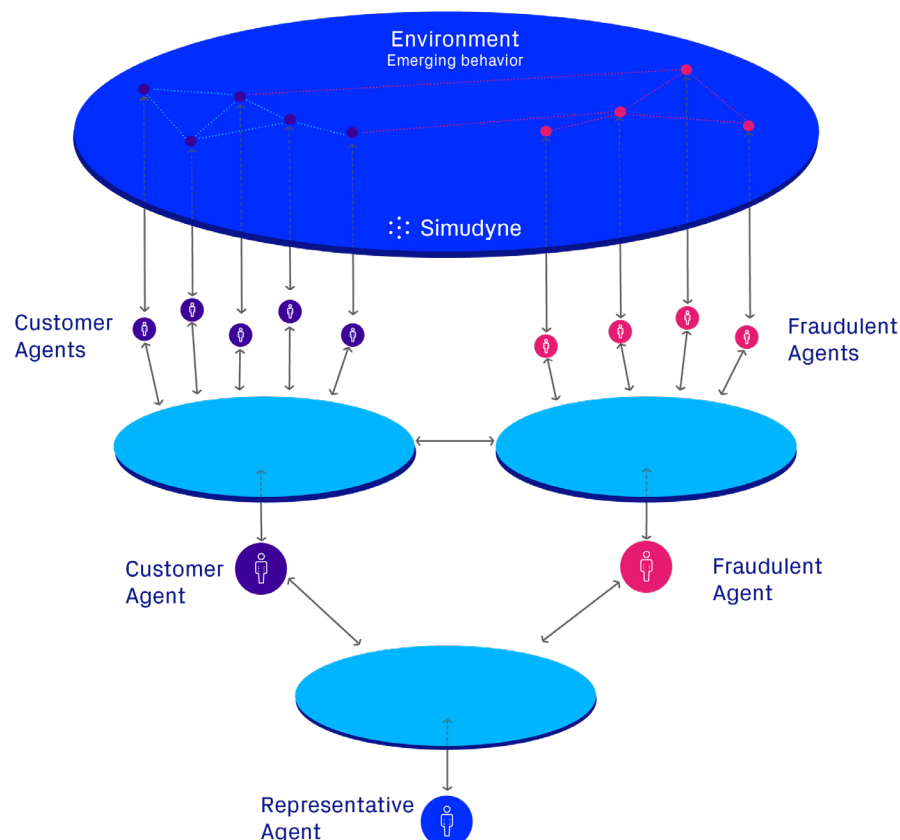


Figure 2: Example of 'bottom-up' approach of ABMs

5.0

Why are ABMs so useful?



Banks are starting to realize that ABMs give a more robust view of the financial system and could enhance existing modeling techniques used across their organization.

Utilizing these techniques often requires a team of quants with machine learning skills and experience in computational simulation—and this isn't an insignificant challenge. Nevertheless, ABMs offer several benefits over traditional modeling techniques which justify the investment.

1. Emergent behavior

The single most powerful feature of ABMs is that the individual actions of the agents combine to produce macroscopic behavior.

The most important example of emergent behavior in economics and financial markets is Adam Smith's metaphor of the invisible hand: how the self-interested actions of real agents in the economy combine to produce socially optimal outcomes.

One of the strengths of ABM is that this invisible hand is made visible and its workings examined. This is in contrast to some other model approaches in which the actions of many individuals are assumed to lead to a particular outcome, often using a single representative agent. This simplification is valid in some cases, but not all combinations of behaviors can be represented by the actions of a single agent.

2. Heterogeneity

As individual agents are modeled, it becomes possible to explore the consequences of the heterogeneity of the agents; that is, agents being different in some way, perhaps by income, preferences, education or productivity.

Incorporating heterogeneity allows for the modeling of much richer behavior. Inequality is a good example: aggregate wealth can increase, but if it is only a small fraction of the population driving this phenomenon, it would suggest very different underlying economic reasons and consequences for a bank than if the entire population were becoming wealthier.

3. Stylized facts

A stylized fact is a term used in economics and finance to refer to empirical findings that are so consistent (for example, across a wide range of instruments, markets and time periods) that they are accepted as truth. Due to their generality, they are often qualitative.

Perhaps the greatest success of ABMs in economics is explaining the stylized facts observed in asset markets. There are a number of phenomena observed empirically in the markets for assets such as bonds or equities which are not explained by traditional economic theory. Some of the two most widely seen across markets are clustered volatility and 'fat tails'. ABMs expose the inner workings of these phenomena to scrutiny by risk managers.

4. Realistic behaviors

The generation of realistic behavior, based on observed behavior, is a strength of agent-based models. Research in behavioral economics shows that people often use *heuristics* when making decisions and that they are not fully rational. There are several models which explore what happens when purely rational options are not available or are too costly, or when agents' environments change over time.

5. Exploring the possibilities

One of the advantages of ABM is that it can very efficiently explore a large number of possibilities. Probabilistic rules applied to each individual agent, in turn, can be a simpler way of exploring scenarios than working out how the entire population of agents should behave together.

6. Complexity, non-linearity, and multiple equilibria

A strength of ABMs is that they describe complex systems. Complex systems are characterized by having many interconnected parts, as well as variables which can change dramatically and which can demonstrate self-organization. Additionally, complex systems can undergo sudden, dramatic transitions, sometimes called phase transitions.

Recent work on agent-based models of the macro-economy has described phase transitions between low and high unemployment. Every economy displays many of the characteristics associated with complex systems.¹

1. <https://www.bankofengland.co.uk/-/media/boe/files/quarterly-bulletin/2016/agent-based-models-understanding-the-economy-from-the-bottom-up.pdf>



6.0

The importance of validation and calibration

ABMs provide an accurate and realistic view of a financial system. The validation and calibration of ABMs is undeniably complex.

A model's ability to reproduce realistic behaviors and the dynamics that generate these behaviors is dependent on quality calibration and validation.

Calibration involves adjusting the model to fit with the known facts; for instance initializing it with empirical data. Validation, on the other hand, is the process of checking that the output of the model is reasonable given what is known, and perhaps cross-checking it with other models or variations in the assumptions.

Why are calibration and validation hard to do?

Calibration is challenging in agent-based modeling because ABMs typically produce stylized facts rather than quantitative forecasts.

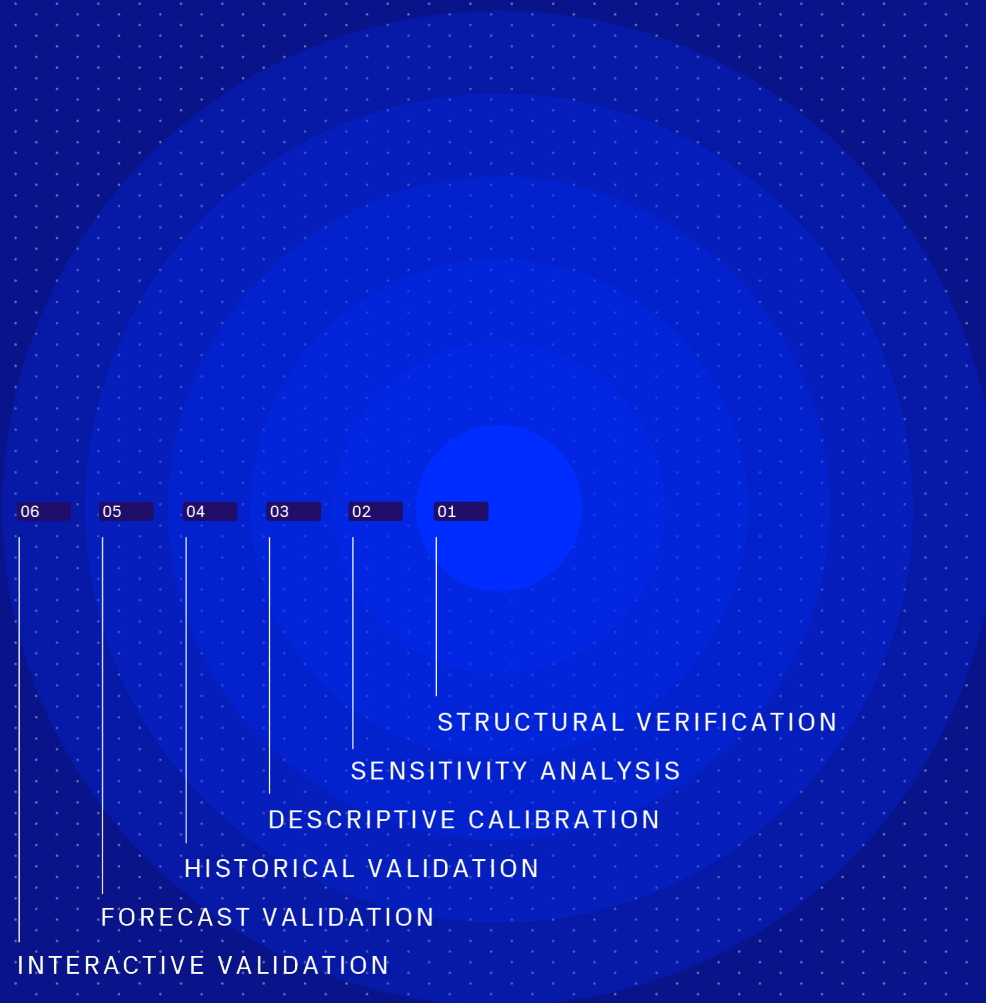
The results of agent-based models can also be difficult to communicate because they must be presented alongside the assumptions used to create them. Although true of all models to some extent, this problem is less acute with models based on historical data alone as they use common statistical techniques.

Additionally, it can be harder in an ABM to understand how changing model inputs affect the model output; this is an unavoidable feature of complex systems but contrasts with more analytically tractable models.

Simudyne's approach

It may be more complex to calibrate an ABM, but that doesn't mean it isn't possible. At Simudyne, our simulators are validated in a 6-step process that ensures the model design and parameters reproduce statistical and behavioral dynamics that match real-world data, which you can see in Figure 3.

Figure 3: Simudyne calibration process



Note: One size doesn't fit all

If a financial institution relies on any one model—ABM or otherwise—they run the risk of failure. An organization should consider using multiple calibrations of multiple models along with a model selection policy as part of their decision-making workflow, so that in a crisis they aren't on the back foot.

8.0

Examples for financial services



Market simulation

Many execution strategies fail to achieve back-test performance benchmarks due to real-world dynamics such as execution fees, market impact, drifting execution latencies, inconsistent liquidity and volatility. Current execution workflows fail to model the true complexities of the entire trade lifecycle.

An ABM greatly expands our understanding of real-world phenomena. By re-creating the low-level interactions that occur between trading entities it is possible to examine systemic effects such as asset price cycles, runs on funding, margin calls and asset fire sales.

In the ABM framework you can represent:

- The different trading strategies of market participants such as high frequency traders (HFT), investment managers or competing execution desks.
- The micro-structure of exchanges including the execution policies, order types, speed bumps or auction rules.
- The latency and locations of different exchanges or execution venues.
- The banks own internal infrastructure including smart-order routers, circuit breakers and trade processing.

Explore extreme market events

The recreation of the low-level interactions between market participants means that these simulations are stochastic in nature and running them multiple times will generate a range of outcomes. By running the same simulation thousands of times, you produce tail events, stressed conditions and market crashes.

Specific scenarios can also be run to produce large amounts of data. Therefore, for the first time, algos can be optimised to run in environments that rarely occur, such as in mini-flash crashes.

Accurately model market impact

Uniquely, ABM allows banks to run the same simulation thousands of times with or without an order.

In comparing the results of the two categories of simulation, it is possible to more accurately quantify the likely impact of a trade.

This same process can be run hundreds of times with different execution strategies to determine which approach is the optimum for every individual execution order.

Create scenarios with no historical precedent

ABMs are calibrated and optimized against historical data to ensure that they produce data that is indistinguishable from the real world.

Once the simulation faithfully creates market data, the underlying parameters are tweaked to produce novel data.

Run many times faster than replay servers

Traditional replay servers are restricted to a maximum throughput that is determined by processing power.

The nature of ABM is that back-tests do not need to run consecutively and can be distributed across machines to run concurrently, speeding the process up by many multiples.

Practical Applications

Algo back-testing

Develop execution algorithms that can be deployed in a wider range of environments, scenarios and customer order types using agent-based modeling and simulation.

Market impact

Agent-based modeling and simulation is a unique way of determining the market impact of any trade by running identical simulations with or without a given trade.

Market slippage

Modeling latency, market volatility or the market impact of news enables the modeling of market slippage.

Market manipulation prevention

By including malicious activities in the model, such as order book spoofing agents and other prohibited trading strategies and simulating their activities, it is possible to develop more effective market manipulation mitigation policies.

Download the full brochure here.



8.0

FinCrime analytics

Banks have spent billions on transaction monitoring systems that scrub their accounts for possible money laundering schemes. This method, however, significantly increases the number of false positives.

With ABM, organizations can simulate millions of potential scenarios to create synthetic data that is essential to identifying the level of fraudulent activity that isn't being picked up by current systems. This type of data set can be used in a sandbox environment for many diverse purposes such as to train detection agents, test current methods and evaluate transaction monitoring systems.

In the ABM framework you can:

- Replicate networks of customers, merchants and banks.
- Run thousands of simulations to generate huge amounts of synthetic data with which you can supercharge your machine learning process.
- Test for fraudulent behavior that has yet to occur.
- Model how fraudulent agents adapt to changing detection rules and thresholds.
- Optimize your thresholds to maximize fraud identification while minimizing false positives.

Generate realistic synthetic datasets

What's so powerful about agent-based modeling and simulation is that it uses real data to create realistic synthetic datasets.

Synthetic data contains no personal information or disclosure of legal or private customer transactions, so it is completely compliant with privacy regulations like GDPR. It has the added benefit of being easier to acquire, faster and at less cost for experimentation, even for those that have access to their own data.

Explore diverse fraud scenarios

For the first time algos can be optimized to run in environments that rarely occur, such as in rare fraud cases or events such as Brexit.

Generating synthetic data would also be helpful where there is a significant change to either the customer's internal policies or regulatory rules.

Create fraud training datasets

Unlike statistic-driven models that are tied to past data, ABMs are capable of using historical data to produce new, unseen and less imbalanced fraud training datasets that reflect the possible behaviour of future events.

Fraudsters adapt their behaviors to side-step heuristics or statistic-driven models and therefore original historical data is less helpful for training new fraud classifiers.

Establish the correct approach

By modeling adaptive adversaries before committing valuable resources, banks and other financial institutions can use precisely the right approach to minimize financial crime, as well as protecting clients from arbitrary, costly and ineffective FinCrime controls.

Looking ahead

Future-proofing your system

Improving FinCrime detection is an ever-evolving process.

As criminals get ever more sophisticated at targeting permeable controls, simulation holds the key to winning the battle against financial crime.

Simulations empower organizations to:

- Continue the development of the simulation model and extend the coverage of fraud and other financial crimes.
- Incorporate simulation as part of their process for improving fraud detection methods.
- Carry out exploratory analysis of foreseen or upcoming scenarios and evaluate current or newly developed fraud controls.
- Continuously calibrate the simulator based on business data to measure the quality of FinCrime controls.
- Prove not only compliance with the law but also engagement in a proactive FinCrime analytics programme inside the bank.

[Download the full paper here.](#)



9.0

Risk management

Risk simulations are complex; traditional ‘top-down’ models often don’t capture the subtleties of the many entities and their interactions that are involved in the risk system. They can correlate variables in historical data but they can’t always identify causal relationships in it, making it difficult to extrapolate and predict a potential future outcome.

Using ABM, banks can enhance the realism of their simulations to account for feedback loops, unusual relationships between agents and complex scenarios that include external factors such as climate impact or economic shocks.

In the ABM framework, banks can:

- Explicitly account for heterogeneity in financial market participants and embrace a view of financial markets as complex adaptive systems.
- Capture non-linear dynamics, knife-edges of risk and tail risk.
- Incorporate the dynamic responses of market participants and recognize that the future will not look like the past.
- Model the impact of behavioral changes of market participants.
- Identify, quantify and manage contagion risk.
- Simulate future events.

Harness Artificial Intelligence (AI)

Artificial intelligence models can be trained using computer-generated data from an ABM simulation. These simulations produce realistic data sets—normally hard to come by in complex systems—which can be used to improve risk management.



Capture causative relationships

Using agent based modeling to simulate network analysis allows us to capture causative relationships — e.g., funding providers pulling funds from banks, CDS contracts linking financial institutions to failing insurers, or suppliers losing out on contracts from failed construction companies — that have the potential to do significant harm.

Simulate IRRBB

Using ABM, banks can model each financial instrument as an agent, and aggregate their cash flows from the bottom-up. Changes in rates will have heterogeneous impacts on different instruments, as determined by the interest rate index, part of the curve they are priced from, or embedded optionality — both contractual and behavioral.

Comply with regulatory metrics

Many banks are still getting to grips with IFRS9—a recent accounting standard that changes the way banks provision against losses.

Agent-based models are ideal for producing IFRS9-consistent outputs. The agent-based set-up is rich enough to capture key concepts embedded in IFRS9.

Use cases in risk

Run multiple what-if scenarios to proactively revise your strategy, mitigate potential risk and optimize bad-debt provisioning in order to make better informed, forward-looking decisions.

Mortgage analytics

Recreate complex dynamics by using a bottom-up approach to model markets and customer behaviors.

Stress testing

Assess whether banks have enough capital to withstand extreme shocks (such as rapid jumps in unemployment rates) and whether they're able to continue to support the economy in such severe conditions.

Contagion risk

Use simulation to link agents together through causal relationships to capture network risks and better understand contagion risk.

Credit risk

Run large numbers of scenarios, model systematic risk and run reverse stress tests to inform lending criteria—all in a massively scalable framework.

Download the full paper here.

10.0



Further reading

[Uncovering Hidden Financial Crime Through Advanced Simulation](#)

Banks have invested, and continue to invest, billions of dollars to stop financial crime. Yet, with the global cost of fraud rising every year, it is clear that we are failing to win the battle. In this blueprint, we look at the inherent problem with developing better controls in the financial crime domain and explore how institutions can make significant improvements using simulation.

[Delivering Algo Performance through Enhanced Market Execution](#)

Sell-side institutions today are facing a highly competitive marketplace that is being constantly reshaped by client demands, new technologies and changing regulatory requirements. This blueprint—based on a survey of 15 sell-side executives involved in algo trading—evaluates current challenges around trade execution and algo testing in simulated market environments and explores their current operating procedures.

[Computational Simulation: The Next Frontier for Better Decision-Making](#)

This paper addresses the challenges faced by the financial services industry and shows how new analytical tools can be used to gain greater insight and produce better decision-making capabilities. We show how platforms, such as those developed by Simudyne, can be combined with the power of the Cloudera modern data platform to provide policymakers, regulators, and other key players in the financial services sector with the necessary tool kit to understand today's financial sector.

[Agent-Based Models: Understanding the Economy from the Bottom-up](#)

This article by the Bank of England considers the strengths of agent-based modeling and the ways that it can be used to help central banks understand the economy. These models provide a complement to more traditional economic modeling which was criticized following the Great Recession.

[Synthetic Reality: Synthetic Market Data Generation at Scale](#)

Agent Based Modeling (ABM) offers advances over traditional model testing applications in the financial markets. This paper, created in partnership with Refinitiv, introduces the topic of agent-based modeling and illustrates its application in a number of scenarios.



About Simudyne

Simudyne is the market leader in agent-based modeling supporting financial services. A partnership with Simudyne ensures your organization has the industry's most advanced simulation technology with support from the leading experts in the field.

Our software provides a robust library of code and examples for frequently used and specialized functions that saves time and reduces the complexity of agent-based modeling. What normally requires several months of engineering and thousands of lines of code can now be delivered at a fraction of the time and cost.

Simudyne uses the firm's existing infrastructure to ensure it is cost effective and easy to deploy. As a deployed solution, the technology is safe and secure. It sits behind the customer's firewall and all the bank's data and models remain proprietary.



Contact

A: 35 Luke Street, London EC2A 4LH

E: info@simudyne.com

W: simudyne.com

TW: twitter.com/simudyne

LI: linkedin.com/company/simudyne

Simudyne is a rapidly growing technology business, harnessing the power of advanced simulation, to help organizations make radically better decisions. Our efficient and scalable simulation platform allows enterprises to create a virtual environment where they can test drive their decisions, fail fast without consequences and create solutions that drive growth.