Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Emergence of Wyckoff Patterns in Financial Markets through Agent-Based Modelling with Asymmetric Information

Francisco Vicente Lasa

lasavf@gmail.com

Universidad Nacional de Educación a Distancia

Research Article

Keywords: agent-based modelling, Wyckoff methodology, Bayesian games, stylized facts, market microstructure, asymmetric information

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11706698/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Emergence of Wyckoff Patterns in Financial Markets through Agent-Based Modelling with Asymmetric Information

Abstract

This paper reproduces a fictitious financial market as a Bayesian game with asymmetric information, formalizing through an agent-based model the hypothesis that the accumulation and distribution structures described by the Wyckoff methodology (price action) can emerge from a reduced set of individual behavioural rules between an institutional agent and a heterogeneous population of retail agents. The model is implemented under two triggering logics of the informed agent, one based on the distance to a reference price and another based on an aggregate flow signal that aims to capture moments of euphoria and capitulation among retail players. Validation relies both on quantitative criteria as well as on the stylized facts that emerge from the simulations themselves. The coincidence between the simulated patterns and the patterns described in the technical literature constitutes sufficient generative evidence, without assuming that this is the definitive causal explanation of price formation in markets.

Keywords: agent-based modelling; Wyckoff methodology; Bayesian games; stylized facts; market microstructure; asymmetric information

JEL Classification: C63; C73; D82; G14; G41

1. Introduction

This paper proposes the modelling of a Bayesian game (Harsanyi, 1967) for the understanding of currency markets, where agent-based simulation acts as a computational resolution tool to explore a system whose complexity prevents a closed analytical solution, and where the model could, in theory, allow behavioural patterns and equilibria to emerge that have not been designed in a deterministic way a priori.

This orientation implies questioning one of the most entrenched assumptions of standard economic theory, namely methodological individualism in its most optimistic version, according to which the rational pursuit of individual interest inevitably leads to collectively efficient outcomes. Game Theory, from the Prisoner's Dilemma to models of asymmetric information, has systematically demonstrated that this is not the case. Individual rationality can generate collectively suboptimal equilibria, and financial markets are perhaps one of the most recurrent examples of this phenomenon. Under this lens, the market is a strategic ecosystem where the heterogeneity of participants and the asymmetry of information are the fundamental drivers of price action.

The persistence of recurring patterns in price series, empirically identified by figures such as Richard Wyckoff and Charles Dow more than a century ago, suggests the existence of an underlying structure that the Efficient Market Hypothesis (Fama, 1970) fails to explain satisfactorily. From the classical models of Cournot (1838) to Bayesian games of incomplete information, Game Theory has attempted to formalize how agents make decisions in the face of uncertainty about others' behaviour. However, in a currency market, the combinatorial explosion of players and indefinite time horizons makes the calculation of a static Nash Equilibrium a practical impossibility. Order, expressed in the form of trends or ranges on charts, is not imposed by a central agent, but rather emerges from the decentralized interaction of participants with asymmetric objectives and information.

The central question we ask in this research is whether strategic interaction among agents with simple heuristic rules and information asymmetry is a sufficient condition for structural market patterns to emerge. The starting hypothesis is that the accumulation, trend, and distribution phases described by the Wyckoff methodology are logical consequences of a game in which the informed player (Smart Money) optimizes its payoff function by managing inventory against a mass of agents with limited information. If this hypothesis is correct, the Wyckoff methodology would not describe graphical patterns but rather transient equilibria of a dynamic Bayesian game, endowing it with a theoretical foundation based on strategic causality.

The model will be considered successful if it generates coherent phase transitions without these having been explicitly encoded, and it will fail if it produces only white noise or if the patterns only emerge through a forced parametrization that nullifies the agents' autonomy. Under this approach, we do not seek price prediction, but rather the characterization of the players' strategies underlying these emergent patterns.

The paper is structured in six sections. The theoretical framework connects Wyckoff phenomenology with Game Theory and complex adaptive systems, arguing why this type of game is not analytically solvable. The formalization of the game expresses that argument in a scheme of players, strategies, and payoffs. The construction of the agent-

based model (ABM)[1] details the computer implementation and the design decisions. The simulations present the results and the quantitative/qualitative analysis, while the discussion and conclusions reflect on the scope of the model, its limitations, and future lines of research.

[1] An interactive version of the simulator is available at https://simulador-patrones-abm.vercel.app/

2. Theoretical Framework

2.1 Financial markets as strategic systems

In a financial market, the outcome of each agent's decisions depends on the decisions of the others, which in turn depend on their own. This circularity, characteristic of interactive decision theory (Game Theory), defines the structure of an *over-the-counter* (OTC) market operating contracts for difference (CFD) as a game in the formal sense of the term, where each participant forms expectations about the behaviour of the others and acts accordingly.

The neoclassical paradigm dominant during the second half of the 20th century, built upon the Efficient Market Hypothesis and the representative agent assumption, assumes that price instantly reflects all available information and that market participants can be treated as a single rational maximizing agent. Within this framework, the price formation process is a black box in which this fictitious walrasian auctioneer[2] finds the equilibrium without friction, and the market would simply be the mechanism that equates supply and demand.

While this simplification can produce clean predictions, the cost is ignoring the heterogeneity that makes a real market interesting, such as the diversity of participants, information asymmetry, transaction costs, unpredictability, chaos, and the strategic nature of every buy or sell decision.

As Kirman (1993) points out in a critique that has gained strength in recent decades, the representative agent is a fundamental distortion because the aggregate behaviour of a population of heterogeneous agents cannot, in general, be represented as the behaviour of a single optimizing agent, however sophisticated it may be. This critique has been formulated in various fields of economics, from consumer theory to macroeconomics, with particular force in financial economics.

The alternative that has been gaining ground since the nineties is that heterogeneity is not noise, but rather an inherent characteristic of markets and their functioning. Participants differ in their information, their time horizons, and their decision rules, and it is these differences that generate order flow, liquidity, and price action. O'Hara (1995) formalizes that trading is not a secondary process but rather the mechanism by which private information dispersed among participants becomes incorporated into the public price. Price would not be an exogenous datum that agents observe, but rather the endogenous result of their strategic interactions.

Contracts for difference are the instrument on which the modern retail FX market has operated since its expansion in the 2000s. They do not involve the physical exchange of assets, but rather the agreement to settle the difference between the opening and closing price of a position, which makes the CFD, by construction, a zero-sum game between

---

[2] Concept introduced by Léon Walras in Éléments d'économie politique pure (1874). An agent external to the market announces prices, collects the quantities demanded and supplied, and adjusts them until finding the vector that simultaneously clears all markets. No transaction is executed until general equilibrium is reached, which implies a total absence of frictions, information costs, and strategic behaviour among participants.

buyer and seller, what one wins is exactly what the other loses (not accounting for the provider's spreads or swaps).

However, this accounting symmetry conceals a deeper information asymmetry, in that the two parties to the contract do not have the same access to relevant information, nor the same resources to process it, nor the same time horizons. The accounting zero sum can thus become a systematic transfer of capital from the less informed to the more informed, a matter that Bayesian Game Theory describes and that institutional trading practice exploits.

This is not a conspiratorial idea, but rather a quantified fact recognized by the European regulator. In 2018, the European Securities and Markets Authority (ESMA) intervened in the retail CFD and forex market through investor protection measures that included, among others, the obligation for each broker to publish the percentage of its clients who lose money (ESMA, 2018)[3].

The regulatory motivation was precisely the accumulated evidence from the competent national authorities of several member states, which consistently showed that between approximately 75% and 90% of retail accounts lose money trading CFDs, with average losses per client ranging between 1,600 and 29,000 euros (ESMA, 2018).

In Spain, the CNMV documented that between January 2015 and September 2016, 82% of clients who traded CFDs closed at a loss, with an aggregate loss of 142 million euros. These figures would be the statistical footprint of a structural asymmetry that the regulator explicitly recognized as a conflict of interest inherent to the product, and although the public disclosure obligation does not change the game, it constitutes a kind of subtle confirmation of the hypothesis this paper aims to formalize.

In the retail CFD and forex market, at least two distinct execution models coexist, known in the industry as A-book and B-book. In the A-book model, the broker acts as a pure intermediary: it passes the client's orders on to external liquidity providers and earns through the spread or commission, without taking its own position. In the B-book model, the broker internalizes the orders and becomes the client's direct counterparty: when the trader buys, the broker sells; when the trader loses, the broker wins. In this case, the orders do not even reach the market but are processed on the intermediary's own internal servers, which acts practically as a croupier.

In that context, the game the retail trader faces is not the one they believe they are playing, it is not smart money against retail in a decentralized market, but rather the client against their own broker, who may know their exact position and stop loss. The B-book model is not a fraudulent or exceptional practice, it is a legal business model in which the retail trader does not operate in a spontaneous market, but within one designed by their counterparty.

Order flow, understood as the net flow of buy and sell orders that at each instant pushes the price in one direction, is the minimum unit of analysis of this system, and it contains signals about the private beliefs of participants that the public price has not yet incorporated (Evans and Lyons, 2002). In a market with a visible order book, zones of accumulated liquidity can be seen, and the institutional operator (whom we will call SM,

[3] See European Securities and Markets Authority (ESMA) Decision (EU) 2018/796 of 22 May 2018 on product intervention measures in relation to CFDs. The studies cited by ESMA gather data from national authorities such as the UK's FCA, France's AMF, and Spain's CNMV, confirming that the vast majority of retail investors suffer systematic losses in these products.

for Smart Money) knows that retail stops are located there. When the price reaches that zone, the stops are triggered, generate a sweep, the SM absorbs that flow, money changes hands, and the price reverses. This is exactly a spring in Wyckoff terms, described in order flow terminology.

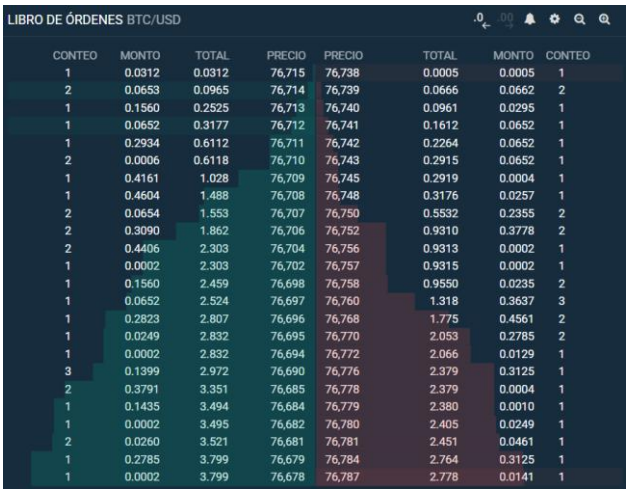

LIBRO DE ÓRDENES BTC/USD

CONTEO	MONTO	TOTAL	PRECIO	PRECIO	TOTAL	MONTO	CONTEO
1	0.0312	0.0312	76,715	76,738	0.0005	0.0005	1
2	0.0653	0.0965	76,714	76,739	0.0666	0.0662	2
1	0.1560	0.2525	76,713	76,740	0.0961	0.0295	1
1	0.0652	0.3177	76,712	76,741	0.1612	0.0652	1
1	0.2934	0.6112	76,711	76,742	0.2264	0.0652	1
2	0.0006	0.6118	76,710	76,743	0.2915	0.0652	1
1	0.4161	1.028	76,709	76,745	0.2919	0.0004	1
1	0.4604	1.488	76,708	76,748	0.3176	0.0257	1
2	0.0654	1.553	76,707	76,750	0.5532	0.2355	2
2	0.3090	1.862	76,706	76,752	0.9310	0.3778	2
2	0.4406	2.303	76,704	76,756	0.9313	0.0002	1
1	0.0002	2.303	76,702	76,757	0.9315	0.0002	1
1	0.1560	2.459	76,698	76,758	0.9550	0.0235	2
1	0.0652	2.524	76,697	76,760	1.318	0.3637	3
1	0.2823	2.807	76,696	76,768	1.775	0.4561	2
1	0.0249	2.832	76,695	76,770	2.053	0.2785	2
1	0.0002	2.832	76,694	76,772	2.066	0.0129	1
3	0.1399	2.972	76,690	76,776	2.379	0.3125	1
2	0.3791	3.351	76,685	76,778	2.379	0.0004	1
1	0.1435	3.494	76,684	76,779	2.380	0.0010	1
1	0.0002	3.495	76,682	76,780	2.405	0.0249	1
2	0.0260	3.521	76,681	76,781	2.451	0.0461	1
1	0.2785	3.799	76,679	76,784	2.764	0.3125	1
1	0.0002	3.799	76,678	76,787	2.778	0.0141	1

Fig.1 Example of an order book accessible to an operator on the Bitfinex crypto asset platform

Naivasha et al. (2025) mathematically formalize the idea that the very structure of the market can act as a predatory player that exploits information asymmetry to systematically extract value from individual traders. Along the argumentative line of this paper, the conclusion is that in a game where one player has superior information about the state of the world, that advantage persistently translates into higher expected payoffs, something that is not surprising from the standpoint of Game Theory but is for the standard view of markets as neutral allocation mechanisms.

This paper proposes replacing the representative agent with an ecosystem of heterogeneous agents, following the tradition of Schelling (1971) in urban segregation and Axtell (2001) in firm size distributions. In the context of financial markets, LeBaron (2006) has extensively documented how populations of agents with heterogeneous strategies produce return distributions with fat tails (very large drops or rises occur more often than a normal distribution would suggest) that representative-agent models with efficient markets cannot replicate.

The financial market, seen from this angle, is neither a pure stochastic process nor an efficient allocation mechanism, but rather a complex adaptive system where the heterogeneity of strategies, information asymmetry, and repeated interaction generate price structures that are simultaneously a consequence of the rules of the game and part of the environment to which the players must adapt. In this sense, price action would not be a random process but rather a kind of language with hidden messages in which agents communicate, reveal, and conceal their strategic intentions. It is in that language that the Wyckoff methodology proposes to read the game.

2.2 The Wyckoff methodology as a description of a game

Price action is one of the most widespread traditional techniques of market analysis. It starts from the idea that price alone, its movements and the graphical patterns it shows, contains all the information necessary to make decisions, without resorting to technical indicators, econometric models, or macroeconomic fundamentals. If price is the aggregate result of the decisions of all participants, its behaviour contains, in compressed form, the state of the game. Reading price action would then amount to interpreting the emotional state, the information, and the balance of forces among market operators.

Within this tradition, the methodology developed by Richard Wyckoff in the early 20th century represents one of its most rigorous systematizations, and one that not only retains its popularity among the trading community but has even been revitalized in the current environment of algorithmic trading and chart pattern recognition using machine learning systems, attesting to its practical operational validity.

Wyckoff, rather than cataloguing visual patterns, built an explanatory framework for why price behaves as it does, based on the logic of large operators. His central premise was that market fluctuations should be studied as if they were the result of a single strategic agent, the Composite Operator, against whom the small retail operator faces off. The term functions as a metaphor for the collective behaviour of institutions that, without explicit coordination, act with similar objectives, accumulate cheap and distribute expensive by exploiting retail predictability. From the standpoint of Game Theory, this logic does not describe a conspiracy, but rather a dynamic equilibrium that emerges from the decisions of agents under conditions of information and power asymmetry.

The Wyckoff system rests on three laws that, read from the perspective of Game Theory, acquire a conceptual dimension that their original formulation does not make explicit:

Law of supply and demand: price rises when demand exceeds supply and falls otherwise. Trivial in appearance, operative when one asks who generates that demand and with what information. Institutional demand is qualitatively different from retail demand, it does not react to price, it anticipates it.

Law of cause and effect: the magnitude of a trend is proportional to the duration and intensity of the sideways range that precedes it. Accumulation is the cause, markup is the effect, which introduces a temporal dimension into the analysis, the informed player patiently builds their position before the price reveals their intention.

Law of effort and result: volume is the effort, price movement is the result. When there is large volume but little movement, someone is absorbing that flow. This divergence would be the visible signature of institutional presence, the informed player involuntarily revealing their position through order flow.

The market cycle produced by these laws unfolds in four phases:

Accumulation: sideways phase following a downtrend. Institutions gradually absorb supply without revealing their intentions. Price does not rise because institutional demand is passive, it buys what retail sells, without pushing.

Markup: uptrend. Supply has been absorbed, demand dominates. Trend-following retail joins the movement once it is already underway.

Distribution: sideways phase following an uptrend. Institutions liquidate positions, transferring them to optimistic retail traders who interpret the sideways movement as consolidation.

Markdown: downtrend. Institutional support disappears. Retail is left trapped with positions bought at high levels.

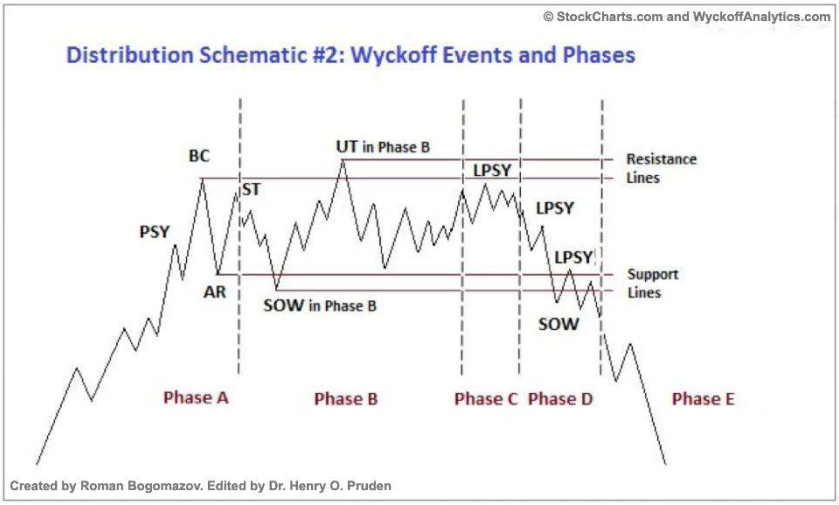

Fig.2 Wyckoff Distribution Schematic showing five phases (A–E) and key events such as Preliminary Supply (PSY), Buying Climax (BC), and Upthrust After Distribution

Within the accumulation and distribution phases, Wyckoff identified an internal microstructure that, in Game Theory terms, corresponds to specific strategic moves. The most relevant for this paper is the spring in accumulation and its symmetric counterpart, the upthrust after distribution in distribution.

The spring is a false bearish breakout below the range's support, where price briefly breaks below a relevant technical level, triggers retail stops, and attracts sellers who interpret the breakout as a bearish signal, then reverses immediately. This would be a strategic move by the informed player to obtain liquidity at more favorable prices, precisely exploiting feelings of panic among retail traders with open long positions, leading the unwary into impulsive selling. Its effectiveness depends on retail placing stops at obvious technical levels, something they tend to do repeatedly by following simple rules. The inverse mechanism operates with the complacency characteristic of FOMO (fear of missing out).

2.3 Taxonomy of the game and limits of analytical resolution

2.3.1 The type of game and its players

The forex market, observed from the standpoint of Game Theory, is a non-cooperative, sequential game repeated tick by tick, in which each agent updates their beliefs about the type and strategy of the others based on observed price and volume, which makes it a Bayesian game in the proper sense. Information asymmetry is not limited to the participants among themselves, the intermediary itself knows the client's exact position and stop loss, information to which no other player has access a priori.

The standard approach in microstructure, faithful to Kyle's (1985) framework, distinguishes only two types of players, the informed and the uninformed. In the real market, however, the group of the uninformed is not a uniform block, but rather a heterogeneous population of agents who follow trends, seek reversals, trade on pure noise, or provide liquidity with no directional opinion. Treating them as a single player would ignore precisely the heterogeneity that, according to this paper's central hypothesis, is a necessary condition for the emergence of Wyckoff structures.

The solution adopted distinguishes two players for purposes of theoretical formalization (informed/ uninformed) and five subtypes for purposes of computational implementation (*SM, retail, contrarian, noise, algo*), supported by the interpretation of a representative player's mixed strategy as the distribution of pure strategies within a heterogeneous population. The $N - player$ problem is thus posed, and this is precisely what simulation via the agent-based model aims to resolve.

2.3.2 The payoff matrix

To illustrate the strategic logic of the game in a concrete case, we will use as an example a situation at a support zone, representing the dilemma between the Smart Money (SM) and retail at a key support zone.

This would correspond to an accumulation phase in which price is below fundamental value and the SM must decide whether to accumulate position while retail decides whether to sell out of fear or buy by following the trend.

Table 1 captures this interaction in normal form, with the payoffs associated with each combination of decisions:

Table 1. Simplified payoff matrix in normal form: Strategic interaction between Smart Money and Retail in a support zone.

	Retail: Sell (Panic/*Stop*)	*Retail*: Hold/Buy
SM: Buy (Accumulate)	(10, -10)	(-5, 5)
SM: Do not act / Wait	(0, -5)	(0, 0)

The matrix captures the central conflict of the game, where the optimal payoff for Smart Money occurs when retail sells in panic (10, -10), the Wyckoff spring scenario, because the SM absorbs that liquidity without market impact. If retail holds or buys at support, the SM who insists on accumulating is forced to chase price upward, a costly outcome due to lack of counterparty (-5, 5). If the SM does not act and retail sells anyway, support breaks without anyone capturing the benefit (0, -5). If both remain passive, there is no transfer of wealth (0, 0).

At the methodological level, this matrix demonstrates that the SM does not have a pure dominant strategy of a passive nature, understood as one that offers the best outcome regardless of what the opponent does, we observe that the SM's best response changes depending on retail's action: if retail sells, it is in the SM's interest to buy (10 > 0), but if retail holds, it is in the SM's interest not to act (0 > -5).

The Wyckoff Spring resolves this indeterminacy through inducement. Instead of waiting for retail's move, the SM generates a false bearish signal that forces retail toward the Sell column, thereby turning its own accumulation strategy into the optimal response and allowing it to absorb that selling flow under favourable conditions.

If we seek to obtain a Nash Equilibrium in this hypothetical game, we will find that such an equilibrium generally does not exist in pure strategies at least. This occurs because in every state, one of the players has an incentive to change their decision, with no cell $(s_1, s_2 = equilibrium)$ where both are simultaneously satisfied.

The solution could be found in a Nash Equilibrium (NE) in mixed strategies, where the equilibrium is not a fixed action but rather a probability distribution. For the SM, the NE consists of being unpredictable, it cannot spring every time, because retail would learn not to place stops or to buy the breakout. The SM must "mix" its strategy: sometimes it accumulates with a sweep (spring), sometimes it accumulates passively, and sometimes it simply does not enter. For retail, the equilibrium consists of its behaviour being random enough that the SM cannot exploit it 100% of the time.

In formal terms, it could be expressed that the Nash Equilibrium for this game would be a pair of probabilities $(p*, q*)$ such that the SM launches a sweep with probability p and retail places its stop with probability q, such that neither can improve their expected outcome by changing those probabilities.

The 2x2 matrix thus serves to identify an equilibrium in mixed strategies where the SM optimizes its payoff through deception, while retail would maximize its payoff by applying a certain degree of randomness to its trading. But beyond how relevant this simplification may be, the reality of the market is that an N number of players operate, with great variability and distribution of actions, strategies, or stops, such that these equilibria depend on a virtually infinite number of interactions among players.

To this it must be added that the most realistic form of representing the game, given that it is a sequential game, would be the extensive tree-type form. The SM moves, retail observes and reacts, the SM absorbs... There is a temporal and informational structure (decision nodes). The problem obviously arises as soon as we introduce more than two players or several rounds of interaction, such that the tree quickly becomes unmanageable.

2.3.3 Limits of analytical resolution

The matrix presented above serves to visualize the conflict of incentives between players but, as noted, as a representation of the real market it proves insufficient. In an environment of asymmetric and dynamic information, the relevant solution concept would be the Bayesian Nash Equilibrium (BNE). Under this framework, a strategy profile is considered to be in equilibrium if no player can improve their expected payoff given their belief about the "type" of opponent they face. Although the theoretical existence of this equilibrium is guaranteed under general conditions, its explicit mathematical calculation is in practice intractable. This is because, while the general calculation of a Nash Equilibrium belongs to the PPAD-complete class (Daskalakis et al., 2009), determining Bayesian equilibria in pure strategies is NP-hard (Conitzer and Sandholm, 2008), and in both cases the complexity grows prohibitively with each added variable[4].

The reasons why the equilibrium is not computable in this context are structural and, moreover, cumulative:

Dimensionality and heterogeneity: the strategy space is continuous and multidimensional. Hundreds of participants operate in the market, whose interactions are not reducible to a simple bilateral competition without losing the richness of the microstructure. The interaction among a heterogeneous population of agents, each with their own rules and risk thresholds, would generate a combinatorial explosion that collapses any attempt at resolution through closed systems of equations.

Internal state and inventory: the SM's payoffs change over time. Unlike a single-round game, here the profit depends on the inventory accumulated over time. We are not facing an isolated event, but rather a dynamic process where the value of each decision is conditioned by the complete history of previous executions.

---

[4] PPAD (Polynomial Parity Arguments on Directed graphs) is a computational complexity class for search problems whose solution is guaranteed by combinatorial arguments, but with no known algorithm that finds it in polynomial time; a PPAD-complete problem is among the most representative of that class. An NP-hard problem, in turn, is at least as hard as any problem in the NP class. In both cases, practical difficulty grows prohibitively with the number of variables in the system.

Non-stationarity and belief updating: retail's beliefs about the nature of price movement mutate with each tick. In each period, the price resulting from order matching modifies agents' expectations, who in turn adjust their strategies for the following period. This feedback loop defines a dynamic, nonlinear, and path-dependent system.

Asymmetry with the intermediary: the presence of a broker or market maker with access to private order and position information introduces a layer of asymmetry that standard microstructure models do not usually contemplate, further removing the system from a predictable equilibrium.

This impossibility of analytical resolution is not a limitation of our computational capacity, but rather the most relevant property of the market. If the equilibrium were analytically solvable, price patterns would be predictable by design and the concept of genuine emergence would not exist.

It is precisely the fact that the game lacks a closed-form solution that justifies the introduction of Agent-Based Modeling (ABM). The ABM does not aim to directly predict market movement, its objective is to simulate agent interactions under simple rules in order to observe whether, from that process, the structures that Wyckoff and other theorists have anticipated theoretically and empirically spontaneously emerge.

2.4 Agent-Based Models and Complexity

We are thus faced with a complexity problem, understood as a system formed by a multitude of interconnected parts that give rise to developments and trajectories that are individually unpredictable, nonlinear, but not for that reason purely random or trivial. Since the nineties, Agent-Based Modeling has established itself as a response to this type of scenario in computational economics, it builds an environment from the bottom up and lets the system evolve. In contrast to the neoclassical approach, which starts from equilibrium and deduces the conditions that produce it, the ABM proceeds inductively from the behavioural rules of individual agents and observes what structures emerge from their interaction.

The agents of an ABM are, in essence, the players of the game, entities with available information, a set of possible actions, and a decision rule that determines how they act given the state of the environment. Their behavioural rules are their strategies, not necessarily optimal in the Bayesian sense, but rather heuristics that aim to capture how real participants decide under uncertainty. The ABM is, in that sense, the laboratory where the game is played.

The works of Thomas Schelling (1971)[5], are considered pioneering in this field, he demonstrated that patterns of extreme urban segregation can emerge from moderate individual preferences. In his initial model, no agent has a strong preference for segregation, but segregation emerges anyway as an unintended consequence of local

[5] Schelling considers agents who have a mild preference for similarity (for example, that at least 30% of their neighbours be of their same type), demonstrating that even such a moderate preference leads to total segregation in the model.

individual decisions. For his part, Robert Axelrod (1984) demonstrated, through strategy tournaments in the repeated prisoner's dilemma, that cooperation can emerge among selfish agents without the need for central authority or explicit agreements. The well-known tit-for-tat strategy[6], of extreme simplicity, systematically survived against more sophisticated strategies. Both works share the property that the global outcome was not contained in any individual rule, but rather emerged from the interaction.

In economics, this methodology has expanded over the last three decades to practically every domain. Axtell and Farmer (2025), in a review published in the Journal of Economic Literature, document its contribution to markets, industrial organization, labor markets, macroeconomics, and public policy, with notable achievements in understanding clustered volatility and systemic risk in financial markets. ABMs have begun to compete with DSGE models even in out-of-sample prediction, Poledna et al. (2022) developed the first macroeconomic ABM capable of competing with VAR and DSGE models, incorporating millions of heterogeneous agents. During the COVID-19 pandemic they were the tool of choice for modeling the economic effects of lockdowns, precisely because the bounded rationality of real agents proves to be a necessary property for capturing real dynamics (Axtell and Farmer, 2025).

The reason this methodology is particularly well suited to the problem at hand is endogenous emergence. In contrast to general equilibrium models, which converge to a stable fixed point, the ABM analyzes systems with continuous state change and emergent properties that cannot be inferred from the individual elements separately. The present model is thus situated at the intersection of Experimental Game Theory and Agent-based Computational Economics (ACE), where the realism of heuristic rules compensates for the absence of closed-form analytical solutions in high-dimensionality systems.

2.5 Equilibria and validation criteria

In this section we ask ourselves, given the game typology described previously, how this game ends, what an equilibrium is, who wins or loses, what the final resolution is. Although classical Game Theory generally considers an ending or equilibrium to be that set of strategies from which no player wants to unilaterally deviate, the truth is that in a real market, the game does not end, rather it transforms indefinitely. Smart Money manages an inventory indefinitely, rotating between accumulation and distribution while increasing it, and retail, taken in aggregate, does not go extinct but rather rotates positions, triggers stops, and re-enters. There is no terminal state, the market is a repeated game of indefinite horizon.

Three conceptions of this resolution are possible, P&L as payoff function, the minimax approach, and the emergence of attractors. The first two are presented here as theoretical grounding, the third is the validation criterion for our simulations.

P&L is the natural payoff function of the game. For Smart Money, success is measured in intertemporal terms, frequently assuming a negative latent P&L during accumulation,

[6] Tit for tat consists of starting by cooperating and, from then on, exactly replicating the opponent's action in the previous round, cooperating if they cooperated, defecting if they defected, which fosters stable long-term cooperation without needing to be permanently vengeful

the profit is realized afterward, during distribution. For retail, by contrast, payoff is measured trade by trade, with greater sensitivity to short-term volatility due to leverage. The logic of interest here is qualitative, the "turning point" at which the SM completes its inventory and retail's latent P&L becomes a loss as its stops are triggered.

The minimax approach, formulated by von Neumann (1928), establishes that in a zero-sum game there exists an optimal strategy for each player that minimizes the maximum possible loss, without the need for coordination or knowledge of the rival's strategy. In the financial literature, this problem is known as optimal execution, how an institution trades large volumes while minimizing its price impact (Almgren and Chriss, 2001), and it can be expressed in our model as a minimax-type strategic game.

Applied to this game, Smart Money's minimax strategy consists of maximizing its position while minimizing the information it reveals, the Wyckoff accumulation phases are essentially that optimization. Retail's defense is the stop-loss, protecting itself behind technical levels to minimize its own maximum loss. The point at which retail minimizes its risk is precisely where the SM finds the liquidity it needs to minimize its buying impact.

This paper does not address a per-agent P&L accounting, which is a matter of scope, not of impossibility. The optimal minimax strategy, on the other hand, inherits the same mathematical intractability of the general equilibrium described previously, it cannot be solved in closed form for this game, such that both logics remain as theoretical framework.

Following the tradition of complex adaptive systems[7], we replace the search for a static equilibrium with the identification of attractors, not as a stopping point, but as a region of the state space toward which the system tends to recurrently return. The system is not at rest at the attractor, but rather orbits around it.

Under this lens, the Wyckoff phases would not be equilibrium states, but rather dynamic attractors. "Accumulation" would be a region of transient stability where price moves sideways while the SM absorbs supply, while "markup" would be the transition trajectory that the system is forced to traverse between the accumulation attractor and the distribution attractor. Therefore, if an analytical equilibrium is computationally intractable, the question is under what microstructure conditions, under what behavioral rules, the system endogenously generates these attractors.

The "resolution" of the game in our ABM will be the emergence of a kind of orbit. The model will fail if it produces only structureless noise or if the patterns only emerge through forced programming that nullifies the agents' autonomy. The model will be considered successful if, upon letting the agents interact, price does not behave like white noise nor becomes locked in a fixed point, but rather begins to trace out accumulation sequences and trend-like transitions. That is, if the synthetic price path shows sideways ranges, false breakouts, and directional movements as a consequence of the internal dynamics,

[7] Complex adaptive systems belong to a broader tradition initiated by Prigogine, who demonstrated that open systems far from equilibrium can spontaneously generate ordered structures known as dissipative structures, instead of degrading into chaos (Prigogine & Stengers, 1984). This idea, transferred to the social sciences and economics by the Santa Fe Institute (founded in 1984, New Mexico, USA), gave rise to the study of systems composed of heterogeneous agents that interact locally, adapt to their environment, and generate collective properties not deducible from individual behaviour, a phenomenon known as emergence.

and if, in addition, these patterns correlate with the inventory state of the institutional SM operator.

3. ABM Development

3.1 Agent archetypes of the model

The model is structured as a synthetic continuous double-auction market in which heterogeneous agents submit orders that are matched in real time. In contrast to the standard binary formulation in market microstructure, which distinguishes between an informed agent and an uninformed agent (Kyle, 1985), this study divides the uninformed side into four subtypes. The resulting mixed strategy of these four subtypes can be interpreted, following the population games literature (Maynard Smith, 1982; Weibull, 1995), as the distribution of pure strategies of a representative uninformed player.

Agents remain fixed throughout the simulation, with no learning or evolutionary selection of strategies. A mechanism in which losers disappear and winners are replicated would be legitimate, in line with Axelrod's approach, but it would shift the focus of the research question and is therefore left as a direction for future research.

The five archetypes have been designed so that each performs a specific function in the price formation process. The detailed technical implementation rules are discussed in the parameterization section, while their general descriptive characteristics are as follows:

Smart Money

This agent embodies the Composite Man, the abstraction with which Wyckoff envisioned a single large operator behind the market. Smart Money is the contemporary term, popularized in the algorithmic trading community and in the ICT (Inner Circle Trader) methodology, for the same real institutional operators (banks, hedge funds, market makers). It is the only agent that knows the fundamental value (V) and operates under an absorption logic, its orders dampen the retail net flow within the same tick rather than chasing the price, and its aggressiveness depends on accumulated inventory, more urgency to distribute the more loaded it is, more patience to accumulate the lighter it is.

Two selectable logics were implemented for triggering the distribution phase. The first triggers distribution when the price moves away by a fixed percentage from the minimum reached during accumulation, a proxy for profit achieved.

The second replaces that proxy with a direct reading of market activity, aggregated over a recent time window, which combines into a single signal the entry of new buyers (euphoria) and the execution of sellers' stops that translate into buy orders (capitulation). Note that under this second logic, the exhaustion of the move is not the signal that triggers the SM's sale but its consequence, it is the SM's own distribution against that incoming flow that exhausts the momentum. Both logics share the same accumulation rule and the same inventory-based aggressiveness reduction mechanism, the difference lies exclusively in what triggers the phase change.

Trend-following retail

It buys because the price rises and sells because it falls, deciding on the basis of recent direction signals and operating with market orders. It is the engine of the markup and markdown phases, once a directional move has begun, its activity feeds and amplifies it. Its relationship with Smart Money is asymmetric, during distribution it buys precisely when the SM sells, becoming the involuntary recipient of the position that the informed agent wants to unwind.

Internal heterogeneity is introduced with respect to the stop loss, 70% operate with a fixed stop and capitulate on adverse moves, the remaining 30% (bagholders) have no forced exit mechanism and accumulate losses, generating trapped positions that act as latent liquidity. This temporal asymmetry in exit flows is one of the mechanisms with the greatest potential to generate spring-type liquidity events.

Contrarian retail

It interprets price extremes as reversal opportunities rather than continuation signals, with a longer time horizon that processes the recent price distribution rather than immediate direction. Its function is to sustain sideways ranges, without it any trend-following retail impulse would turn into an unbounded trend and accumulation and distribution ranges could not form.

The tension between the two profiles depends on the speed of the move, if the price drifts away slowly the contrarian dominates and the price reverts, if it moves away fast enough the trend follower interprets the move as a signal and reinforces the breakout before the contrarian can contain it. False breakouts and spring phases can emerge spontaneously from that tension. It shares with the trend follower the same stop loss heterogeneity, 70% with a fixed stop, 30% bagholders.

Noise trader

It buys or sells with no apparent reason, with small orders and no temporal pattern. It represents uninformed retail flow. Its presence is necessary for the system to function, without noise the informed agents would have no anonymous counterparty behind which to hide and the emergent patterns would reflect only the deterministic mechanics of the programmed rules. That Wyckoff structures appear despite the noise is more convincing evidence of endogenous emergence than if they appeared in an artificially clean environment.

Medium-algorithmic agent

Represents participants with non-directional quantitative strategies, neither informed like the SM nor reactive to price like the rest of retail. Given the variability of real algorithmic strategies, it is modelled as a simplified market maker that captures their net effect, it provides liquidity on both sides of the price and withdraws when the environment becomes uncertain, under normal conditions it stabilizes the market and under high volatility it widens margins and reduces exposure. It is the most dispensable agent with respect to the central hypothesis, its omission does not alter the SM-retail mechanism,

and it serves mainly as a contrast variable to assess the effect of its liquidity on the emergent patterns.

3.2 Model operation scheme

The model is a synthetic market where N heterogeneous agents interact in discrete time over T ticks, with no real data and no exogenous shocks. Everything that happens to the price is a consequence of the interaction between agents and a stochastic noise component, so that any structures that emerge can only be attributed to the system's internal dynamics.

In each tick a four-step process is executed. First, each agent checks whether its position has exhausted the maximum allowed time, in which case it is closed automatically. If not, it evaluates its decision rule and issues an order. Second, once all agents have been processed, retail positions that have hit their stop level are liquidated. Third, the matching engine aggregates the net flow resulting from all orders in the tick, both discretionary and forced. Fourth, the price is updated as a function of that net flow and a noise component:

$$P_{t+1} = P_t + \alpha \cdot (\text{Net flow}) + \epsilon$$

where α weights the impact of the flow according to its nature (the exact form of α, which distinguishes the impact of aggressive orders from that of the SM's silent absorption). The price of the next tick is not an input but a variable that emerges from this process.

The primary output is a price series P which is aggregated into OHLC candles for visual analysis. Alongside it, the model records the SM's net inventory at each tick, which allows checking whether the SM accumulates while the price moves sideways or falls, and distributes while it rises.

3.3 Model Architecture and Parameters

3.3.1 Parameters and initial conditions

The simulation takes place in a closed and purely endogenous market environment, where the price emerges exclusively from the interaction of the agents and a stochastic component. The duration $(T = 5000\, ticks)$, the initial price and the fundamental value $(P_0 = V = 100)$, and the population size $(N = 200)$ are fixed constants of the model.

The remaining parameters are grouped into five blocks, population, price dynamics, common retail behaviour, Smart Money logic, and exit management. The model implements two selectable triggering logics for the SM's distribution phase, the "distance logic" operates according to the relative distance of the price from the accumulation minimum, and the "flow logic" replaces that criterion with an aggregated market activity

signal over a recent time window. Several parameters are exclusive to one branch or the other, indicated in Table 2 below:

Table 2: Configurable model parameters

Parameter	Default Value	Description
seed	67	Random generator seed
prop_smart_money	5%	Proportion of SM agents (n=10)
prop_contrarian	20%	Proportion of contrarian agents (n=40)
prop_retail	75%	Resulting proportion of retail (n=150).
α (market)	0.005	Price impact of aggressive (market taker)
α (absorption)	0.0005	Price impact of the SM's passive absorption
σ_noise	0.03	Standard deviation of stochastic noise. Alternative levels, Medium 0.05, High 0.10
size_trend / size_cont	1.0 / 1.0	Order volume, trend-following and contrarian retail
size_noise / size_algo	0.5 / 0.5	Order volume, noise and algorithmic
k_trend	5	Memory window of trend-following retail
k_contrarian	20	Memory window of contrarian retail
θ	1.0	Activation threshold (z-score) for the contrarian
stop_dist	1.0	Retail stop loss distance
prob_stop_loss	70%	Probability that a retail agent has an active stop
duracion_max_posicion	20	Time-stop, forced closure by time
factor_inv_sm	50	Multiplier for the SM's inventory limit (maximum = factor × n_SM = 500)
factor_agresividad_sm	2.5	SM order size multiplier
logica_sm	—	Selector, "original" or "flow"
umbral_dist_pct	3%	Price distance from the accumulation minimum that triggers distribution
k_flujo	10	Time window for the aggregated flow signal
umbral_flujo	75	Flow signal threshold that triggers distribution
tolerancia_pct	0.01%	Asymmetric extension of the SM's buying range above V
capital_inicial	10	Initial capital of TREND/CONT agents, allows bankruptcy if it reaches zero. Inert (infinite) under the original logic

3.3.2 Decision rules by type

Each agent profile implements an independent decision function, which evaluates the recent price history and the agent's internal state (inventory, position) to return the direction of the flow and the order size. Smart Money is the only profile with two interchangeable variants, the rest of the agents (trend follower, contrarian, noise, algorithmic) have a single decision rule, described below.

3.3.2.1 Institutional agent or Smart Money

The SM operates in two phases, accumulation and distribution, with a common rule for the order size (n) that modulates aggressiveness according to inventory:

$$R = 1 - \frac{I_t}{I_{\max}} \text{ in accumulation, } R = \frac{I_t}{I_{\max}} \text{ in distribution}$$

As it approaches its inventory limit, it reduces the size of its orders. The net order volume (n) is calculated as:

$$n = \max(1, \text{round}(\text{k} \cdot R))$$

where I_{t} is the SM's accumulated inventory, I_max its limit $(\text{factor_inv_sm} \times \text{n_SM})$ and k the aggressiveness multiplier.

What distinguishes the two implemented and selectable logics is the buying condition in accumulation and the criterion that triggers the switch to distribution, the rest of the mechanics is identical.

The distance logic buys while $P_{\text{t}} < V$, with k fixed at 2.5, and triggers distribution when the percentage distance between the price and the minimum reached during accumulation exceeds a threshold (umbral_dist_pct, 3% by default).

The flow logic buys while $P_t \leq V \cdot (1 + "tolerancia_pct")$, a slightly wider range asymmetric to the upside, with configurable k (factor_agresividad_sm), and triggers distribution when an aggregated flow signal exceeds a threshold:

$$S_t = \sum_{i=t-k_{\text{flujo}}+1}^{t} \text{flujo_mercado}_i$$

The market flow at each tick aggregates all aggressive orders in the system plus those generated by stop execution, capturing in a single variable the euphoria, the entry of new buyers chasing the move, and the capitulation, the forced closing of shorts whose stop

is hit and which executes as a buy order indistinguishable from the previous one. S_t sums that signal over a window of k_flujo ticks so that the trigger responds to sustained pressure rather than an isolated noise spike.

Under this logic, the exhaustion of the move is not the signal that triggers the SM's sale but its consequence, it is the SM's own distribution against that incoming flow that exhausts it. Under both logics, the SM returns to the accumulation phase when its inventory is depleted $(I_t \leq 0)$.

3.3.2.2 Trend-following retail agent

This agent represents the mass psychology that chases the move once confirmed. Its logic is based on short-term momentum, it uses a short memory window (k_trend) and calculates the direction of recent price changes, if the sum of the signs is clearly positive, it assumes an uptrend and buys, and vice versa.

$$Signal_t = \sum \text{sign}(P_{t-i} - P_{t-i-1})$$

Buy: if $Signal_t > 0$; Sell: if $Signal_t < 0$

To simulate the retail investor's limited capital and avoid infinite overexposure that would distort the model, the agent can only open a position if its current state is inactive. Once in a position, it cannot add further volume until the position is closed, either by time or by stop.

It operates through market orders, with size defined by size_trend. In the matching engine, these orders are the ones that drive the price up or down.

3.3.2.3 Contrarian retail agent

This agent represents the operator who looks for tops and bottoms, betting on mean reversion. Its role is crucial in generating resistance at the extremes of Wyckoff ranges.

It uses a medium memory window (k_contrarian), wider than that of the trend follower. It calculates the standard deviation σ and the mean μ of prices within that window to determine the level of overbought or oversold conditions by means of a z-score,

$$z = \frac{P_t - \mu}{\sigma}$$

Buy if $z < -\theta$, Sell: if $z > \theta$ that is, it sells when the price has deviated upward beyond the threshold, and buys when it has deviated downward.

By selling when the price rises fast and buying when it falls, this agent brakes the trend follower's moves, making it easier for Smart Money to find a counterparty for its orders

without the price running away prematurely. Like the trend follower, it only operates if its state is inactive, with order size defined by size_contrarian.

The asymmetry between k_trend (short) and k_contrarian (wider) is what allows the transition between sideways movement and trend. Faced with a slow move, the contrarian reacts in time and contains it, faced with a fast move, the trend follower breaks the range before the contrarian's moving average has time to shift.

3.3.2.4 Noise trader agent

This agent has no data-based strategy, its function is purely stochastic. It does not consult the price history or its own inventory state, at each tick it decides to buy or sell with symmetric 50% probability, with size defined by size_noise. Direction: $\{1, -1\}$ with probability $(P = 0.5)$

Its function is not to provide direction but background noise, without it the informed agents would have no anonymous counterparty behind which to hide, and any emergent structure would be attributable only to the deterministic mechanics of the other agents' rules.

3.3.2.5 Medium-algorithmic agent

For this agent, the logic does not depend on where the price is going but on how much it is moving. It monitors the standard deviation of the last 20 prices, if it exceeds a critical threshold of 0.5, it goes into total inactivity and stops issuing orders, if not, it remains active. Window and threshold are not configurable parameters in the current version of the code, unlike k_trend and k_contrarian.

When active, it operates analogously to the noise trader, providing market liquidity on both sides of the price through random orders, with size defined by size_algo.

The disconnection rule could be key to the emergence of structures, if the disappearance of part of the intermediate flow when the market becomes unstable contributes to a liquidity vacuum, favouring more violent price moves in the face of imbalances between the SM and the rest of retail.

3.3.3 Stop execution

Stop execution is treated as an independent module, separate from each agent's decision rules for the sake of clarity. It acts after capturing the tick's discretionary orders but before matching, so that retail's forced capitulation enters the same price formation process as the rest of the orders, and it is precisely that capitulation that provides the liquidity the SM absorbs.

For each agent with an active position and stop enabled, the exit level is calculated as:

$$Level = Entry_{Price} - (Side \cdot StopDist)$$

Where *Side* takes a value of $(+1)$ for buys and (-1) for sells, so that the stop sits below the entry price on longs and above it on shorts.

When the price crosses that level, the agent issues a market order in the direction opposite to its original position, consuming liquidity and pushing the price in the direction of the stop, which can feed order cascades. When the stop is executed, the agent becomes inactive and its position is reset, in addition, under the flow logic, the closure settles the position's PnL against the agent's capital, if that capital reaches zero, the agent is permanently excluded from the rest of the simulation.

There is a hypothesis that the triggering of retail stops provides the SM with part of the liquidity it needs to complete its inventory without chasing the price, to the extent that such triggering coincides temporally with its zone of activity. This is not imposed by design, the SM's buying zone is defined solely in relation to the fundamental value V, with no direct reference to stop_dist, so the coincidence, if it exists, would be an emergent effect of the interaction between parameters rather than an explicit mechanism of the model. This has not been verified against the runs used for the final figures.

3.3.4 Matching, price formation and main loop

This block of code consolidates all execution intentions at each tick to determine the new equilibrium point. The process does not use an explicit order book *(Limit Order Book)*, but is instead based on the interaction between net flows and a price impact function. In turn, agents' inventory is updated at the moment the order is issued.

First, the system segregates the orders received into two functional categories according to their behaviour with respect to price:

- ***Market taker flow*** $(\boldsymbol{M})$**:** issued by retail agents (trend followers and contrarians), noise traders, and stop execution. These orders are executed "at market", representing the immediate pressure that moves the price.
- ***Market maker flow*** $(\boldsymbol{L})$**:** issued mainly by the SM. Although they are integrated into the tick's matching, they function as limit orders, since they are only activated when the price reaches predefined value zones. Their role is to provide the counterparty needed to absorb excess retail flow.

The formation of the new price P_{t+1} id governed by the following expression:

$$P_{t+1} = P_t + \alpha_{\text{market}} \cdot \left(\sum M_i\right) + \alpha_{\text{absorption}} \cdot \left(\sum L_j\right) + \varepsilon$$

Where:

- P_t: price of the previous tick.
- α_market: liquidity coefficient (price sensitivity to net flow).
- $\alpha_absortion$: impact coefficient of the SM's maker flow, substantially lower than α_market by default.
- $\sum M_i$: sum of *taker flow* (buys minus sells).
- $\sum L_j$: sum of passive absorption (maker) provided by the SM.
- ϵ: stochastic White noise defined by $\epsilon \sim N(0, \sigma_{noise})$

Unlike a real order book where transactions are bilateral, our model applies a simplification and executes a global balance at each tick, summing all buy intentions and subtracting sells.

The key to the model lies in the fact that Smart Money uses its passive flow to counteract the direction of retail. When there is an excess of retail supply in a support zone, the SM generates a flow L_j of opposite sign, resulting in:

$$\left|\sum M_i + \sum L_j\right| < \left|\sum M_i\right|$$

This reduction in net impact is what produces the accumulation phase. Instead of a linear move, the price moves sideways due to the absorption of supply. The SM thus completes its inventory silently, the price does not rise because the SM absorbs retail sales, maintaining equilibrium until the inventory is satisfied, the SM ceases its passive activity, and the next aggressive (taker) flow breaks the range.

The loop goes through the agents one by one, each agent first checks whether its position has exhausted the maximum time, closing if so, or evaluates its decision rule if not. Once all agents have been processed, stops are executed on the positions that remain open, and finally matching updates the price. This order guarantees strict causality, no decision at tick t is made with information from tick t itself, all of them are based on the state resulting from tick $t-1$. The complete pseudocode and initialization details are documented in the project repository.

3.3.5 Post-processing and visualization

The raw price series P_t generated by the loop is aggregated into Japanese candlesticks (OHLC) to facilitate visual validation of Wyckoff structures. The model aggregates the series into a fixed number of 200 OHLC candles, each composed of $m = T/200 = 25$ ticks, with open and close defined by the first and last price of the block, and high and low determining the wicks. Each candle also records the SM's inventory at close, which allows overlaying the latent institutional positioning variable on the price chart.

Validation relies on two panels. The upper one shows price evolution through the candles, with a reference line at the fundamental value V, there the SM's footprints are sought, springs (long lower wicks that break previous supports), upthrusts or UTADs (upper wicks that trap buyers before a decline), or SOS (sequences of large-range bullish-bodied candles following an accumulation phase). The lower one shows the SM's accumulated net position, the validation sought is a characteristic divergence between sideways price and growing inventory, evidence of silent accumulation by the informed agent.

3.3.6 Use of generative AI

During the preparation of this manuscript, the author used Claude (Anthropic) and DeepSeek for general queries, code assistance, and writing support. All content was reviewed and edited by the author, who takes full responsibility for the final version of the manuscript.

4. Models and Simulations

The validation criterion rests on two complementary approaches. The first is quantitative and the second is graphical. Regarding quantitative validation, in the accumulation phase, defined by sustained growth in the SM's inventory, a positive slope of the inventory is expected, $\Delta Inv_SM_t > 0$, consistently across the window, accompanied by a price that remains lateral, with reduced dispersion, $\sigma(P_t)$ in the accumulation window lower than or comparable to the mean volatility of the rest of the series. This reflects that the SM's absorption does not move the price against it, a necessary condition for accumulation to be silent.

In the distribution phase, by contrast, the inventory drops sharply and the price moves in the same direction, not the opposite one, so the expected correlation is positive, $\text{Corr}(\Delta P_t, \Delta\text{InvSM}) > 0$ within that window.

A time-lag criterion is added to this, calculating the correlation between ΔP_t and $\Delta\text{Inv}_{SM,t-k}$ for different values of k, in order to determine whether the SM's selling precedes the price drop or whether, conversely, the SM reacts to a price movement already under way.

Note that the patterns sought, when they appear, occupy a stretch of the series, one that also varies from one simulation to another, not the entire 200-candle series (this can be seen in the samples shown below in 4.1 and 4.2). Applying a correlation, dispersion, or other statistical treatment over the whole sample therefore dilutes the target signal with the rest of the candles where the mechanism is not active. For this reason we chose to sample window by window across a limited number of runs, since generating a massive statistical sweep would first require solving the problem of automatically detecting where, within each run, the accumulation/distribution window is located, a task with an identity of its own that falls outside the scope of this work and is noted as a future line of research.

As stated above, the second approach is visual, the inspection of the price series in OHLC candlesticks in search of the structures described by Wyckoff, accumulation followed by markup, distribution followed by markdown, false breakouts in stop zones. This approach has an established role in the pattern-oriented validation customary in ABM (Grimm et al., 2005) and in the literature on stylized facts in simulated markets (Cramer & Trimborn, 2019); it provides evidence — though not proof — and it serves the qualitative exploration of the model's behaviour.

4.1 Distance logic

The distance logic models the Smart Money as an agent with privileged information who adjusts its position according to the percentage distance between the price and a recent accumulation low.

Figure 3 shows a representative simulation (seed 67), with the price series in OHLC candlesticks and the evolution of the SM's inventory over the same time interval.



Fig.3 Simulation of the distance-logic version, full cycle (seed 67, SM Inventory Factor=35)

The inventory grows gradually between candles 0 and 150, while the price stays within a relatively narrow range, with no clear directional trend, which is the expected pattern of silent accumulation. Quantitative validation confirms this, between candles 128 and 150 the price's standard deviation is 0.24 against 0.53 for the simulation as a whole (45% of mean volatility), while the inventory grows with a mean slope of 4.6 units per tick.

The inventory stays at its maximum until candle 185, at which point it drops abruptly together with the price, consistent with a distribution concentrated over a few candles. Between candles 185 and 192, the correlation between ΔP and $\Delta InvSM$ is positive (0.26) and decreases from 0.28 at $k = 1$ to 0.22 at $k = 5$, consistent with the SM's selling preceding the price movement, although the magnitude is moderate and should not be read as strong confirmation of causality, only as an indication consistent with the hypothesis.

To check that the pattern does not depend on a particular seed, the same analysis is repeated with a second simulation (seed 41). Figure 4 shows an additional full cycle, with a first phase of accumulation and distribution between candles 0 and 90 approximately, followed by a second accumulation phase toward the end of the series.



Fig.4 Simulation of the distance-logic version, full cycle (seed 41, SM Inventory Factor=50)

The inventory grows steadily, with an additional lateral phase until candle 40, while the price stays within a narrow range, and remains at its maximum until it drops abruptly coinciding with the price peak of the entire series (≈102.3 at candle 90). After an extended period of zero inventory, a second accumulation cycle begins, which does not reach a completed distribution within the observed interval.

Quantitative validation reinforces this reading, above all in accumulation. Across the three windows identified so far (seed 67, candles 128-150, and seed 41, candles 0-40 and 145-180), the price's standard deviation sits consistently between 43% and 48% of mean volatility, a regularity that reinforces price laterality as a robust signature of the mechanism, rather than a feature of a single run. In distribution, the correlation keeps the same positive sign but with a more modest magnitude (0.11 against 0.26) and a less defined lag (peaking at $k = 2$ instead of $k = 1$), which confirms that distribution is, of the two phases, the weaker and more variable signature across runs.

4.2 Flow logic

The flow logic models the Smart Money as an agent that adjusts its position according to an aggregate flow signal, calculated over a window of recent ticks, rather than the distance to the accumulation low. Unlike the distance logic, here retail operates with finite capital, which allows breakdowns to be observed when market pressure exhausts its margin, but above all it aims to capture the signal of buyer euphoria and seller capitulation on the SM's part.

Figure 5 shows a representative simulation with seed 67 and flow_threshold=100 set by default:



Fig.5 Simulation of the flow-logic version, full cycle (seed 67)

Unlike the distance logic, several minor phase changes appear before the main cycle, which begins around candle 115, with the inventory growing to a maximum near 3,800 units around candle 150; unlike the previous cases, the price does not stay strictly lateral, it shows a moderate upward trend that intensifies before the distribution, abrupt between candles 185 and 195, coinciding with the series' price peak.

Quantitative validation confirms, broadly speaking, the same structure observed under the distance logic. During accumulation (candles 115-150), the price's standard deviation is 0.23, against 0.57 for the simulation as a whole, approximately 40% of mean volatility, a ratio consistent with the range obtained under the distance logic (43-48%), although here a slight directional trend in the price within the window is more clearly visible, one that reduced dispersion alone does not fully capture.

During distribution (candles 190-197), the correlation between ΔP and $\Delta InvSM$ is 0.31, the highest obtained across the whole analysis, and it remains positive and decreasing with the time lag, from 0.26 at $k = 1$ to 0.21 at $k = 5$. This pattern, more pronounced than under the distance logic, is consistent with SM selling pressure preceding the price movement in the short term, reinforcing the reading of distribution as the moment when the informed position is transferred to the market more directly.

Next, another simulation is examined with a different seed to check for the emergence of the same pattern (seed 23, flow_threshold=100), shown in figure 6:



Fig.6 Simulation of the flow-logic version, full cycle (seed 23)

The inventory grows quickly between candles 8 and 20, after a brief, unstable initial stretch, until stabilizing around 1,430 units, a considerably lower level than the maximum reached under the previous seed. This plateau holds until approximately candle 80, at which point a rapid distribution occurs that returns the inventory to zero within a few candles. Unlike the previous simulations, the price does not stay lateral during accumulation, it shows a sustained upward trend that continues even after the distribution; even though the SM has taken its profit (selling higher than it bought), this is a more demanding case for the hypothesis, as it coincides with an underlying trend unrelated to the SM's activity.

Here, accumulation (candles 8-20) shows a dispersion ratio of 32.5% (0.34 against 1.05), the lowest observed so far, and distribution (candles 77-82) shows a correlation of 0.14, with a virtually flat lag pattern, without the monotonic decline observed in the previous simulations.

The pattern holds, in the sense that both metrics keep the expected sign, but with a clearly smaller magnitude. This suggests that the quantitative signature of the accumulation-distribution mechanism is sharper when the price is not subject to an underlying trend unrelated to the SM's activity, a condition that the flow logic, unlike the distance logic, does not impose by design.

4.3 Robustness and calibration

To assess the extent to which the accumulation-distribution pattern depends on a specific calibration, two complementary exercises are examined. The first tests different seeds with no adjustment at all, the second fixes a seed and progressively adjusts the model's parameters.

Grid 1, simulations with random seed

Five simulations are generated with randomly chosen seeds, without fixing or adjusting any parameter relative to the default values in table 3. The aim is to check whether the structure emerges spontaneously in a non-selected sample, or whether it requires a particular combination of initial conditions on its own.

Table 3: Pattern emergence in simulations with random seed

	Inventory with a sawtooth pattern, with no clear sustained plateau phase. The price, moreover, is in an underlying downward trend. It does not resemble a recognizable accumulation/distribution structure, it looks more like noise with repeated entries and exits.
	Small, scattered inventory in brief spikes, with no real accumulation plateau at any point. The price moves within ranges without the inventory giving a clear reading of which phase corresponds to which stretch.
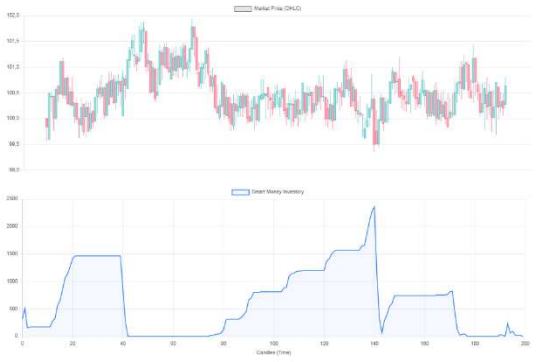	Two cycles are observed, the second more convincing than the first. The first accumulation (candles 0-40) resolves into a clean distribution between candles 40 and 45. The second, longer and more complex (candles 75-140), ends in an abrupt inventory drop at candle 140, coinciding with the price peak of the entire series, a clear example of distribution concentrated over a few candles, the same signature quantified in 4.1 and 4.2.

	Accumulation occurs in two steps (candles 0-3 and 60-65), with the price rising clearly during the first phase, with no laterality. The inventory stays at its maximum until the end of the window, without distributing within the 200 candles. This is a case of accumulation without the cycle closing, useful as an example of an ongoing process, not a complete pattern.
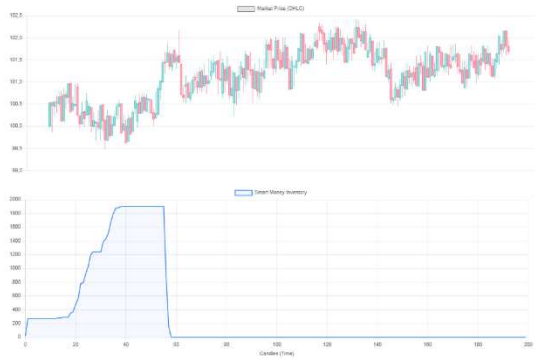	Accumulation holds until approximately candle 55 (the inventory goes from 0 to 1,900 units, with the price lateral between 99.5 and 101), followed by an abrupt distribution and breakout between candles 55 and 60, where the inventory drops to zero and the price jumps to a higher level that it maintains for the rest of the series. This is the structure of accumulation followed by a sustained breakout.

Of the five simulations, two show no recognizable structure, and the remaining three do reproduce, with varying degrees of clarity, the pattern described in the theoretical framework. The result suggests that the pattern's emergence does not depend on a special calibration, it appears with some regularity even under non-selected conditions, although its clarity varies from one run to another.

Grid 2, fixed seed with progressive calibration

The second exercise reverses the approach, a seed is fixed (seed 11) and the SM's trigger parameters and the contrarian's sensitivity are progressively adjusted, to check whether calibration allows the underlying microstructure to be revealed more clearly, which we observe in table 4 below:

Table 4: Effect of progressive calibration on the pattern's clarity (seed 11)

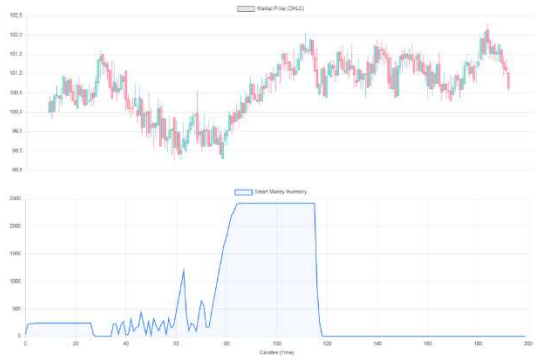	With the default parameters, the pattern turns out reasonably complete. After a noisy initial stretch (candles 0-60) with no real accumulation, the inventory grows steadily until candle 90, with the price contained within a range, and distributes abruptly around candle 115.
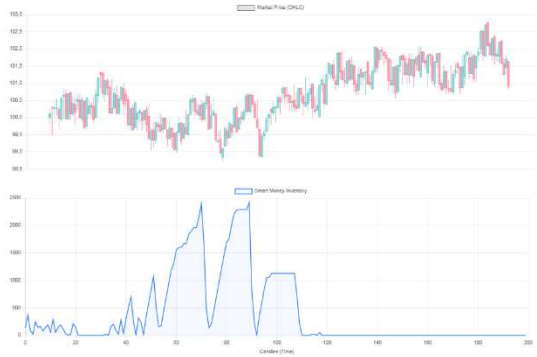	When the distribution threshold (flow threshold) is raised, two rapid accumulation phases appear, separated by an intermediate drop around candle 75, the first up to a moderate maximum around candle 30, the second somewhat larger around candle 68, before distributing around candle 115. The result suggests also raising the trigger threshold and the contrarians' sensitivity, to mitigate that early distribution.
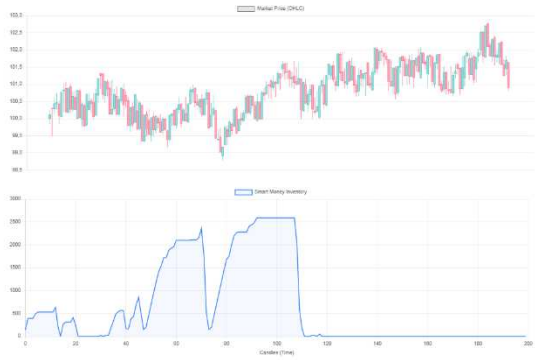	The second accumulation stretch reaches a higher maximum (candle ~105) and holds for longer before distributing, but the fragmentation into two phases persists. We insist on raising both parameters at once to try to merge the two phases into one.
	With the trigger threshold and the contrarians' sensitivity now adjusted jointly, the two accumulation phases merge into one, continuous between candles 0 and 115, followed by a single distribution. The price, however, does not return to its previous range after the drop, it stabilizes at a lower level for the rest of the series.

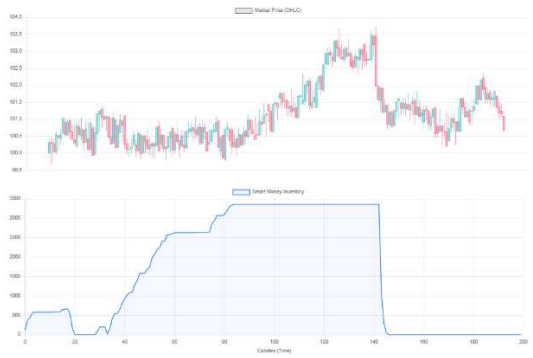

With the joint adjustment of both parameters, accumulation becomes practically continuous, it holds without partial discharges, and distributes abruptly at the end, coinciding with the price peak. The cleanest result of the series.

This example result (with seed 11, k_contrarian=40, flow_threshold=160, tolerance=0.002, k_flujo=20, capital_inicial=20) reproduces the accumulation-distribution structure described in the theoretical framework with notable fidelity, without any rule of the model imposing that shape directly.

One can observe the accumulation cycle carried out by the informed player and the distribution, abrupt as it is, where it realizes its profit, displaying a recognizable chart pattern known as a "double top" including a deception in the form of a second wick to trap its opponents in the game.

5. Discussion and conclusions

The model formalizes, by means of a Bayesian game with asymmetric information between an institutional agent and a heterogeneous population of retail agents, the hypothesis that the accumulation and distribution structures described by the Wyckoff methodology can emerge from a small set of individual behavioral rules, with none of them directly imposing the aggregate pattern.

The results presented show that this hypothesis is consistent with structures where the SM's inventory increases while the price stays in a lateral range over time, followed by the onset of an upward phase, and a subsequent distribution and sell-off with declines in both price and inventory, clearly evoking the cycles described in the theoretical framework regarding Wyckoff-type price action patterns. Traps such as *spikes* or *upthrust after distribution* even appear, with candles leaving elongated shadows before the drops.

While the formation of these figures can be observed across a large number of runs, in many of the random or fixed-seed simulations the emergence of patterns could not be observed, which is accepted as expected given that the stochastic component is present in the simulation at all times.

The proposed set of rules produces, consistently across different seeds and under both variants of the informed agent's programmed trigger logic, accumulation phases with reduced price dispersion and distribution phases with price movement in the same direction as the inventory.

Note that different sets of rules, with different behavioral assumptions, could generate graphically similar structures, an equifinality problem described in the pattern-oriented validation of agent-based models (Grimm et al., 2005), so the match between the simulated pattern and the pattern described by the technical literature shows that these mechanisms are sufficient to produce it, not that they are the only ones capable of doing so.

The quantitative validation criteria show that during silent accumulation the price stays within a reduced dispersion range while the inventory grows, and that the positive correlation with time lag in distribution offers evidence of a causality that precedes the effect, with the informed agent's selling running ahead of the price drop. This second signal is, however, weaker and more uneven across simulations, with correlation values ranging between 0.14 and 0.31 depending on the seed, and a lag that in some cases decreases clearly and in others stays practically flat. In some of these distribution phases, abrupt drops in price and inventory are also observed within a very short interval, a pattern that points to a sharper triggering mechanism than correlation alone describes, and which is flagged here as an aspect to explore further.

The minimax equilibrium is retained solely as theoretical motivation for why an agent with privileged information would have an incentive to behave as modeled; its direct computation is intractable at this scale and it is not used as an operational validation criterion. The zero-sum game's counterpart, measurable losses for the uninformed agent, is only partially observable, and only under the model variant with finite capital for retail agents, where market pressure can go as far as exhausting their margin.

The model has limitations that open up concrete lines of future work. The informed agent's phase-change criterion, based on a fixed percentage distance to the price or on an aggregate flow signal over a time window, is a functional but exogenous proxy, it does not arise from an observable market condition. The variant with incoming flow driven by buyer euphoria and seller capitulation partially explores an alternative to this, besides adding the finite-capital variable, allowing retail agents' capacity to react to erode with the outcome of their own trading.

An additional simplification remains an open line, which is that retail agents' exit from positions is resolved through a fixed time criterion, a practical concession that departs from realistic risk management and that could be replaced by a criterion tied to the evolution of capital or price momentum.

One line of future work consists of inverting the calibration problem, starting from price patterns already documented in the technical literature and systematically searching for which set of behavioral rules produces them, rather than sweeping parameters over a fixed logic. That is, given that we already know and routinely observe charts and price action, the task is to find the logical rules that produce them, to reveal them, rather than sweeping parameters of a few a priori programmed logics. This approach connects with mechanism design and simulation-based inference, and could also guide the search for quantitative signals, for instance through econometric methods applied to the liquidity imbalances that precede trending moves, or through the definition of price zones that act as attractors.

References

Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3, 5-39.

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.

Axtell, R. L. (2001). Zipf distribution of U.S. firm sizes. *Science*, 293(5536), 1818-1820.

Axtell, R. L., & Farmer, J. D. (2025). Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature*, 63(1), 197-287.

Conitzer, V., & Sandholm, T. (2008). New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2), 621-641.

Cramer, S., & Trimborn, T. (2019). *Stylized facts and agent-based modeling* (arXiv:1912.02684). arXiv. https://doi.org/10.48550/arXiv.1912.02684

Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 195-259.

European Securities and Markets Authority (ESMA). (2018). *Decision (EU) 2018/796 of 22 May 2018 to temporarily restrict contracts for differences in the Union in accordance with Article 40 of Regulation (EU) No 600/2014 of the European Parliament and of the Council*. Official Journal of the European Union, L 136, 50-56.

Evans, M. D. D., & Lyons, R. K. (2002). Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1), 170-180.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383-417.

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., Thulke, H.-H., Weiner, J., Wiegand, T., & DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, 310(5750), 987-991.

Harsanyi, J. C. (1967). Games with incomplete information played by "Bayesian" players, I-III. Part I. The basic model. *Management Science*, 14(3), 159-182.

Kirman, A. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, 108(1), 137-156.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315-1335.

LeBaron, B. (2006). Agent-based computational finance. En L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics* (Vol. 2, pp. 1187-1233). North-Holland.

Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press.

Naivasha, P., Musumba, G., Gikunda, P., & Wandeto, J. (2025). A game-theoretic model of forex trading with stochastic strategies and information asymmetry. *2025 IST-Africa Conference (IST-Africa)*, Nairobi, Kenya, 1-9. https://doi.org/10.23919/IST-Africa67297.2025.11060511

O'Hara, M. (1995). *Market microstructure theory*. Blackwell.

Poledna, S., Miess, M. G., Hommes, C., & Rabitsch, K. (2022). Economic forecasting with an agent-based model. *European Economic Review*, 151, Article 104306. https://doi.org/10.1016/j.euroecorev.2022.104306

Prigogine, I., & Stengers, I. (1984). *Order out of chaos: Man's new dialogue with nature*. Bantam Books.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2), 143-186.

von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100, 295-320.

Walras, L. (1874). *Éléments d'économie politique pure, ou théorie de la richesse sociale*. L. Corbaz.

Weibull, J. W. (1995). *Evolutionary game theory*. MIT Press.

Statements & Declarations

Consent for publication

Not applicable.

Data availability

No external datasets were used in this study. All data analysed are generated endogenously by the agent-based model during simulation and can be reproduced using the code and parameters described in Section 3, available at the repository cited above.

Code availability

The code necessary to reproduce the simulation results presented in this study is available at *[link removed for double-anonymous review]*. An interactive web-based version of the simulator can be accessed at https://simulador-patrones-abm.vercel.app/.