DOI: 10.46541/978-86-7233-439-5_531

XXXI International Scientific Conference
Strategic Management
and Decision Support Systems
in Strategic Management
Sm2026

Subotica (Serbia), May 18, 2026

Ognjen Radović
Faculty of Economics, University of Niš
Niš, Serbia

e-mail ognjen.radovic@ekonomski.rs

Jelena Radojičić
Faculty of Economics, University of Niš
Niš, Serbia

e-mail jelena.radojicic@ekonomski.rs

Participation (direct/virtual): direct
JEL: G14; G41; C63

THE IMPACT OF THE FEAR AND GREED INDEX ON LIMIT ORDER BOOK MICROSTRUCTURE AND EXECUTION COSTS: AN AGENT-BASED SIMULATION OF CRYPTO MARKETS

Abstract: This paper investigates whether the Fear and Greed Index (FGI) influences cryptocurrency market microstructure and execution quality in an agent-based limit order book environment. The NetLogo model incorporates explicit buy and sell order books, a price-priority matching mechanism, mixed Poisson/Hawkes-like order arrivals, heavy-tailed order sizes, and rule-based cancellations. Sentiment enters the model parsimoniously, with the probability of a buy order set equal to FGI/100. The analysis is based on a single simulation horizon of 50,000 executed trades in event time. The results show that higher FGI regimes generate a pronounced directional shift in order flow, with the buy-order share increasing from 25.64% in the 20–35 FGI bin to 73.05% in the 65–80 bin, while average order-book imbalance shifts from -0.5175 to 0.4424. Quoted spread changes only modestly across sentiment regimes, whereas effective spread rises from 3.0029 to 3.9329, indicating materially poorer execution quality under greed-dominated conditions. The model also reproduces key stylized facts, including fat-tailed return distributions and persistent volatility clustering in mid-price returns. Overall, the findings suggest that sentiment affects market quality primarily through directional order-flow pressure and liquidity depletion rather than through a simple widening of top-of-book quotes.

Keywords: fear & greed index; limit order book; crypto markets; agent-based simulation; execution costs; volatility clustering

1. INTRODUCTION

Cryptocurrency markets provide a particularly useful environment for studying the interaction between investor sentiment and market microstructure (Almeida & Gonçalves, 2023). They operate continuously, are highly order-driven, and often exhibit rapid switches between fear- and greed-dominated trading regimes. In these markets, the limit order book (LOB) is the fundamental mechanism through which sentiment influences quotes, depth, imbalance, and execution costs. While the LOB literature (Hasbrouck, 2007; Gould et al., 2013; Abergel et al., 2016) demonstrates that order book states contain critical information for price formation, the specific link between aggregate sentiment proxies and granular LOB dynamics remains under-explored.

Furthermore, market quality research emphasizes that quoted spreads alone are insufficient metrics, as execution costs often deteriorate due to directional pressure. This study addresses this gap by employing an Agent-Based Modeling (ABM) framework to isolate sentiment as a structural driver of market outcomes. Unlike exogenous models, ABM allows liquidity and volatility to emerge endogenously from order submission and matching rules (Axtell & Farmer, 2022), consistent with the stylized facts of financial time series such as heavy tails and volatility clustering (Cont, 2001; Bacry et al., 2015). By incorporating self-exciting event dynamics (Bacry et al., 2015) and correlated order flow (Bouchaud et al., 2009), the model captures the clustering mechanisms typical of high-frequency activity, ensuring that emerging patterns are consistent with stylized facts such as heavy tails and persistent volatility (Cont, 2001).

The objective of this paper is to examine how the Fear and Greed Index (FGI) affects LOB microstructure and execution costs by integrating a sentiment proxy directly into a controlled NetLogo simulation (Wilensky, 1999). Specifically, the analysis addresses four questions: (1) whether higher FGI values systematically shift the direction of order flow, (2) whether they alter order-book imbalance and quoted spread, (3) whether they increase effective execution costs, and (4) whether the model reproduces stylized facts consistent with realistic market dynamics.
The remainder of the paper is organized as follows. Section 2 outlines the methodology and model design. Section 3 presents the simulation results. Section 4 discusses the findings and their implications. Section 5 concludes.

2. LITERATURE REVIEW

Recent literature suggests that sentiment is a central driver of trading patterns in cryptocurrency markets rather than a peripheral factor. Almeida and Gonçalves (2023) show that crypto-asset investment intentions are strongly shaped by social influence, often generating herding and speculative bubbles during periods of high volatility. In this context, the Fear and Greed Index (FGI) serves as a composite proxy for market sentiment by aggregating momentum, volatility, and social media activity (Kyriazis and Corbet, 2026). Its economic relevance is supported by evidence that sentiment-driven shifts affect the cross-section of cryptocurrency returns (Chowdhury et al., 2024). Although the index may partly reflect past price dynamics, which raises concerns about endogeneity (Bokhari, 2025), it remains useful in agent-based modeling as a regime indicator. Gessie (2026) further shows that Bitcoin returns Granger-cause changes in the FGI, suggesting that sentiment operates within a feedback loop rather than as a purely exogenous force.

Agent-based modeling is particularly well suited to studying such mechanisms in markets populated by heterogeneous agents with bounded rationality (Axtell and Farmer, 2022). Consistent with heuristic-based limit order book models, ABM allows stylized facts such as heavy tails and volatility clustering to emerge endogenously from local interactions rather than from externally imposed shocks (Gould et al., 2013). Fratrič et al. (2022) demonstrate this in a Bitcoin market setting, showing that polarized sentiment can increase fragility and create conditions conducive to manipulation. However, an important gap remains in the literature: sentiment studies typically focus on aggregate outcomes such as returns and volatility, whereas microstructure research emphasizes fine-scale limit order book variables such as quotes, depth, and spreads. This study addresses that gap by introducing a sentiment proxy directly into the order-submission structure of an explicit limit order book simulation, allowing an examination of how sentiment affects directional order flow, liquidity conditions, execution quality, and the emergence of stylized facts.

3. METHOD

3.1. Model design

The simulation was implemented in NetLogo (Wilensky, 1999) as an order-driven market with 100 trader agents, two explicit order books, and a price-priority matching mechanism. Each agent holds cash and inventory and can submit buy or sell limit orders. Orders are represented by trader identifier, limit price, quantity, side, and submission time. This explicit representation follows the central logic of continuous double-auction markets described in the market microstructure literature (Hasbrouck, 2007; Gould et al., 2013).
The model uses event time for order arrival and trade sequencing, while simulation ticks are used for snapshots and interval aggregation. This distinction is methodologically useful because it separates the stochastic clock of market events from the reporting clock used for tables, figures, and rolling diagnostics. In documentation terms, the model is presented in a compact form consistent with the ODD orientation toward transparent, reproducible simulation design (Grimm et al., 2020).
At the structural level, the model belongs to the family of agent-based LOB simulations in which aggregate market outcomes emerge from repeated order submission, cancellation, and matching rather than from an exogenously specified price process. This makes it suitable for examining whether a simple sentiment rule can be translated into measurable changes in spreads, imbalance, volatility and execution costs.

3.2. Order flow, sentiment and cancellation rules

Sentiment is introduced through a simple but transparent rule: the probability that the next order is a buy order equals FGI/100. Thus, if FGI = 25, the probability of a buy order is 0.25; if FGI = 75, the probability is 0.75. This parsimonious specification makes the economic interpretation of results straightforward because sentiment affects the market through the directional composition of order flow rather than through an externally forced price path.
Order arrival is generated by a mixed Poisson/Hawkes-like process. This mechanism ensures that order arrivals are not purely random but exhibit endogenous bursts, mimicking the latency and reaction patterns of algorithmic trading in crypto-markets. The Poisson component provides a baseline of independent arrivals, while the Hawkes-like component is implemented as a self-exciting intensity process where the arrival rate of new orders is temporarily increased by preceding market events. Although the implementation is deliberately simplified and avoids full parameter estimation to

remain parsimonious, this endogenous feedback loop is designed to capture the clustered nature of high-frequency trading activity and the resulting volatility persistence observed in empirical crypto-markets (Hawkes, 1971; Bacry, Mastromatteo, & Muzy, 2015).

Order prices are drawn around reference buy and sell prices, and order sizes can be generated under a heavy-tail mechanism based on a capped log-normal distribution. This choice is consistent with the broader empirical observation that financial and trading-related distributions are often heavy-tailed (Cont, 2001). It is also consistent with order-book research showing that order placement and cancellation rules can generate realistic spread and volatility patterns even in comparatively simple behavioral settings (Mike & Farmer, 2008).

The model also includes several cancellation policies: random agent-level cancellations, cancellations of orders too far from a reference price, cancellations of orders close to the reference price, and age-based removal of stale orders. These rules keep the book dynamic and prevent the accumulation of unrealistic inactive depth. They also provide a practical way to maintain a continuously evolving book without introducing explicit market-maker agents.

3.3. Key metrics and experiment design

The main microstructure metrics are quoted spread (S_t), midpoint (M_t), book depth, and order-book imbalance (It). If a_t denotes the best ask and b_t the best bid, the metrics are defined as follows:

$$S_t = a_t - b_t$$

$$M_t = \frac{a_t + b_t}{2}.$$

$$I_t = \frac{\mathrm{BuyDepth}_t - \mathrm{SellDepth}_t}{\mathrm{BuyDepth}_t + \mathrm{SellDepth}_t}.$$

Execution quality is evaluated primarily using the effective spread, $\mathrm{ES}_t = 2|\mathrm{p}_t - \mathrm{Mid}_t|$ where p_t is transaction price. This metric captures the realized distance between execution and the contemporaneous midpoint more accurately than the quoted spread, which may remain stable despite significant directional pressure (Bessembinder, 2003; Hasbrouck, 2007).

The reported experiment consists of a single 50,000-trade simulation. This horizon was selected to ensure the convergence of stylized facts and to provide a sufficiently large sample size across all FGI bins, thereby minimizing the impact of transient stochastic noise. The analysis draws from multiple output layers, including trade-level, interval-level, and sentiment-bin files, as well as rolling diagnostics for volatility clustering and agent-level performance. The following sections emphasize a compact set of tables and figures that capture baseline dynamics, FGI-dependent microstructure shifts, and evidence of empirical regularities such as fat tails. The core design choices and reporting structure are summarized in Table 1

Table 1: Core model components and experiment settings

Component	Implementation in the simulation
Market structure	Order-driven crypto market with explicit buy and sell order books and price-priority matching
Agents	100 trader agents with cash and inventory positions
Sentiment rule	Probability of a buy order = FGI/100
Arrival process	Mixed Poisson/Hawkes-like event arrivals
Order sizes	Baseline random quantities with optional capped log-normal heavy-tail mechanism
Cancellation rules	Random, far-from-reference, near-reference, and age-based cancellations
Main metrics	Quoted spread, midpoint, imbalance, VWAP, effective spread, rolling stylized facts
Reported experiment	Single run with 50,000 executed trades and multiple CSV output layers

Source: Author's systematization of the simulation design.

3.4. Output structure and validation strategy

The simulation generates several complementary output layers to ensure a granular analysis of market dynamics. Trade-level data records execution price, quantity, midpoint, spread, and sentiment conditions at the millisecond of execution. Interval-level output aggregates these events into discrete time bins, reporting Volume Weighted Average Price (VWAP), average spreads, and order-book imbalance. Additionally, sentiment-bin summaries aggregate data across predefined FGI regimes (e.g., Extreme Fear to Extreme Greed), allowing for a cross-sectional analysis of market quality. This categorical approach is consistent with recent evidence that sentiment effects, particularly those driven by heightened fear, operate non-linearly across different market states and impact broader investor behavior (Kyriazis & Corbet, 2026). Validation follows a dual-track strategy focusing on both microstructural coherence and statistical plausibility. The model should generate economically sensible relationships between sentiment, order-flow direction, imbalance, and execution costs. Statistical plausibility is assessed by verifying whether the simulated data reproduce a meaningful subset of stylized facts, most notably fat tails and volatility clustering. In this context, the application of rolling windows, midpoint diagnostics, and Ljung-Box statistics aligns the simulation with established empirical benchmarks in financial econometrics (Cont, 2001; Engle, 1982).

The mapping between research questions and observables is straightforward. The first question is evaluated through buy-order share across FGI bins. The second is evaluated through quoted spread and order-book imbalance. The third is evaluated through effective spread, which is closer to realized trading conditions than visible quotes alone. The fourth is evaluated through rolling excess kurtosis, return autocorrelation, midpoint-based volatility clustering, and the dispersion of final agent outcomes. Together, these observables allow the paper to move from sentiment input to market-quality output without leaving the order-book environment.

4. RESULTS

4.1. Baseline market behavior

The 50,000-trade simulation generated a total trading volume of 313,269 units, an average transaction price of 103.5479, and an average trade size of 6.2654 units. The discrepancy between the mean and the median trade size (4 units) indicates a right-skewed distribution consistent with the heavy-tail order-size mechanism. At the interval level, the model produced 4,549 reporting bins, with an average of 10.99 trades and 68.87 units of volume per bin. The average Volume Weighted Average Price (VWAP) across bins was 103.4046. The main baseline indicators are summarized in Table 2.

Table 2: Descriptive statistics for the 50,000-trade simulation

Indicator	Value
Number of trades	50,000
Total volume	313,269
Average transaction price	103.5479
Median transaction price	103.9133
Price standard deviation	2.6969
Average trade size	6.2654
Median trade size	4
Average VWAP across bins	103.4046
Average quoted spread at trade time	1.2721
Average end-of-bin spread	0.7735

Source: Author's calculation based on NetLogo simulation outputs

Quoted spreads remained moderate on average. The mean spread measured at trade time was 1.2721, while the mean end-of-bin spread was lower at 0.7735, reflecting the continuous replenishment of the limit order book. Average end-of-bin imbalance was close to zero (-0.0313), suggesting that the simulated market does not suffer from permanent directional bias; instead, it remains dynamically rebalanced through the interplay of order submissions and cancellations. While local sentiment regimes introduce significant temporary asymmetries, the market architecture ensures a return to structural equilibrium over the long term.

The overall price trajectory serves as a primary indicator of the model's structural stability. Figure 1 shows that transaction prices fluctuate within a relatively narrow but persistent corridor around the 100-105 range, with repeated short-lived deviations. This price path is stationary enough to sustain active matching, yet sufficiently variable to produce non-trivial return dynamics. This balance is desirable in an exploratory ABM, as it ensures that the simulation avoids the pitfalls of both explosive instability and trivial price constancy.

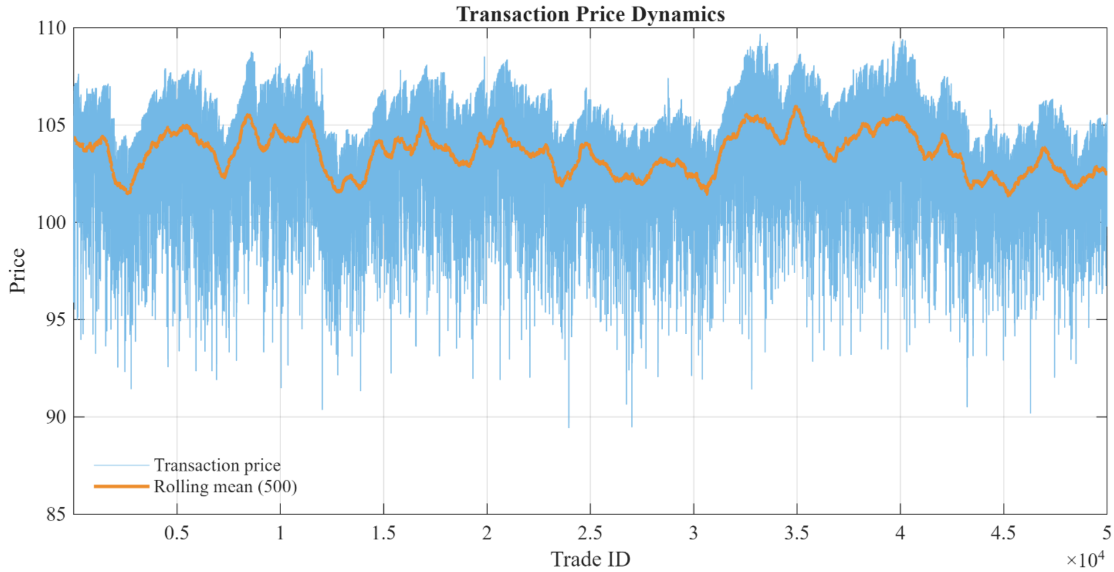

Figure 1: Transaction price dynamics over 50,000 trades
Source: Author's calculation based on NetLogo simulation outputs.

4.2. FGI effects on order flow, imbalance, and execution costs

The strongest result concerns the directional composition of order flow. As the FGI increases from the lowest-sentiment bin (20-35) to the high-sentiment bin (65-80) the share of buy orders rises monotonically from 25.64% to 73.05%. In parallel, average order-book imbalance shifts from -0.5175 to 0.4424. This means that low-FGI regimes are clearly sell-dominated, whereas high-FGI regimes are clearly buy-dominated. The sentiment-bin results are reported in Table 3.
An analysis of liquidity indicators reveals a notable divergence between visible and realized costs.Quoted spread reacts only mildly to sentiment. Mean spread falls from 1.3432 in the lowest FGI bin to 1.2726 in the 50-65 bin and then slightly increases to 1.2824 in the highest bin. By contrast, effective spread rises steadily from 3.0029 to 3.9329. This distinction is theoretically significant because it suggests that extreme sentiment does not merely widen visible top-of-the-book quotes, but instead increases realized execution costs through depth depletion. In economic terms, greed-dominated regimes appear to intensify directional pressure and liquidity depletion more than they mechanically affect the visible spread, highlighting the importance of using effective spread as a proxy for true market friction.

Table 3: FGI effects on microstructure and execution costs

FGI bin	Buy-order share	Trades	Quoted spread	Effective spread	Imbalance
20-35	25.64%	12,950	1.3432	3.0029	-0.5175
35-50	42.49%	11,700	1.3191	3.2789	-0.1673
50-65	57.37%	15,000	1.2726	3.5606	0.1933
65-80	73.05%	10,350	1.2824	3.9329	0.4424

Source: Author's calculation based on NetLogo simulation outputs
Note: Sample sizes vary across bins due to stochastic FGI distribution but remain statistically significant for all regimes (N > 10,000).

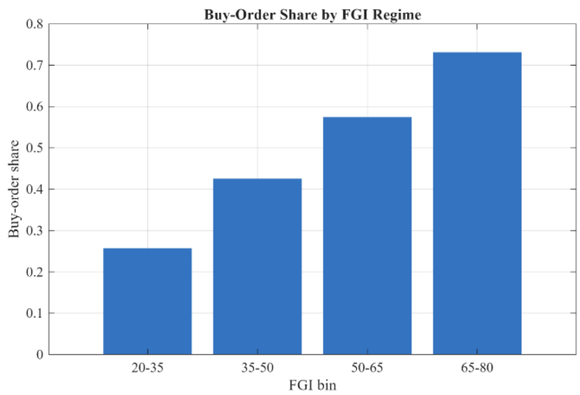

Figure 2: Buy-order share by FGI regime
Source: Author's calculation.

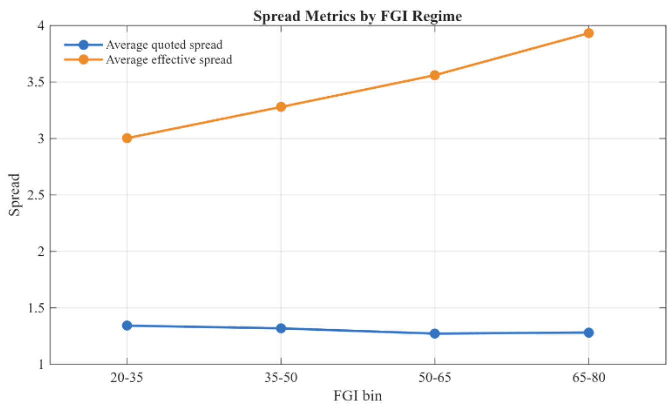

Figure 3: Spread metrics by FGI regime
Source: Author's calculation.

Figures 2 and 3 summarize the main sentiment effect visually. The monotonic increase in buy-order share confirms that the sentiment module transmits directly into order-flow direction. At the same time, the effective spread curve rises much more clearly than the quoted spread curve, reinforcing the conclusion that execution costs become worse under stronger greed even when the top of the book does not deteriorate proportionally. This divergence confirms that intense greed deteriorates execution quality through increased liquidity consumption, meaning visible quoted spreads may mask the true cost of trading during high-sentiment regimes.

4.3. Stylized facts and agent-level outcomes

The simulation successfully reproduces key empirical regularities of cryptocurrency markets. Across rolling windows, the average excess kurtosis of log returns equals 4.2885, indicating fat tails. The average first-lag autocorrelation of raw log returns is -0.3010, which is plausible in an order-book environment affected by microstructure bounce. More importantly, the average first-lag autocorrelation of absolute log returns is 0.1334, showing persistence in volatility-related measures. Mid-price-based diagnostics provide even stronger evidence of volatility clustering. This persistence is visually summarized in Figure 4, which tracks the clustering diagnostics over the simulation horizon.

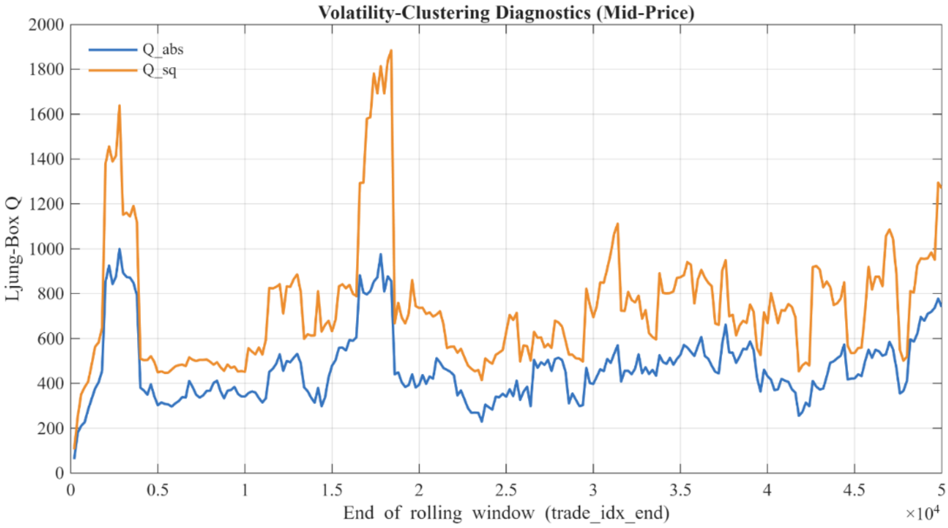

Figure 4: Volatility-clustering diagnostics based on mid-price returns
Source: Author's calculation based on NetLogo simulation outputs.

Across rolling windows, the average first-lag autocorrelation equals 0.4153 for absolute mid-price returns and 0.5375 for squared mid-price returns, while the mean Ljung-Box statistics are 467.89 for absolute returns and 740.79 for squared returns. These values strongly suggest that volatility is serially clustered even when the analysis moves from transaction prices to the cleaner mid-price series.

Beyond aggregate dynamics, agent-level outcomes remain heterogeneous. On the final snapshot, mean PnL is 7.61, median PnL is 13.61, the standard deviation is 1,087.12, and 52% of agents are profitable. The best agent reaches a PnL of 2,763.14, while the worst agent records -2,627.23. Although agent profitability is not the central focus of this paper, these results show that the same microstructure environment generates substantial dispersion in realized opportunities and risks. Table 4 consolidates the most important stylized-facts and agent-level diagnostics.

Table 4: Stylized-facts and agent-outcome diagnostics

Indicator	Value
Excess kurtosis of log returns	4.2885
ACF(1) of log returns	-0.3010
ACF(1) of absolute log returns	0.1334
ACF(1) of absolute mid returns	0.4153
ACF(1) of squared mid returns	0.5375
Ljung-Box Q for |mid returns|	467.89
Ljung-Box Q for squared mid returns	740.79
Mean final PnL	7.61
Median final PnL	13.61
Profitable agents	52%

Source: Author's calculation based on NetLogo simulation outputs.

5. DISCUSSION

The results demonstrate a coherent transmission mechanism where sentiment, captured by the FGI, materially alters market quality through directional order-flow pressure rather than exogenous price paths. This highlights the advantage

of using agent-based modeling (ABM) to capture how micro-behaviors aggregate into market-wide phenomena (Axtell & Farmer, 2022). The monotonic rise in buy-order share and the corresponding shift in imbalance confirm that FGI signals translate directly into one-sided liquidity pressure. A critical finding is the divergence between quoted and effective spreads; while the former remains stable, the latter increases significantly. This suggests a "walking the book" effect, where execution costs rise due to depth depletion, a phenomenon where visible liquidity masks actual deterioration in execution conditions (Hasbrouck, 2007; Bessembinder, 2003).
The main contribution of this paper is empirical rather than architectural. It demonstrates that even a robustly implemented, simple sentiment rule can materially alter order-flow direction, order-book pressure, and effective execution costs within an agent-based limit order book framework. This reinforces the idea that sentiment is a significant factor in shaping market friction in crypto-assets. This finding is consistent with empirical evidence highlighting the role of investor sentiment in shaping trading behavior and liquidity dynamics in cryptocurrency markets (Chowdhury et al., 2024). Furthermore, the observed sensitivity to sentiment regimes (e.g., Extreme Fear) aligns with recent evidence suggesting that cryptocurrency fear acts as a critical information channel for market participants (Kyriazis & Corbet, 2026). While some literature explores if returns themselves drive sentiment (Gessie, 2026), our simulation isolates the reverse path, showing that sentiment alone is a sufficient driver of liquidity friction.
A key implication is that quoted spread and effective spread capture different dimensions of market quality. From a practical perspective, the results suggest that execution strategies relying solely on quoted spreads are insufficient in sentiment-sensitive environments. More informative signals include the joint evolution of order-book imbalance and effective spread. For market participants, this implies that execution risk increases even when quoted liquidity appears stable; for exchange design, it highlights that apparent stability may coexist with hidden inefficiencies under directional stress (Bouchaud et al., 2009).
Furthermore, the model endogenously reproduces key stylized facts, such as fat-tailed returns and volatility clustering. These emerge from the interaction of order arrivals and heavy-tailed order sizes, consistent with established literature where such properties are emergent outcomes of heterogeneous agent interaction (Cont, 2001; Bacry et al., 2015). Importantly, a high FGI should not be viewed solely as a directional predictor, but also as a signal of market quality indicating increased execution costs due to concentrated order flow on one side of the book. This aligns with a systematic understanding of investor behavior in digital assets, where attention and sentiment often dictate market structure more than fundamental value (Almeida & Gonçalves, 2023; Aslanidis et al., 2022).
Table 5 provides a structured overview of the central empirical findings and their implications for understanding the relationship between sentiment and market microstructure.

Table 5: Interpretation of the main empirical findings

Finding block	Interpretation
Order-flow composition	Higher FGI values shift submitted order flow strongly toward the buy side.
LOB state	Sentiment pressure is expressed clearly in imbalance and only moderately in quoted spread.
Execution quality	Effective spread responds more strongly than quoted spread, indicating worse trading conditions under greed-dominated regimes.
Statistical plausibility	Fat tails, volatility clustering, and dispersed agent outcomes suggest that the sentiment effect emerges in a non-trivial market environment.

Source: Author's interpretation of simulation results

The paper nevertheless has several limitations. The implementation relies on price priority rather than full price–time priority (Abergel et al., 2016),, uses a simplified Hawkes-like process rather than a fully calibrated model and excludes explicit exchange fee structures. Accordingly, the simulation should be interpreted as a controlled experimental environment designed to isolate the sentiment transmission mechanism, rather than as a full replication of a specific exchange microstructure.

6. CONCLUSION

This study employs an agent-based simulation to investigate how market sentiment, proxied by the Crypto Fear and Greed Index (FGI), shapes the microstructure and execution costs of cryptocurrency limit order books.
The results show that a parsimonious sentiment rule generates economically meaningful microstructure effects. As the Fear & Greed Index increases, the simulated market becomes progressively more buy-driven, order-book imbalance shifts toward the demand side, and effective execution costs rise, even though quoted spread remains relatively stable.
The central empirical finding is that sentiment affects market quality primarily through directional pressure within the order book, rather than through direct price adjustments. This reinforces the view that sentiment should be incorporated into microstructure-based analyses of cryptocurrency markets, particularly when evaluating execution conditions and trading costs.
The validity of these findings is supported by the model's ability to endogenously reproduce key stylized facts, including fat-tailed returns and volatility clustering. This confirms that the observed sentiment effects are emergent properties of a non-trivial market environment rather than artifacts of a constrained price process. By highlighting that sentiment

materially alters execution quality, this research reinforces the necessity of incorporating behavioral indicators into microstructure-based analyses of digital assets.

Future research may extend this framework by implementing full price–time priority, increasing agent heterogeneity, and conducting broader multi-run sensitivity analyses to further refine the understanding of sentiment transmission in decentralized markets.

REFERENCES

Abergel, F., Chakrabarti, B. K., Chakraborti, A., & Mitra, M. (2016). *Limit Order Books*. Cambridge University Press.

Almeida, J., & Gonçalves, T. C. (2023). A systematic literature review of investor behavior in the cryptocurrency markets. *Journal of Behavioral and Experimental Finance*, 37, 100785. https://doi.org/10.1016/j.jbef.2022.100785

Aslanidis, N., Bariviera, A.F. and López, Ó.G. (2022) 'The link between cryptocurrencies and Google Trends attention', *Finance Research Letters*, 47, 102654. https://doi.org/10.1016/j.frl.2021.102654

Axtell, R. L., & Farmer, J. D. (2022). Agent-Based Modeling in Economics and Finance: Past, Present, and Future. *INET Oxford Working Paper* No. 2022-10. Institute for New Economic Thinking at the Oxford Martin School

Bacry, E., Mastromatteo, I., & Muzy, J. F. (2015). Hawkes processes in finance. Market Microstructure and Liquidity, 1(01), 1550005. https://doi.org/10.1142/S2382626615500057

Bessembinder, H. (2003). Issues in assessing trade execution costs. *Journal of Financial Markets*, 6 (3), 233-257. https://doi.org/10.1016/S1386-4181(02)00064-2

Bokhari, A. (2025). Unpacking Herding in Crypto Markets: An Agent-Based Study of True and Spurious Herding Dynamics', *Procedia Computer Science*, 274, pp. 463–476. https://doi.org/10.1016/j.procs.2025.12.046

Chowdhury, M. I. H., Hasan, M., Bouri, E., & Tang, Y. (2024). Emotional spillovers in the cryptocurrency market. *Journal of Behavioral and Experimental Finance*, 41, 100878. https://doi.org/10.1016/j.jbef.2023.100878

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223-236. https://doi.org/10.1080/713665670

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 987-1007. DOI:10.2307/1912773

Fratrič, P., Sileno, G., Klous, S. et al. Manipulation of the Bitcoin market: an agent-based study. Financ Innov 8, 60 (2022). https://doi.org/10.1186/s40854-022-00364-3

Gessie, L. (2026). Do Bitcoin Returns Move Sentiment? Evidence from the Crypto Fear & Greed Index. *Finance Research Open,* 2 (2026) 100094. https://doi.org/10.1016/j.finr.2026.100094

Gould, M. D., Porter, M. A., Williams, S., McDonald, M., Fenn, D. J., & Howison, S. D. (2013). Limit order books. *Quantitative Finance*, 13(11), 1709–1742. https://doi.org/10.1080/14697688.2013.803148

Grimm, V., Railsback, S. F., Vincenot, C., Berger, U., Gallagher, C., DeAngelis, D., Edmonds, B., et al. (2020). The ODD protocol for describing agent-based and other simulation models: A second update to improve clarity, replication, and structural realism. *Journal of Artificial Societies and Social Simulation*, 23 (2), 7. https://doi.org/10.18564/jasss.4259

Hasbrouck, J. (2007). *Empirical Market Microstructure*. Oxford University Press

Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58 (1), 83-90. https://doi.org/10.2307/2334319

Kyriazis, N., & Corbet, S. (2026). Can cryptocurrency fear influence technology firm investors? *International Review of Economics & Finance*, 107, 105043. https://doi.org/10.1016/j.iref.2026.105043

Mike, S., & Farmer, J. D. (2008). An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control,* 32(1), 200-234. https://doi.org/10.1016/j.jedc.2007.01.025

Wilensky, U. (1999) NetLogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston. http://ccl.northwestern.edu/netlogo/