QUANTITATIVE FINANCE

Rule-Based Trading on an Order-Driven Exchange: A Reassessment

Journal:	*Quantitative Finance*
Manuscript ID	RQUF-2022-0240
Manuscript Category:	Research Article
Date Submitted by the Author:	11-Aug-2022
Complete List of Authors:	Isaac, Alan; American University, Economics Ramaswamy, Vasudeva; American University, Economics
Keywords:	Market Microstructure, Limit Orders, Agent Based Modelling, Bound Rationality, Price Formation
JEL Code:	G12, G17, C63
Abstract:	This paper presents a computational model of an order-driven exchange and reexamines the contribution of trading strategies to market outcomes. Our agent-based model of market microstructure assembles familiar components: near-zero-intelligence (NZI) agents adopt a mix of rule-based strategies and interact indirectly through a limit order book. To lend focus to this exploration, we reconsider the findings an early and well-known exposition of the NZI framework. We show that some previously reported findings are incorrect. We offer a surprising clarification of the contributions of various trading strategies to market behavior, and we suggest improved visualizations to this behavior. These visualizations shed new light on the volume-volatility relationship in this class of models.

Rule-Based Trading on an Order-Driven Exchange: A Reassessment*

Alan G. Isaac
Department of Economics
American University

Vasudeva Ramaswamy
Department of Economics
American University

August 11, 2022

Abstract

This paper presents a computational model of an order-driven exchange and reexamines the contribution of trading strategies to market outcomes. Our agent-based model of market microstructure assembles familiar components: near-zero-intelligence (NZI) agents adopt a mix of rule-based strategies and interact indirectly through a limit order book. To lend focus to this exploration, we reconsider the findings an early and well-known exposition of the NZI framework. We show that some previously reported findings are incorrect. We offer a surprising clarification of the contributions of various trading strategies to market behavior, and we suggest improved visualizations to this behavior. These visualizations shed new light on the volume-volatility relationship in this class of models.

Keywords: market microstructure, order book, trade volume, return volatility, NZI agents, price shading

JEL: G12, G17, C63

*Equal authorship; authors are in alphabetical order. We thank Blake LeBaron for useful discussions.

1 Introduction

An important strand of the computational finance literature explores exchange-mediated interactions between traders with diverse strategies. The work of Gode and Sunder (1993) on zero-intelligence (ZI) traders kick-started research in this area by demonstrating that trading institutions can promote market efficiency, even in the absence of trader rationality. That result focused substantial research attention on the implications of market microstructure for financial-market outcomes.

Much of the subsequent research in agent-based computational finance incorporates near-zero-intelligence (NZI) trading strategies, which are simple rules that embody stylized representations of trading behaviors observed in securities markets. This is a relatively simple and tractable approach to characterizing the substantial variety in observed trader behaviors. In an early guide to agent-based models in computational finance, LeBaron (2000) argues that simple trading behaviors aid understanding, since complex models can obscure the mechanism through which observed outcomes are produced. Nevertheless, achieving an adequate understanding of the market implications of even simple trading behaviors remains challenging.

This paper elucidates some of these challenges by reporting on an attempted replication of the model developed in Chiarella and Iori (2002). Unexpectedly, implementing the CI model proves to be less than straightforward. This is surprising because the CI model is a paradigm of the NZI research strategy, with hundreds of citations. However, many important implementation details are missing or unclear in its original exposition. We explicate some replication problems and show how to achieve a near replication. We discover and explain a surprising result about the contribution of trading behaviors, and provide a correction to some previously reported results. We present some new results and suggest additional, more useful representations of market outcomes.

The paper proceeds as follows. Section 2 provides an overview of order-driven exchanges, and then describes the market microstructure in the CI model. We fill a long-standing gap in the literature by providing a painstakingly detailed exposition of this model. Section 3 describes the market behaviors implied by this model. We verify some of the findings of Chiarella and Iori (2002), correct others, and offer some new findings. Finally, section 4 concludes.

2 Background and Model Description

This section provides basic background on how the agent-based modeling and simulation literature has approached order-driven exchanges. It then describes our computational model of an order-driven exchange for a financial security.

2.1 Order-Driven Exchanges

An order-driven securities exchange brokers trader orders subject to institutional rules. These rules typically include price-time priority and often include a fixed minimum price change (the tick size). An order-driven exchange derives liquidity from limit orders, which are passive commitments to buy or sell specified quantities at specified prices (Gould *et al.*, 2013; Bouchaud, 2018). At any point in time, the order book of the exchange contains the active limit orders.

Briefly, the core attributes of a limit order are its price, quantity, and expiry. An order may be a bid, specifying the maximum unit price to pay for the demanded quantity, or an ask, which specifies the minimum unit price for the offered quantity. The bid side of the book is a collection of unexpired bids, and the ask side of the book is a collection of unexpired asks. The current best bid is the highest bid price currently in the book, and the current best ask is the lowest asking price currently in the book. The best bid and best ask are the top of the book, and the distance between them is the bid-ask spread.

An order-driven system automatically produces an order execution whenever a new order (bid or ask) hits a quote on the opposite side (ask or bid). It follows that for the limit-order book of a properly functioning exchange, the current best bid must be below the current best ask. An unmatched limit order remains active in the order book until it expires. If an order expires before it can be executed, it is removed from the book. However, a new marketable order will be executed at the price currently at the top of the other side of the book. For example, an arriving bid may be filled at the current best ask. (If there is more than one ask at that price, the earlier order has priority.)

As Glosten (1994) famously predicted, electronic open limit-order books have come to dominate securities exchanges. As this dominance grew, research attention increased accordingly. Since the mediated interactions between individual traders on the exchange determine the volatility of the market price and the volume of trade, trader behavior is a focus of the related empirical, theoretical, and computational research.

Research on the contribution of trader strategies to market outcomes has a long lineage. Zeeman (1973, 1974) suggested the chartist-fundamentalist model of financial markets as an application of catastrophe theory. Early mathematical models include Beja and Goldman (1980), Frankel and Froot (1987), and Bowden (1990). Chiarella (1992) added 'noise' effects to the set-up outlined in Beja and Goldman (1980). Agent-based implementations appear in the early artificial stock markets (ASMs), including the Santa Fe Institute ASM (Palmer *et al.*, 1994) and the Genoa ASM (Raberto *et al.*, 2001). The work of Lux and Marchesi (1998) encouraged a transition to small scale models.

Three decades ago, Gode and Sunder (1993) highlighted the importance of the institutional features of securities markets by finding that good market efficiency occurred even with so-called zero-intelligence (ZI) agents. However, researchers found that in models of order-driven exchanges, ZI agents do not seem adequate to the set of financial-market stylized facts. They also exhibit counterfactual degrees of volatility and other undesirable behaviors (Chakraborti *et al.*, 2011; Maslov, 2000). This inadequacy led to increased interest in NZI traders, who respond to current market conditions, including perceived uncertainty in the execution price. Varied approaches are now common in the agent-based behavioral finance literature, and traders often implement a mixture of chartist, fundamentalist, and noise trader approaches.

Two decades ago, Chiarella and Iori (2002) provided a famous demonstration of the explanatory potential of varied NZI traders with mixed strategies in an agent-based model of an order-driven exchange. Most importantly, they exposed implications of market participation, order duration, and market tick size for the behavior of asset prices and market returns. The CI model remains an exemplar in agent-based computational finance of tying market outcomes to trader strategies on an order-driven exchange. Correspondingly, their model is a focus of the present paper.

2.2 Limit Orders and Limit-Order Books

Our computational model encompasses the CI model and easily allows parameterized extensions to it. At the core of this model is a dynamic limit-order book. It stores two collections of active (unexpired, unfilled, and uncanceled) limit orders that correspond to the two basic transaction types: bid orders and ask orders. Traders in the model generate limit orders, and an unfilled order passively remains in the order book until it is filled or it expires. Orders have the typical features of a limit order on a trading exchange. An order of either transaction type specifies a limit price, an order size, and an expiry. The order is also assigned a time of receipt, which determines time priority. Figure 1 illustrates these key features.

LimitOrder
transactionType: String {bid | ask} price: Price quantity: Quantity datetime: DateTime expiry: DateTime traderID: Integer
fill(:Price, :Quantity, :DateTime)

Figure 1: Features of a Limit Order
A `Price` is a positive real number. A `Quantity` type is a nonnegative integer.

Figure 2 displays an activity diagram for the core simulation schedule. Since the market is an order-driven exchange for a financial asset, activity commences when a trader places an order with the exchange. A matching engine immediately compares the order price to the top of the other side of the order book. If the prices cross (e.g., if a new bid exceeds the best ask), the trade executes (at the prior price) and the order book updates. Otherwise, the new order is entered on the order book, awaiting possible future execution (or expiration). This is standard behavior for an order book.

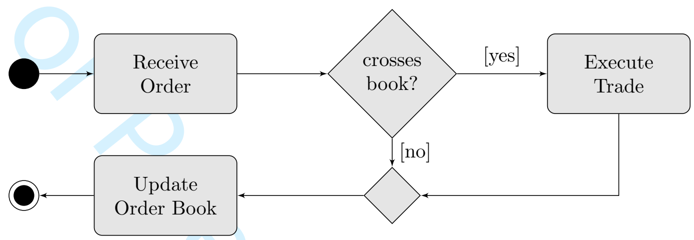

Figure 2: Core of Simulation Schedule

Order matching and trade execution reflect the usual price-time priority. Orders with better prices always have a higher priority. Among orders with the same limit price, the oldest order has priority.[1] If execution of a new order is possible, it happens immediately. The matching engine immediately removes executed orders from the order book.

To illustrate in more detail, consider the handling of a new bid. The new bid is received by the market. If it does not beat the existing best bid, it simply joins the order book. However, if it beats the book's best bid, it is compared to the best ask. If it crosses the book, a trade executes and the order book updates accordingly. New asks are handled symmetrically.

[1] If all floating point prices were acceptable, the market would not see common prices across orders. However, a market in a security typically has a tick size, which is the minimal price increment. In addition, order prices outside an acceptable (wide) trading range are typically rejected. Chiarella and Iori (2002, p.348) address this by introducing a pre-specified grid of possible prices, based on the tick size (Δ). (Unfortunately, they do not document the minimum and maximum values of this grid.) We follow this practice, specifying a (wide) range of possible prices, from 1% to 200% of the reference fundamental price.

2.3 Trader Behavior

Any agent-based model of a securities exchange must characterize the behavior of participating traders. This section briefly reviews trader behavior in the CI model.

2.3.1 Order Construction

Each time step, a random trader may choose to submit a limit order, as illustrated in Figure 3. An active trader constructs a target price, which reflects the trader's pricing strategy and trading horizon. The decision to place a bid or an ask rests on whether the trader's target price (p^*) is above or below the observed price (p^o). In the CI model, the order is always a limit order, with a specified price and duration. (In the CI model, the quantity is fixed at one unit of the security.) Finally, the trader submits the constructed order to the exchange, thereby initiating the order processing events already illustrated by Figure 2.

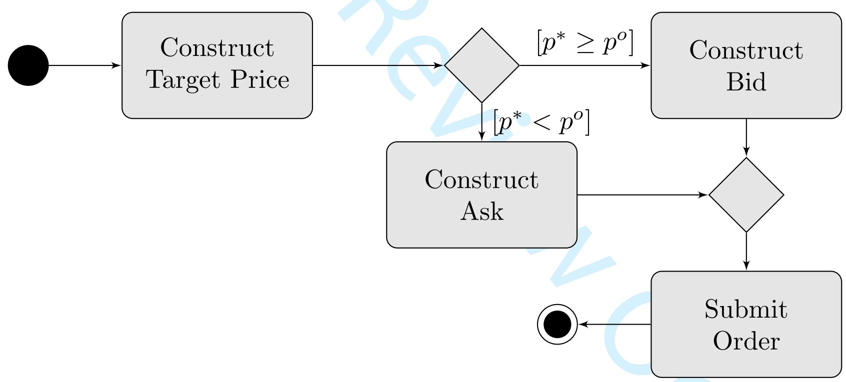

Figure 3: Overview of Trader Activity
The trader's target price is p^*; the observed price is p^o.

2.3.2 Pricing Strategy

Each trader has a trading horizon of τ time periods. (This is a parameter of the CI model; see section 3 for details.) Each trader's strategy for constructing a target price may respond to up to three components: a contemporaneous market-related shock or idiosyncratic influence (ϵ), the deviation of the observed price from the fundamental value ($p^f/p^o - 1$), and the average spot returns over a recent past ($\bar{r}^L$). A trader projects average spot returns over the τ-period horizon by weighting these components.

$$r_t^* = g_0 \epsilon_t + g_1 \frac{p^f - p^o}{p^o} + g_2 \bar{r}_t^L \tag{1}$$

The weights (g_0, g_1, and g_2) are trader specific, reflecting the diverse strategies taken by different traders. (See section 3 for details.) Projected returns (r^*) and the observed price (p^o) imply a target price (p^*), as follows.

$$p_t^* = p_t^o e^{r_t^* \tau} \tag{2}$$

The weights (g_0, g_1, and g_2) placed on each component lie at the core of a trader's strategy. A *noise trader* weights only the market shock. The literature typically treats this as a simplest, zero-intelligence (ZI) strategy. Traders who weight the other components are considered to be near-zero-intelligence (NZI) traders, because they use simple rule-based strategies. A *fundamentalist* responds only to deviations from the perceived fundamental price. A *momentum trader* responds only to the recent behavior of the market price.

Strategy
pf: Real weights: Real[3] L: Integer
pTarget(shock: Real, priceHistory: Real[*]): Real

Figure 4: Features of a Trader Strategy
Fixed attributes of a strategy are the trader's perceived fundamental price, the weights on the return components, and the length of the relevant price history. The target price responds to a contemporaneous shock and the recent price history.

As illustrated by Figure 4, we find it useful to bundle these considerations into a trader-specific strategy. This includes the trader's perceived fundamental price (p^f), the sub-strategy weights (on noise, deviation from fundamental, and past momentum), and the trader's historical time horizon (L). In the CI model, a history of trade prices remains visible to all traders, and computation of the momentum component is based on the asset's price history. Price momentum is captured by the average of the spot returns over over a recent history of market returns.

$$\bar{r}^L = \frac{1}{L}\sum_{j=1}^{L}\frac{p_{t-j} - p_{t-j-1}}{p_{t-j-1}} \tag{3}$$

Price Shading Recall from Figure 3 that the relationship between the target price and the current price determines whether the agent submits a bid or ask. A trader's pricing strategy in the CI model has a second component: traders engage in price shading. A trader's order price implies an anticipated profit based on the trader's target price.

The usual intuition for price shading draws on the bid-shading strategies developed for avoiding the winner's curse in first-price auctions (Riley and Samuelson, 1981; Lillo, 2007; Zhan and Friedman, 2007). In the CI model, NZI traders manifest risk aversion as a rule-based tendency to shade prices. A trader's bid (p^b) is a bit less than the trader's target price, making it less likely that a trader bids more than necessary to acquire the security. Symmetrically, a trader's ask (p^a) is a bit more than the trader's target price, making it less likely that a trader sells at a lower price than necessary to sell the security. The amount of shading is trader specific. Equation (4) characterizes shading for trader with target price p^* and shading parameter $k \in [0.0, 1.0]$.

$$\begin{aligned} p_t^b &= p_t^*(1-k) \quad \text{if } p^* \geq p^o \\ p_t^a &= p_t^*(1+k) \quad \text{if } p^* < p^o \end{aligned} \tag{4}$$

3 Parameterization and Results

This section extracts a baseline parameterization for our computational model from Chiarella and Iori (2002). Therefore, in principle, the baseline simulation constitutes an attempted replication of their results. As it turns out, we reach some novel conclusions.

3.1 Baseline Parameterization and Initialization

Table 1 lists key model parameters and their baseline values. In the CI model, agents agree about the fundamental price of the asset (p^f). The market participation rate (λ) is the probability that, each time step, a randomly chosen trader will choose to place an order. An order's time in force (τ) contributes to the price projection in (2).[2] The exchange for this security imposes a fixed tick size (Δ).

Table 1: Key Baseline Parameters

Name	Type	Value	Description
p^f	Real $\{> 0\}$	1000.00	perceived fundamental price
λ	Real $\{\text{in } [0\,..\,1]\}$	0.50	market participation rate
τ	Real $\{> 0\}$	2.00	order duration (in periods)
Δ	Real $\{> 0\}$	0.01	tick size

Data source: Chiarella and Iori (2002). A supplement to this paper documents all model parameters.

Strategy weights are trader-specific, and the CI model parameterizes their distributions. For each trader i, weights are initialized as follows, where N indicates a normal distribution (as a function of the mean and standard deviation).

$$g_{0_i} \sim N[0, \sigma_0] \qquad g_{1_i} \sim |N[0, \sigma_1]| \qquad g_{2_i} \sim N[0, \sigma_2] \tag{5}$$

In the baseline parameterization, $\sigma_0 = 3.0$, $\sigma_1 = 1.0$, and $\sigma_2 = 1.4$. The time-step-specific market shock (ϵ) is also drawn from a mean-zero normal distribution, with $\sigma_e = 1.0$ in the baseline parameterization.

$$\epsilon_t \sim N[0, \sigma_e] \tag{6}$$

[2]We talk about time in two different ways, which we call time *steps* and time *periods*. A time step is just a single simulation step. There are 100 time steps in a time period. Baseline order duration is $\tau = 2$ time periods (i.e., 200 time steps). We elaborate on this in the results subsection.

Additionally, each trader has a trader-specific shading rate, $k_i \sim U[0, k_{\max}]$, where the baseline value is $k_{\max} = 0.5$. Finally, each trader has an idiosyncratic horizon for past prices. Each agent i has a relevant price history L_i that is drawn from the discrete uniform distribution on the integer interval $[1 \mathinner{..} L_{\max}]$. In the baseline parameterization, $L_{\max} = 100$ time steps (i.e., one time period).

3.1.1 Price History

Momentum traders in the CI model make order-pricing decisions based on historical price information. This history specifies a price at each time step, whether or not a trade is executed (Chiarella and Iori, 2002, p.348). Figure 5 illustrates the construction of this price history. If a trade occurs, this is the trade price. If not, the price is a function of the state of the order book. If there is a bid and ask in the book, then the price is the midpoint price (i.e., the mean of the best bid and the best ask). If the book has zero depth on either side, then the price is the best bid or the best ask, whichever exists. If the book has zero depth on both sides, then the price is considered unchanged from the previous time step.[3]

3.2 Baseline Simulation

This section discusses the simulation of the model under the baseline parameterization. In principle, under this parameterization, our computational model should readily replicate the findings of Chiarella and Iori (2002) by producing qualitatively similar financial-market features. In practice, challenges arise and some important differences emerge.

3.2.1 Replication Challenges

Attempts to replicate the findings of Chiarella and Iori (2002) stumble on the absence of key simulation details.[4] In order to achieve a near replication, this paper adds some clarifications

[3]Chiarella and Iori (2002) do not specify an initialization of the price history, which momentum traders require when the model starts. This paper uses the fundamental price to initialize the price history. Results with a random initial price history are similar, so the model appears robust to this choice.

[4]The original code is unavailable (G. Iori, personal communication, 2020).

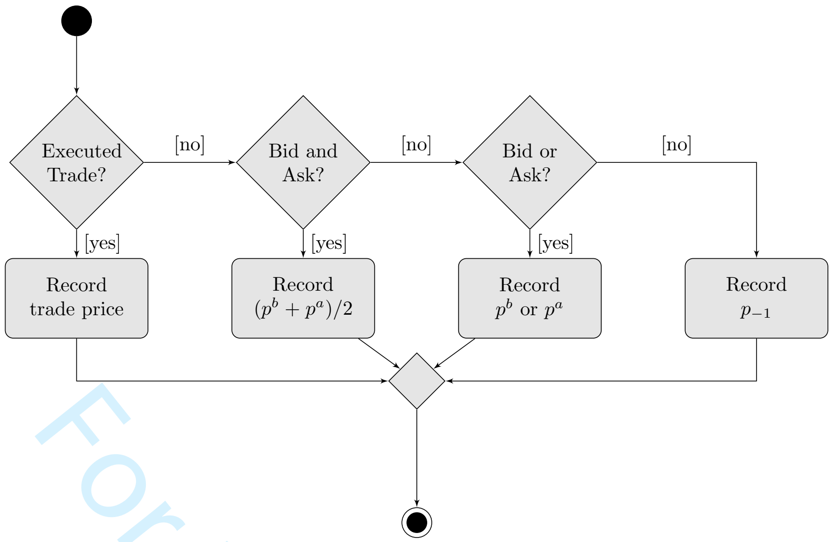

Figure 5: Price History Construction
The best bid is p^b; the best ask is p^a; the last recorded price is p_{-1}.

and one reinterpretation of the baseline parameterization.

Most importantly, our explorations lead us to conclude that CI's implementation in code of their paper's model must have contained a time-accounting error. This proved very hard to sort out, partly due to a terminological inconsistency. The original paper uses the terms *step*, *period*, *time unit*, and *daily*. The term *daily* is a synonym for *per period*. It is clear that a period is synonymous with a time unit and that a period is 100 time steps (i.e., iterations of the core simulation schedule). However, CI report τ values in two different units: time periods (e.g., in their figure 4) and time steps (e.g., in the text on the same page).

This creates an ambiguity in their discussion of the size of the time horizon (τ) appearing in (2). They define $T = 100$ time steps, and $\tau = 2T$ (p.350). However, using a value of $\tau = 200$ in (2) makes replication impossible; the target price fluctuations are far too large. Since their figure 4 shows $\tau \in [0 \mathinner{..} 3]$, we explored using these smaller values in (2). This offered some improvement, but the target-price fluctuations remained far too large. After much exploration, we concluded that they must have accidentally deflated τ twice: scaling τ by 1/100 as would be required to go from 200 (as in their text) to 2 (as in their Figure 4), but failing to notice that τ had already been parameterized as 2 (as in the figure) instead of

200. Deflating twice finally allowed us to approximate many of their reported results.[5]

The remaining barriers to replication were less important. We had to deduce the duration of the simulation runs from the `Time` axis on the CI charts. The units on this axis correspond to time periods of 100 time steps each, where each time step is one iteration of the core model schedule (Chiarella and Iori, 2002, p.350). Since the CI charts encompass $1,000$ periods, each simulation runs for $100,000$ time steps.

As a final clarification, we explicitly characterize multiple concepts of return volatility. To begin, Chiarella and Iori (2002, Figure 2) distinguishes between the *time-step* and the *time-period* volatility of returns. Their formula for time-period volatility is (7), where t varies over the time steps of a single period of length T.

$$\sigma^T = \frac{1}{T} \sum_{t=1}^{T} \left| \frac{p_t - p_{t-1}}{p_{t-1}} \right| \sqrt{T} \tag{7}$$

There is no separate algebraic formula for time-step volatility, and CI's volatility plots are unfortunately 'rescaled and shifted.' To comport with (7), we compute time-step volatility as the absolute one-step price returns.

This does not quite close out the matter of volatility computation, because CI's final figures plot a volatility measure that produces only one observation per simulation run. For this measure, we use (8), so that each observation represents the average of the time-step

[5]Scaling τ in (2) is equivalent to scaling the weights in (1), which is our implementation.

absolute returns over the entire simulation.[6]

$$\frac{1}{NT}\sum_{t=1}^{NT}\left|\frac{p_t}{p_{t-1}}-1\right| = \frac{1}{NT}\sum_{n=1}^{N}\sum_{i=1}^{T}\left|\frac{p_{(n-1)*T+i}}{p_{(n-1)*T+i-1}}-1\right| \tag{8}$$

With these corrections and clarifications in hand, we are ready to produce the simulation results below. The remainder of this section attempts to replicate and elaborate on the market features illustrated in Chiarella and Iori (2002, Figures 1 and 2). These figures descriptively summarize a mix of stock and flow variables but focus on the behavior of prices and returns in this order-driven market. We extend the discussion to cover some descriptive measures that we find more helpful for understanding the market behavior.

3.2.2 Behavior of Market Prices

Recalling the weight distributions in (5), Figure 6 varies σ_1 and σ_2 systematically. From top to bottom, σ_1 changes from 0.0, to 1.0, and then to 5.0. From left to right, the σ_2 values are 0.0 and 1.4.[7] The market price may be measured at each time step, but Figure 6 requires prices only periodically. We use end-of-period prices. This figure demonstrates that our implementation and baseline parameterization can produce price behavior similar to Chiarella and Iori (2002, Figure 1). The associated return series, discussed in the next

[6] We base this choice on Chiarella and Iori (2002, p.351–352), which describes the average spot volatility as being in the neighborhood of 2×10^{-4}. This appears to be roughly the value of volatility indicated at around $\lambda = 0.5$ in Chiarella and Iori (2002, Figure 3). However, there are two other imaginable measures for a simulation of N periods of T time steps per period. One is a simple average of CI's time-period volatilities, as described by 7. That produces

$$\frac{1}{N}\sum_{n=1}^{N}\left(\frac{1}{\sqrt{T}}\sum_{i=1}^{T}\left|\frac{p_{(n-1)*T+i}}{p_{(n-1)*T+i-1}}-1\right|\right) = \frac{1}{N\sqrt{T}}\sum_{n=1}^{N}\sum_{i=1}^{T}\left|\frac{p_{(n-1)*T+i}}{p_{(n-1)*T+i-1}}-1\right|$$

The other is to apply (7) directly to the entire price trajectory.

$$\frac{1}{\sqrt{NT}}\sum_{t=1}^{NT}\left|\frac{p_t}{p_{t-1}}-1\right| = \frac{1}{\sqrt{NT}}\sum_{n=1}^{N}\sum_{i=1}^{T}\left|\frac{p_{(n-1)*T+i}}{p_{(n-1)*T+i-1}}-1\right|$$

Clearly the only difference between these measures is scaling.

[7] This means that the first, fifth, and fourth subfigures correspond to the left column of Chiarella and Iori (2002, Figure 1).

subsection, is also similar.

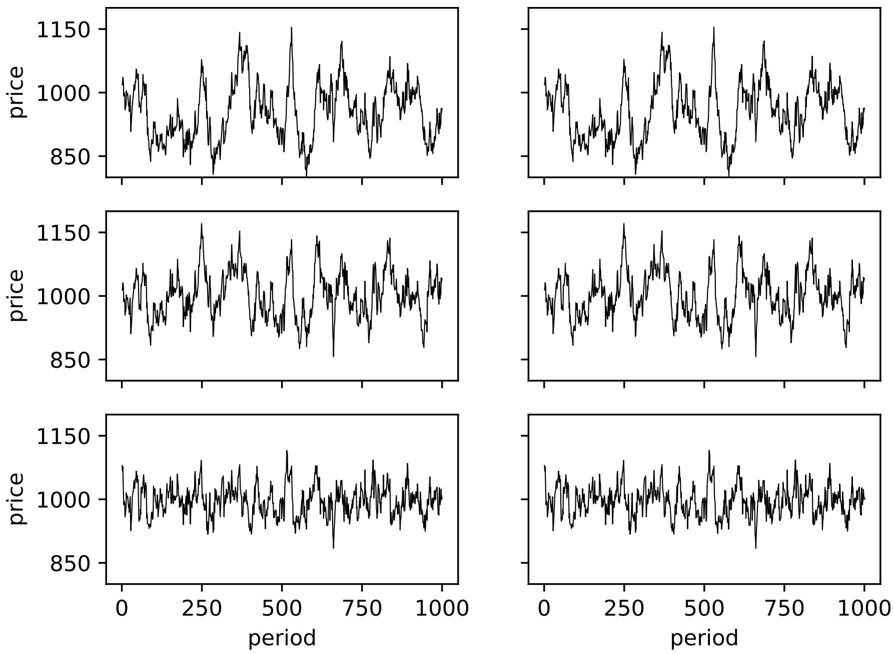

Figure 6: Market Price over Time
Top to bottom, rows correspond to σ_1 equal to 0.0, 1.0, and 5.0. Left to right, columns correspond to σ_2 equal to 0.0 and 1.4. Baseline parameterization, except for σ_1 and σ_2.

However Figure 6 need not be representative of the behavior of the market. Each subfigure represents a trajectory from a single run of the model, and as shown in section 3.3, behavior can vary across runs. In particular, when the market has only noise traders (top-left subfigure), the time series outcomes bear a close resemblance to a random walk.[8]

Figure 6 also exposes a surprising discovery about the role of momentum trading in the CI model. To facilitate comparisons across subfigures, we initialize the random number generator for every simulation run with a common seed. Comparison of the left and right columns shows that the inclusion of momentum trading is essentially irrelevant to the behavior of prices. That is, in this figure it is evident that the price series when traders eschew any momentum strategy ($\sigma_2 = 0.0$) is nearly identical to the result with baseline weights ($\sigma_2 = 1.4$).

[8]More precisely, aside from artefacts of the series construction, the logarithm of the market price should resemble a random walk in this case.

Contrary to the claims of Chiarella and Iori (2002), it appears that momentum traders who react to an average of past returns cannot effectively introduce a role for chartism in this model.

To explore this surprising result, recall that a trader's target price responds to three components: noise (ϵ_t), price gap ($p^f/p^o - 1$), and return momentum ($\bar{r}_t^L$). By design, noise has a zero mean and a unit standard deviation. While momentum is also centered on zero, its standard deviation is orders of magnitude smaller (roughly 0.0003). So it is always very close to zero.[9] As illustrated by Figure 6, the absolute price gap is highly autocorrelated and can be quite large—orders of magnitude greater than the average noise contribution in any single period. Due to the substantial autocorrelation in the price gap, however, its overall mean value provides little insight into the tendency of this magnitude.

3.2.3 Return Volatility and Persistence

Chiarella and Iori (2002, Figure 1, column 2) also plot a time series of periodic market returns, computed as the percentage change in the end-of-period price. We do not find these plots to be particularly informative, but Chiarella and Iori (2002, p.351) perceive volatility clustering. The first subfigure in Figure 7 is our corresponding return plot for the baseline simulation.[10] This provides similar, very slight visual hints of volatility clustering. In addition, the series appears to imply somewhat fat tails in the return distribution.

To explore this further, consider some additional descriptive results. The second subfigure of Figure 7 provides direct evidence concerning the return distribution via a kernel density line plot of the return distribution (the black line). This exposes fat tails when compared with a fitted normal distribution (the gray background). The third subfigure plots the autocorrelation function for returns and absolute returns, with the 95% confidence in-

[9]For this reason, altering the weighting schema to emphasize chartism fails, even when imposing an unreasonably large relative weight on chartism. The supplement to this paper provides additional exploration of these issues.

[10]Attentive readers will note the scale change in the last subfigure of Chiarella and Iori (2002, Figure 1). See the supplement to this paper for further explorations.

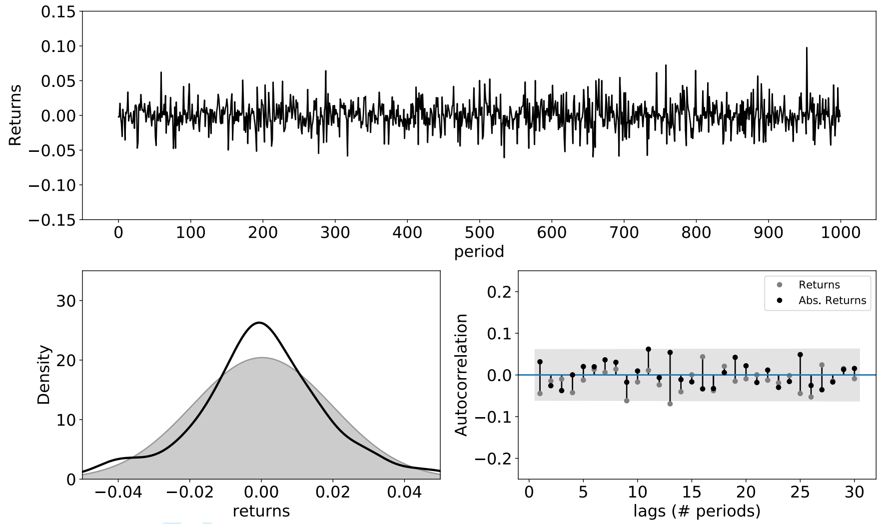

Figure 7: Market Returns
The first panel plots a time-series of period returns. The second subfigure displays a kernel density plot of returns, which has fat tails. (Gray area demarcates a fitted normal distribution.) The third figure shows little autocorrelation in returns or absolute returns. (Gray area demarcates the 95% confidence band.) Baseline parameterization.

terval (calculated with Bartlett's formula) as a shaded region. We find little evidence of autocorrelation in either, undermining CI's claim of volatility clustering in periodic returns.

In order to provide a rough point of comparison, Figure 8 displays the same information for the daily returns of an arbitrarily chosen security. The top panel plots the daily returns of Apple common stock for a thousand days.[11] Some volatility clustering is visually evident. The second panel (bottom left) shows the kernel density plot of returns (the black line), indicating a high degree of kurtosis relative to a fitted normal distribution (the gray background). The third panel (bottom right) plots the autocorrelation function for both daily returns and absolute daily returns, with the 95% confidence interval (calculated with Bartlett's formula) as a shaded region. The autocorrelation of returns is not significant at beyond one lag, while absolute returns remain significantly autocorrelated for more than 40 days. This figure displays slow-decaying autocorrelation in absolute returns, evincing an empirical stylized fact of stock prices (Cont, 2001). This combined outcomes for returns and

[11]The returns are calculated on the closing prices of Apple Inc. stock (ticker: AAPL) over 1,001 trading days between May 4th, 2016 to April 24, 2020.

absolute returns is a typical empirical manifestation of volatility clustering (Cont, 2007). Inability to produce such stylized volatility clustering is one clear shortcoming of the CI model.

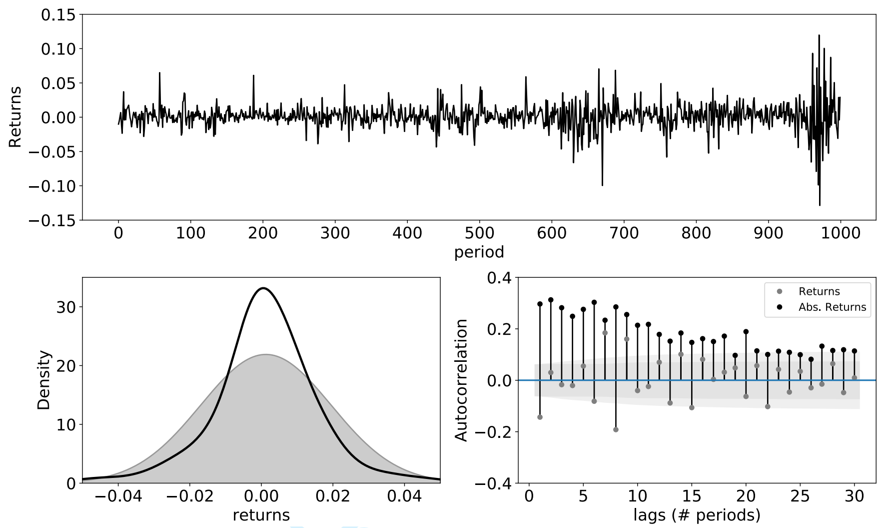

Figure 8: Daily Returns on AAPL
The first subfigure plots a time-series of daily returns (based on closing prices). The second subfigure displays a kernel density plot (black line), which confirms the apparent fat-tailed returns. (Gray area demarcates a fitted normal distribution.) The third figure shows little autocorrelation in returns but substantial, slowly declining autocorrelation in absolute returns. (Gray area demarcates the 95% confidence band.) Baseline parameterization.

Table 2 provides descriptive statistics supporting the above findings. The simulated distribution of returns displays substantial kurtosis, supporting the finding of fat-tails evident in Figure 7. A Jarque and Bera (1980) test correspondingly rejects the null hypothesis of normally distributed returns. In line with Figure 7, a Ljung and Box (1978) test fails to reject the absence of linear autocorrelations, and the McLeod and Li (1983) test fails to reject the absence of non-linear autocorrelations.[12] For ready comparison, the table also reports the statistics corresponding to Figure 8.

[12]We use a lag length of 10 periods for both tests. However, we also deployed the test using the optimal lag selection methodology prescribed in Escanciano and Lobato (2009). The optimal lag for the Ljung-Box test for the simulated returns was 1, and for AAPL stock was 30. The optimal lag for the McLeod-Li test for the simulated returns was 6, and for AAPL stock was 29. The results were qualitatively the same in all cases as those reported in Table 2.

Table 2: Compare Simulated Returns to Apple Stock

	Baseline	**AAPL**
Variance	3.8×10^{-4}	3.3×10^{-4}
Kurtosis	4.45	12.13
Jarque-Bera (1980)	91.43	$3,476.39$
	(0.00)	(0.00)
Ljung-Box (1978)	8.83	131.06
	(0.55)	(0.00)
McLeod-Li (1983)	4.60	750.36
	(0.92)	(0.00)

P-values are in parentheses.

Trade Volume and Return Volatility Periodic trade volume is the sum of all trades that execute during the period. The distribution of periodic trade volume in this model is roughly Poisson. Figure 9 illustrates this by plotting the proportion of periods producing each realized level of trade volume in the baseline simulation. For reference, this figure adds a line plot linking the outcomes predicted by a fitted Poisson distribution.

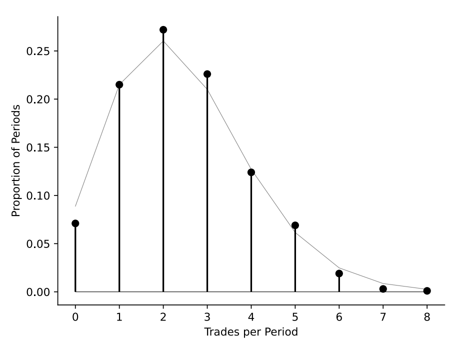

Figure 9: Distribution of Periodic Trade Volume
Black markers designate simulation outcomes. A gray line joins fitted Poisson values. Baseline parameterization.

Figure 10 explores the relationship between periodic trade volume and return volatility. (Equation (7) gives the volatility formula.) This figure's left-side panels mirror Chiarella and Iori (2002, Figure 2, left panel). For 200 periods of the baseline simulation, it illustrates the relationship between the time series of trade volume and return volatility. This particular window on our baseline simulation is only mildly suggestive of a relation between the two

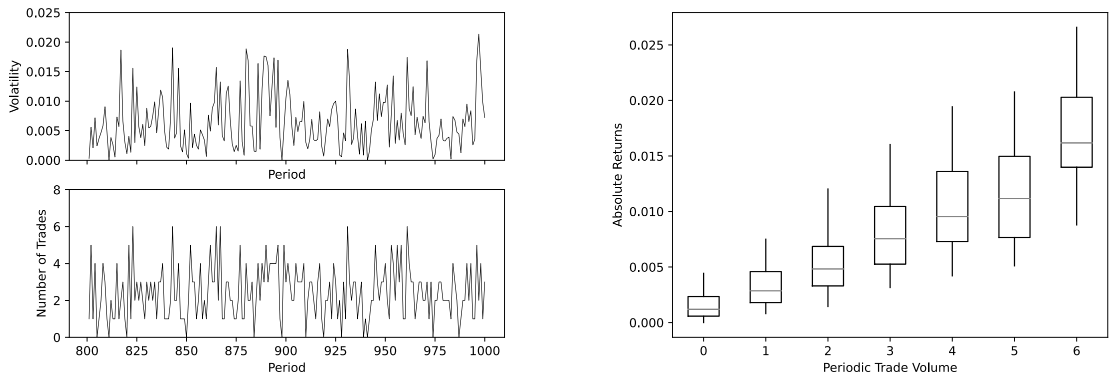

(a) Volatility and Volume (b) Absolute Returns and Trade Volume

Figure 10: Volatility and volume in the baseline simulation.
On the left, find periodic volatility (top) and periodic volume (bottom) for the last 200 periods of the baseline simulation. The right panel provides, for each level of trade volume, a box plot of periodic returns with 90% whiskers.

series.

However, the right panel provides a more informative window on the relationship between trade volume and return volatility. Using data from the entire baseline simulation, for each observed level of trade volume, it provides a box plot of period returns with 90% whiskers.[13] This figure lends new support to CI's claim that periodic volatility is correlated with periodic trade volume (Chiarella and Iori, 2002, p.351). Absolute returns tend to be higher in higher volume periods.

Bid-Ask Spread and Return Volatility Figure 11 zooms in further to illustrate the step-level behavior of returns and spreads. To facilitate comparison with Chiarella and Iori (2002, Figure 2, right panel), it displays the outcomes for 100 periods (i.e., $10,000$ time steps) of the baseline simulation. Here, the time-step return volatility is the absolute proportional change in the market price, as described in section 3.2.1. While there is naturally some variation in behavior across simulation runs, here we see little evidence of the volatility clustering apparent in CI's figure. There is also much less of a visual suggestion that the bid-ask spread correlates with return volatility. For the simulation as a whole, we find the

[13]Boxplots are for trade counts seen in at least 10 periods.

two series to have only a modest positive contemporaneous correlation (as measured by a significant Pearson correlation coefficient of around 0.18).

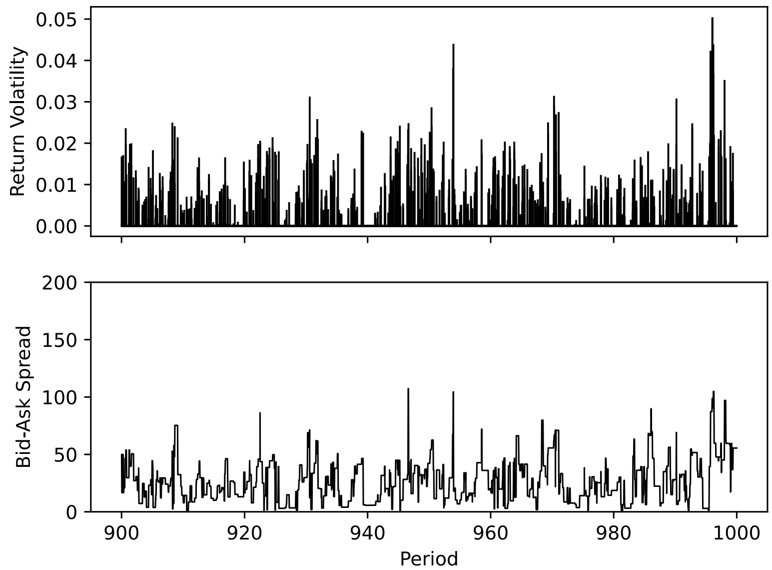

Figure 11: Volatility and Spread
Time-step volatility (top) and bid–ask spread (bottom). Baseline parameterization.

3.3 Sweep Experiments

Since the baseline simulation is not empirically calibrated, robustness checks are important. This section explores the influence of market participation (λ), order duration (τ), and tick size (Δ) on key market behaviors. We examine the effect of parameter variations on the bid-ask spread, market volatility, trading volume, and number of bids and asks in the order book. We discuss the implausibility of some findings reported in Chiarella and Iori (2002), and we show how our results paint a clearer picture of order flow dynamics.

3.3.1 Bids, Asks, and Spreads

Each step, the parameter λ controls the likelihood that a trader will participate in the market. A greater likelihood of market participation should place more bids and asks in the book, increasing market liquidity. In accord with Chiarella and Iori (2002, Figure 3), our Figure 12 confirms this expected effect.

Increased market participation naturally increases the number of trades (as the next section confirms). Whether the bid-ask spread increases or decreases depends on whether liquidity is replenished near the middle of the book faster than trades remove it. Chiarella and Iori (2002, Figure 3) find, rather implausibly, that a very thin market has very low bid-ask spreads, and that spreads rise as market participation increases. That suggests that the trade effect of increased market participation dominates the liquidity effect. In contrast, our Figure 12 shows the expected large spreads in thin markets. As market participation increases, the spread initially falls. Eventually however, at high levels of λ, the trade effect is balanced by the liquidity effect, causing spreads to stabilize.

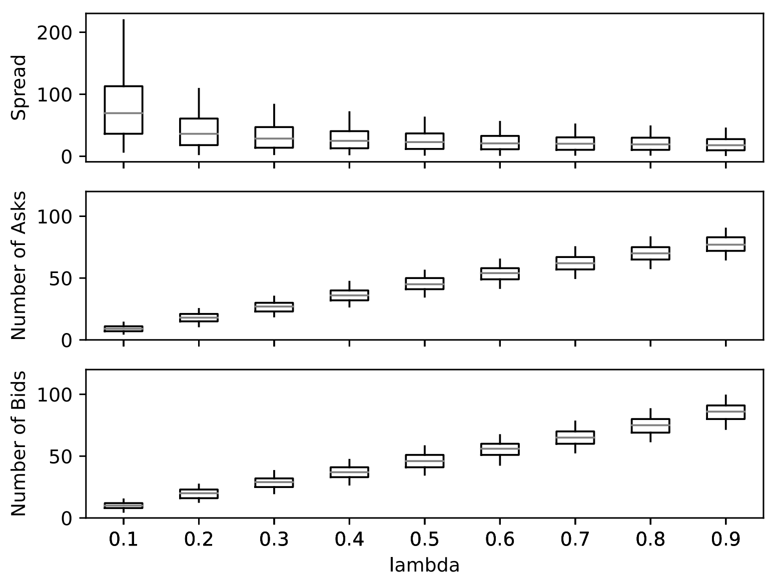

Figure 12: Market Participation and Bid-Ask Spreads
Number of bids, number of asks, and the bid–ask spread as a function of the market participation rate (λ). Interquartile boxplots with median markers and 90% whiskers. Baseline parameterization, except for the participation rate.

Next, consider the influence of the order-duration parameter (τ). This is the time to expiry (in periods) of a new order. A higher τ therefore implies a more liquid order book, as order remain longer in the book. This in turn should reduce the bid-ask spread, which is the finding of Chiarella and Iori (2002, Figure 4). Our Figure 13 corroborates this finding.

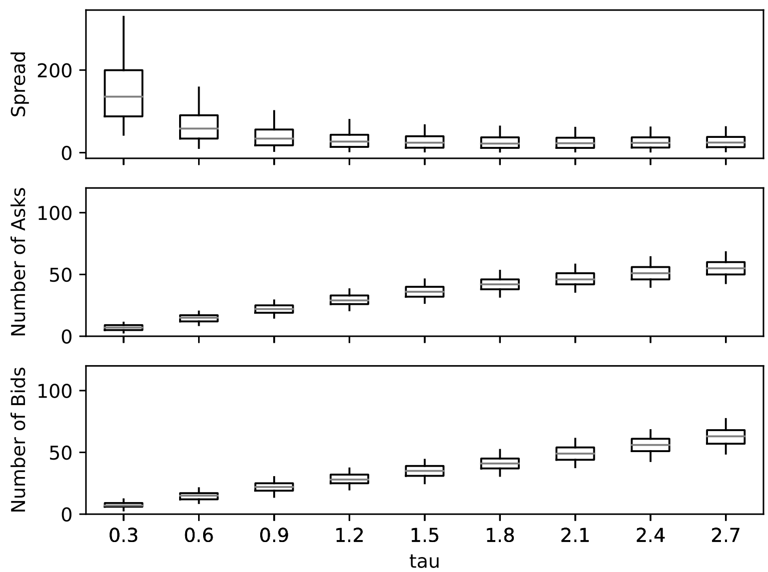

Figure 13: Order Duration and Bid-Ask Spreads
Bids, asks, and bid–ask spread as a function of order duration (τ). Interquartile boxplots with median markers and 90% whiskers. Baseline parameterization, except for the order duration.

Finally, consider the influence of the tick-size parameter (Δ), which is the minimum tick size in this market. The necessary effect on the spread is plain: a higher Δ must increase the spread (to at least match the tick size). This is the finding of Chiarella and Iori (2002, Figure 5), and because it is obvious we do not duplicate it here. (For completeness, the supplement to this paper includes a supportive chart.)

3.3.2 Volume and Volatility

This subsection reconsiders the same parameter sweeps, this time exploring the changing relationship between volume and volatility. First consider variations in the market participation rate (λ). More trade volume is a natural consequence of increased market participation, and removing liquidity from the center of the book can produce a positive volume–volatility relationship. If this *trade effect* dominates, the expected effect on volatility would also be positive.[14] However, more participation could also place more orders in the book, potentially narrowing the spread. This *liquidity effect* can lower measured volatility. The net result on volatility depends on order flow dynamics and is indeterminate.

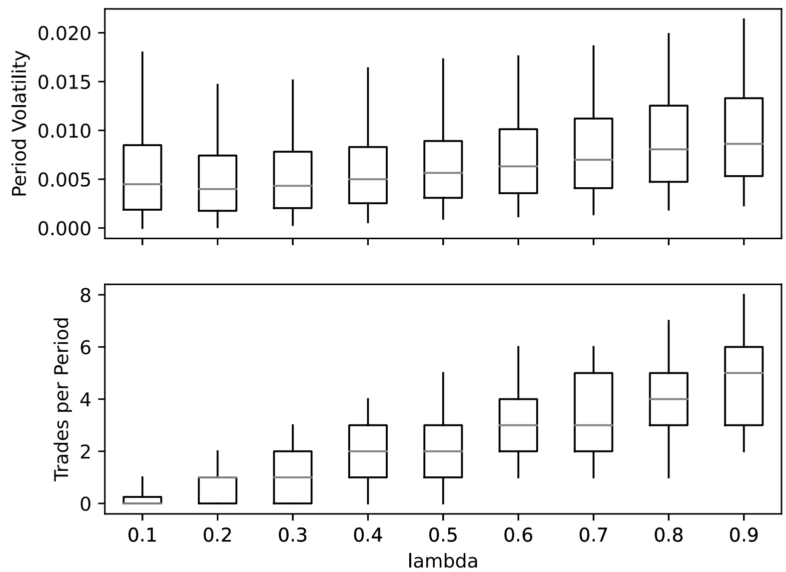

Figure 14: Market Participation, Volume, and Volatility
Trade volume and price volatility as a function of the market participation rate (λ). Interquartile boxplots with median markers and 90% whiskers. Baseline parameterization, except for the participation rate.

Nevertheless, the results reported in Chiarella and Iori (2002, Figure 3) are rather puzzling. First they find that an increase in market participation increases trade volume and increases the spread, suggesting that the trade effect dominates. But they also report a decline in measured volatility, which seems unlikely to be true if both number of trades and the spread are increasing. The conundrum is resolved when we observe that they find an oddly low average number of trades per period—about 1 trade per period at the baseline parameterization. Additionally, they find that the response of trade volume to increased market participation is quite small: increasing λ from 0.2 to 1.0 only moves mean trades

[14]Figure 5 describes the price series, which uses the book's midpoint price if no new trade executes. More market participation makes it more likely that a new bid or ask will beat the top of the book, which increases measured volatility.

per period from ~ 0.7 to ~ 1.3. Thus, the dynamics of volatility are disconnected from the trading frequency, and instead reflect the other drivers of price changes (see section 3.1.1).

Our Figure 14 provides a more helpful window on the volume–volatility relationship. We find a more robust increase in the number of trades as market participation increases. A very modest upward trend in median periodic volatility results. (The supplement to this paper provides additional evidence.) With eventually steady spreads, the increase in volatility is driven by the increase in the number of trades, producing a weakly positive volume–volatility relationship.

Next consider variations in the order-duration parameter (τ). Equation (2) implies that, relative to the last observed price, an increase in τ encourages more aggressive orders. This should produce more volume and more volatility. Separately, however, an increase in order duration implies a more liquid order book, which could lead to smoother price transitions and lower volatility. As illustrated by Figure 15, the outcomes are complicated. Volume does demonstrate the expected positive relationship to τ, but period volatility declines before starting to rise. (See the supplement to this paper for additional evidence.) Our result conflicts with the puzzling weak effect on volume and the strong decline in volatility reported by Chiarella and Iori (2002, Figure 4).

Finally, consider the influence of tick size (Δ). We see little reason for modest changes in tick size to affect volume. A large enough increase, however, could harm the ability to trade within relatively small price ranges. For similar reasons, a modest increase in measured price volatility is possible, because a larger tick size increases the required price jumps for trades to take place at different prices. These expectations match Figure 16. However, they conflict with Chiarella and Iori (2002, Figure 5), which associates substantial increases in volume and orders of magnitude decreases in volatility with large increases in tick size. This is counterintuitive, and we could find no way to produce such a result.

However, CI's volatility vs tick size subfigure is oddly labeled. We suspect that they actually found *rising* volatility, as they claim in the text of the paper. If so, their result may

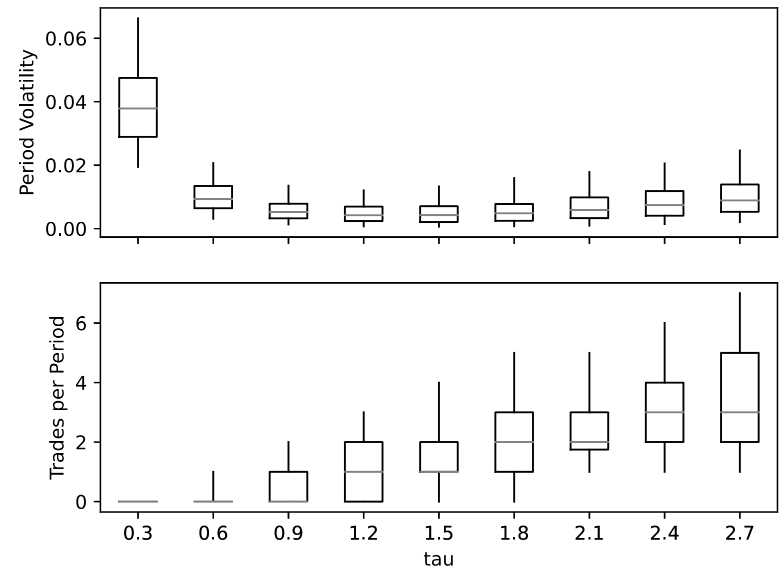

Figure 15: Order Duration, Volume, and Volatility
Trade volume and price volatility as a function of order duration (τ). Interquartile boxplots with median markers and 90% whiskers. Baseline parameterization, except for the order duration.

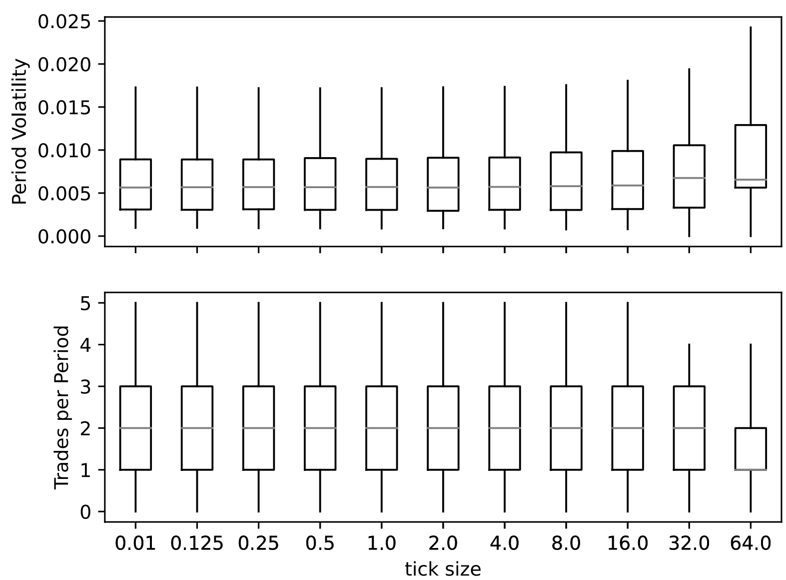

Figure 16: Tick Size, Volume, and Volatility
Trade volume and price volatility as a function of tick size (Δ). Interquartile boxplots with median markers and 90% whiskers. Baseline parameterization, except for the tick size.

hinge on a rather subtle modeling choice in the rounding of bids and ask to ticked prices. As shown in the supplement to this paper, we can produce rising volume and volatility by rounding trader prices to the nearest tick. In contrast, our results use directional rounding (i.e., floor rounding for bids, and ceiling rounding for asks), which is more appropriate. Without directional rounding, a rounded bid could exceed the traders target price, despite price shading. The high volume and volatility at larger tick sizes may still seem anomalous, but when the tick size precludes trading at p^f, this can induce fundamentalist-driven dynamics in the CI model. Fortunately, this happens only at tick sizes that lack real-world relevance.

4 Conclusion

Research in agent-based computational finance accelerated when Gode and Sunder (1993) suggested that common institutional constraints could lead zero-intelligence (ZI) traders toward efficient market outcomes. Subsequent research shifted towards near-zero-intelligence agents, who adopt stylized trading rules distilled from observed practices. In this vein, Chiarella and Iori (2002) claim to produce realistic returns series and market dynamics in an order-driven exchange where traders use noise-trading, fundamentalist, and chartist strategies. The CI model has hundreds of citations and has motivated a substantial body of research in agent-based finance.

The CI model was instrumental in stimulating research in order-driven markets populated by a variety of trader types. Key claims of the original paper may be summarized as follows. Chartist behavior must supplement fundamentalist and noise trader behaviors in order produce a returns series with realistic features. Asset returns display volatility clustering. Volatility is positively correlated with volume and the bid-ask spread. And, the spread, volatility, volume and liquidity exhibit particular nonlinearities in the relationship market participation (λ), order duration (τ), and tick size (Δ).

Despite its important role in the literature of agent-based finance, key elements of the CI model's specification are missing in the original exposition, and its original implementation in code is unavailable. We therefore revisit the CI model with the goals of providing a detailed specification, re-examining key findings, and exposing new insights. We identify missing details, discuss alternative implementation approaches, and suggest solutions that enable replication of key results of the original study. After overcoming substantial barriers to replication, we are able to confirm a number of their results. Nevertheless, we contest several of their claims, and we shed new light on others.

Our most startling result is that momentum trading is not important in the CI framework. The core reason derives from a natural feature of the simulation data: the expected returns from the traders' momentum strategy are always very close to zero. As a result, momentum trading in the CI model has negligible effects on the price and return series of the model.

Additionally, we show that some of the reported properties of the CI model that originally appeared surprising or counterintuitive are indeed not replicable. For example, the original study reports that a higher trader participation rate results in both a higher spread and a higher trading volume but also that it leads to lower volatility. The former is potentially explicable, if higher volume somehow guts liquidity near the midpoint of the order book, but then measured volatility should also increase. Our implementation of the CI model produces more reasonable results. Excepting minimal levels of participation, we find that that increases in market participation decrease spreads while increasing volume (as should be expected). Eventually the increase in volume dominates the decline in spreads, leading to a small increases in measured volatility. We find a similar result when varying order duration. Our implementation thereby resolves apparent puzzles in the original study's findings.

Our paper additionally exposes some alignments and deviations between outcomes in the CI model of the behavior of actual asset prices. This suggests directions for future research on two fronts. On the theoretical front, since chartist behavior is well documented in actual markets (and well supported by many trading platforms), the CI framework needs a better formulation of momentum trading. On the empirical front, research needs to clarify which extensions of the CI model can better approximate the actual behavior of assets that trade on order-driven exchanges. This may required expanding the scope of the model to include features such as variable order quantities, variable time in force, explicit characterizations of order cancellation strategies, and interactions between multiple order types.

References

Beja, A. and Goldman, M.B., On the Dynamic Behavior of Prices in Disequilibrium. *The Journal of Finance*, 1980, **35**, 235–248.

Bouchaud, J.P., Agent-Based Models for Market Impact and Volatility. In *Handbook of Computational Economics*, Vol. 4, pp. 393–436, 2018, Elsevier.

Bowden, R.J., Predictive Disequilibria and the Short Run Dynamics of Asset Prices. *Australian Journal of Management*, 1990, **15**, 65–87.

Chakraborti, A., Toke, I.M., Patriarca, M. and Abergel, F., Econophysics Review: II. Agent-Based Models. *Quantitative Finance*, 2011, **11**, 1013–1041.

Chiarella, C., The Dynamics of Speculative Behaviour. *Quantitative Finance*, 1992, **37**, 101–123.

Chiarella, C. and Iori, G., A Simulation Analysis of the Microstructure of Double Auction Markets. *Quantitative Finance*, 2002, **2**, 346–353.

Cont, R., Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, 2001, **1**, 223–236.

Cont, R., Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models. In *Long Memory in Economics*, edited by G. Teyssière and A.P. Kirman, pp. 289–309, 2007 (Springer Berlin Heidelberg: Berlin, Heidelberg).

Escanciano, J.C. and Lobato, I.N., An Automatic Portmanteau Test for Serial Correlation. *Journal of Econometrics*, 2009, **151**, 140–149.

Frankel, J. and Froot, K., Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations. *American Economic Review*, 1987, **77**, 133–53.

Glosten, L.R., Is the Electronic Open Limit Order Book Inevitable?. *The Journal of Finance*, 1994, **49**, 1127–1161.

Gode, D.K. and Sunder, S., Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 1993, **101**, 119–137.

Gould, M.D., Porter, M.A., Williams, S., McDonald, M., Fenn, D.J. and Howison, S.D., Limit Order Books. *Quantitative Finance*, 2013, **13**, 1709–1742.

Jarque, C.M. and Bera, A.K., Efficient Tests for Normality, Homoscedasticity and Serial Independence of Regression Residuals. *Economics Letters*, 1980, **6**, 255–259.

LeBaron, B., Agent-Based Computational Finance: Suggested Readings and Early Research. *Journal of Economic Dynamics and Control*, 2000, **24**, 679–702.

Lillo, F., Limit Order Placement as an Utility Maximization Problem and the Origin of Power Law Distribution of Limit Order Prices. *The European Physical Journal B*, 2007, **55**, 453–459.

Ljung, G.M. and Box, G.E., On a Measure of Lack of Fit in Time Series Models. *Biometrika*, 1978, **65**, 297–303.

Lux, T. and Marchesi, M., Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents. *IFAC Proceedings Volumes*, 1998, **31**, 7–10 IFAC Symposium on Computation in Economics, Finance and Engineering: Economic Systems, Cambridge, UK, 29 June - 1 July.

Maslov, S., Simple Model of a Limit Order-Driven Market. *Physica A: Statistical Mechanics and its Applications*, 2000, **278**, 571–578.

McLeod, A.I. and Li, W.K., Diagnostic Checking ARMA Time Series Models using Squared-Residual Autocorrelations. *Journal of Time Series Analysis*, 1983, **4**, 269–273.

Palmer, R., Brian Arthur, W., Holland, J.H., LeBaron, B. and Tayler, P., Artificial Economic Life: A Simple Model of a Stockmarket. *Physica D: Nonlinear Phenomena*, 1994, **75**, 264–274.

Raberto, M., Cincotti, S., Focardi, S.M. and Marchesi, M., Agent-Based Simulation of a Financial Market. *Physica A: Statistical Mechanics and its Applications*, 2001, **299**, 319–327 Application of Physics in Economic Modelling.

Riley, J.G. and Samuelson, W.F., Optimal Auctions. *The American Economic Review*, 1981, **71**, 381–392.

Zeeman, E., Differential Equations for the Heartbeat and Nerve Impulse. In *Dynamical Systems*, edited by M. Peixoto, pp. 683–741, 1973 (Academic Press: Salvador, Brasil).

Zeeman, E., On the Unstable Behaviour of Stock Exchanges. *Journal of Mathematical Economics*, 1974, **1**, 39–49.

Zhan, W. and Friedman, D., Markups in Double Auction Markets. *Journal of Economic Dynamics and Control*, 2007, **31**, 2984–3005.

A Paper Supplement

The following material is informational only. It is *not* part of the paper. The code for this paper is available upon request.

A.1 Outline of Paper and Supplement

A.2 Baseline Parameterization: Details

```
#Baseline parameterization for CI
[nTraders]
value = 1000
description = "Number of traders"
source = "CI p.350"
[maxOrderSize]
value = 1
description = "maximum order size"
source = "CI p.348"
[minOrderSize]
value = 1
description = "minimum order size"
source = "CI p.348"
[lamb]
value = 0.5
description = "probability a trader is activated (chosen to trade), each tick"
source = "CI p.350; CI call this liquidity"
[stdWeightPf]
value = 1.0
description = "std of fundamentalist weight (mean is 0)"
source = "CI p.347 & 351"
[stdWeightChart]
value = 1.4
description = "std of chartist weight (mean is 0)"
source = "CI p.347 & 351"
[stdWeightNoise]
value = 3.0
description = "std of noise-trader weight (mean is 0)"
source = "CI p.350 (called n in Fig 1 or n0 on p.349)"
[weightsType]
value = "CI"
description = "weights style (CI|LP|random)"
source = "New parameter."
note = "introduced to compare CI vs LP"

# Market-wide attributes
[pf]
value = 1000.0
description = "fundamental price (shared by all traders)"
source = "CI p.350"
[ticksize]
value = 0.01
description = "mkt tick size (price points at which valid orders are placed)"
source = "CI p.350 (names this Delta)"
[pctMinPrice]
```

```
value = 1
description = "minimum allowable order price, as pct of fundamental"
source = "CI grid is unspecified; we allow large deviations from pf"
note = "60 in LP (L email 08/11/20)"
[pctMaxPrice]
value = 200
description = "maximum allowable order price, as pct of fundamental"
source = "CI grid is unspecified; we allow large deviations from pf"
note = "140 in LP (L email 08/11/20)"
[stdNoise]
value = 1.0
description = "std of mkt noise in CI"
source = "CI p.349"
[orderQSettings]
value = [1.0, 1.0, 0.0, 0.0]
description = "in order these represent: baseline size of liq. consuming and
   providing orders, volatility weights for each"
source = "new parameter. At baseline values, quantity=1 as in CI."
[directedTickRounding]
value = true
description = "Bid price rounds down to nearest tick; ask rounds up."
source = "Set to false to reproduce CI fig 5, using simple rounding."

# Determinants of agent-specific attributes
[maxlagShortVolatility]
value = 5
description = "Horizon over which short volatility is calculated (in periods)"
source = "not in CI; 500 in LP"
note = "ignored by baseline parameterization"
[maxlagMomentum]
value = 100
description = "Maximum price history (in time steps) used by traders"
source = "CI p.350 sets their Lmax to 100 time steps (= 1 period)"
note = "Called Lmax in our paper, as in CI."
[tau]
value = 2.0
description = "Time in force for orders (in time periods of 100 steps)."
source = "CI p.350"
[rhoLP]
value = 0.0
description = "Multiperiod persistence (AR) in convergence to fundamental price."
source = "Not present & tf effectively 0 in CI; 0.995 in LP."
note = "LP's representation of multiperiod persistence (AR) in convergence to
   fundamental price; not recommended. See paper."
[maxPShade]
value = 0.5
description = "Maximum price shading that traders will apply"
```

```
source = "CI p.350 (as kmax)"
[minPShade]
value = 0.0
description = "Minimum price shading that traders will apply"
source = "not in CI but implicitly 0 on p.350"
note = "LP use -maxPShade."
[maxQShade]
value = 0.5
description = "Maximum quantity shading that traders will apply"
source = "New parameter; ignored in baseline."
[minQShade]
value = 0.0
description = "Minimum quantity shading that traders will apply"
source = "New parameter; ignored in baseline."
[bPshading]
value = false
description = "Volatility based price shading"
source = "LP p.11; not in CI; ignored in baseline."
note = "LP's extension to generate volatility persistence."
[aggQshading]
value = false
description = "Aggression based quantity shading"
source = "New parameter; ignored in baseline."
[volQshading]
value = false
description = "Volatility based quantity shading"
source = "New parameter; ignored in baseline."
[orderTypes]
value = 1
description = "determines allowable order types (limit or market)"
source = "New model parameter."
[fixedQ]
value = 1
description = "Fixed order size (1 in CI); 0 means not fixed"
source = "CI p.350"

## simulation parameters
[nPeriodLength]
value = 100
description = "number of time steps (half minutes) in a time unit (period or hour
   ); also see maxlagMomentum"
source = "CI p.348"
note = "Experiments that change this will of course change period aggregates!"
[nPeriods]
value = 1000
description = "Number of periods per scenario (parameterization)"
source = "Based on x-axis values on Figures 1 and 2, CI p.347, 348."
```

```
[seed]
value = 40
description = "Initial seeds for prng for each experiment."
source = "arbitrary"
```

A.3 CI Figure 1

Figure 17 is our reproduction of Chiarella and Iori (2002, Figure 1). We include it here just to show that we can produce very similar results. CI use this figure to suggest that producing realistic price and returns series will require traders to employ fundamentalist, noise trader, and chartist strategies. However, as shown in our paper, the set of weights chosen for each strategy in this figure make it difficult to draw this conclusion satisfactorily. In fact, our paper overturns their claim.

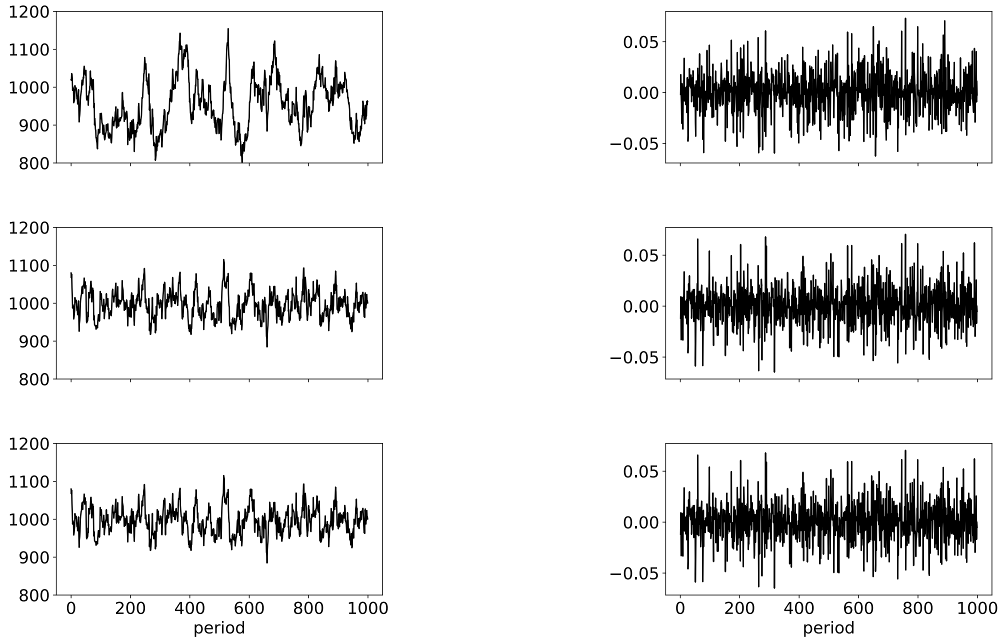

Figure 17: Replicating CI Figure 1
Prices (left) and returns (right). Top row: only noise traders participate in the markets ($\sigma_1 = \sigma_2 = 0$). Middle row: fundamentalists and noise traders participate in the markets ($\sigma_1 = 5.0$, $\sigma_2 = 0$). Bottom row: fundamentalists, chartists, and noise traders participate in the market ($\sigma_1 = 1.0$, $\sigma_2 = 1.4$). The other parameters have baseline values.

The figure plots prices (left) and returns (right). The first and second panels isolate the impact of the noise strategy on the price and returns series. The third and fourth panels allow the reader to identify the marginal effect of adding a strong fundamentalist strategy to the noise strategy. In these panels, raising the variance σ_1 to 5.0 causes the price series to display a strong mean-reverting tendency, a finding we replicate. The fifth and sixth panels add the chartist strategy to the traders' strategy mix, while also reducing the mean-reverting tendency by lowering σ_1 to 1.0.

This choice of parameterization unfortunately does not allow a direct comparison of the three strategy components, and particularly obscures the marginal effect of the chartist strategy on the price and returns series. Performing such an exercise, as we do in our paper, attests to the negligible influence of chartists in the CI model.

A.4 Asset Returns over Time

Figure 18 presents the asset returns associated with the prices in Figure 6 of the paper. Correspondingly, each return series in Figure 18 is a proportional change in end-of-period prices. Compare this figure to the right column of Chiarella and Iori (2002, Figure 1). Figure 18 varies σ_1 and σ_2 in exactly the same manner as Figure 6. This means that the first, fifth, and fourth figures correspond to the right column of CI's Figure 1. (When making a comparison, carefully note the change in scale in CI's final subfigure.)

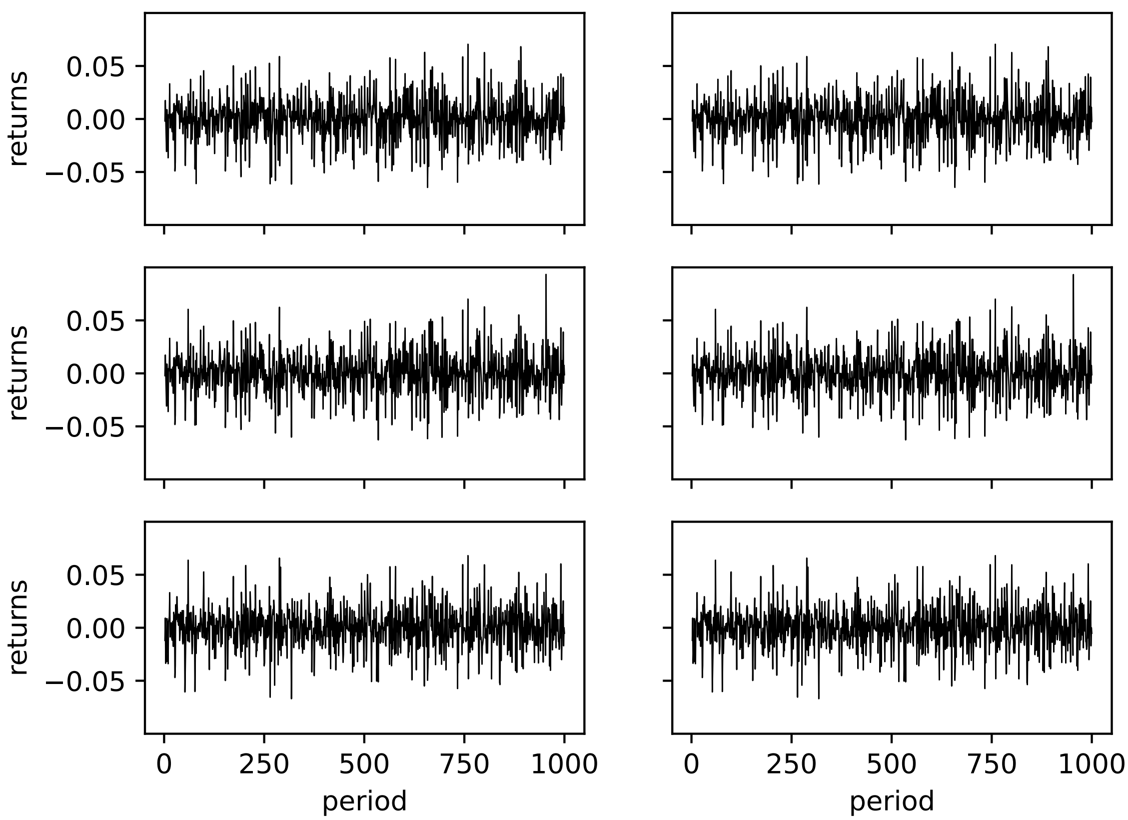

Figure 18: Asset Returns over Time
Top to bottom, rows correspond to σ_1 equal to 0.0, 1.0, and 5.0. Left to right, columns correspond to σ_2 equal to 0.0 and 1.4. (So the first, fifth, and fourth figures correspond to the right column of CI's Figure 1.) End-of-period returns. Baseline parameterization, except for σ_1 and σ_2.

A.5 Strategy Components

For all orders placed during a simulation run, the top panel of Figure 19 displays the kernel density estimates of these components. The bottom panel shows the kernel density estimates of the *weighted* components—that is, $\sigma_0\epsilon_t$, $\sigma_1 r_t^f$, and $\sigma_2\bar{r}_t^L$, where the weights $(\sigma_0, \sigma_1, \sigma_2)$ are the strategy weights of the trader placing the order.

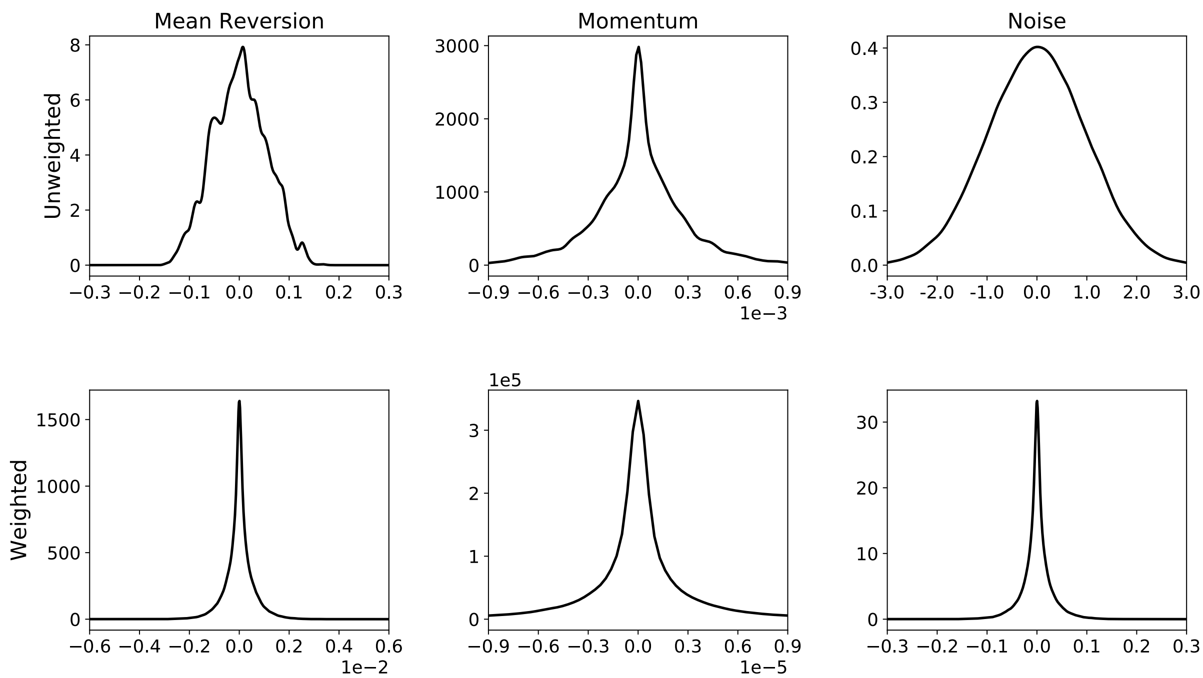

Figure 19: Kernel density estimate plot of unweighted (top panel) and weighted (bottom panel) components of final predicted returns. Simulation conducted under baseline CI parameterization.

Figure 19 makes two features of this market immediately clear: first, that the predicted returns of traders is, on average, dominated by the noise component; and second, that the contribution of the chartist component is extremely small.

The unweighted noise component is drawn from a standard normal distribution, as described in section 3.1, while the unweighted fundamentalist component is centered around a positive value. The fundamentalist component is about 1 order of magnitude smaller before weighting. Note, however, that this illustration underemphasizes the impact of the fundamentalist component on the price series. While the return values are smaller *on average*, they exhibit high autocorrelation. The impact is apparent in Figure 6 of the paper and Figure 17 above.

In contrast, the returns from the momentum component are about 4 orders of magnitude less than those from the noise component, and 3 orders of magnitude less than those from the fundamentalist component. Further, the median return value for the momentum component is 0.0. Weighting does not meaningfully change this relative importance. Thus, the momentum component barely contributes anything to the final predicted return of the trader.

The irrelevance of the momentum component can be overcome to some extent by weighting it much more heavily. Even then, two additional problems remain. First, it is not clear that the method of averaging past returns is a good characterization of chartist behavior in the real world. Second, the trader-specific weight assigned to the momentum component in the CI model can be either positive (indicating the trader is trend-following) or negative (indicating the trader is contrarian). Recall that these weights are drawn from a symmetric distribution. This implies that half of all traders are always undermining the trend — whose magnitude, as we established, is already very small. These issues point to a deeper problem for characterizing chartism in the CI model.

A.6 CI Sweep Experiments

In the accompanying paper, we choose quite different visualizations of the outcomes of the parameter sweep experiments than the ones favored by Chiarella and Iori (2002, Figures 3–5). The paper explains the advantages of our approach. Nevertheless, for ease of comparison, this supplement presents more easily comparable visualizations.

In producing these visualizations, we encountered another replication difficulty. CI do not document their strategy for sampling the parameter space. They appear to use some kind of adaptive random sampling, but its nature is undocumented and it produced very odd sampling gaps.[1] In this supplement, we approach random sampling in a more straightforward way by choosing values from a uniform distribution in the relevant interval. This proves to be at least as revealing as the CI approach (although still inferior to the deterministic sampling used in our paper).

One question that can naturally arise when comparing our sweep results to the CI sweep results concerns the relatively smooth relationships we report. The answer is simple: we identically seed the random-number generator for each scenario, so we avoid introducing spurious volatility into the results. As our paper indicated when comparing our figure 6 to Chiarella and Iori (2002, Figure 1), failure to use a common seed can prove extremely misleading.

A.6.1 CI Figure 3

Figure 20 attempts to reproduce the findings presented in Chiarella and Iori (2002, Figure 3). We sample 200 values for the market participation rate (λ) from a standard uniform distribution and then run a complete simulation for each value. As before, each plot marker represents data from one full simulation run. (That is $1,000$ periods, or equivalently $100,000$ simulation steps).

[1] For example, in Figure 3 on page 349, there exist gaps in the values for λ approximately between 0.025 and 0.125, 0.5 and 0.7, and again between 0.7 and 0.8.

Our results are qualitatively similar to those reported by CI. Nevertheless, the are a few important differences, and they are in line with the discussion in our paper. (See section 3.3.) Most obviously, the scale of the bid-ask spread in our simulation exceeds theirs. Additionally, at values for λ nar 0.1, we see strongly nonlinear spread responses to the value of λ, while they have a sampling gap in this region.

With respect to volatility, we find a similar strongly nonlinear response at low values for λ, but we find volatility is ultimately increasing in λ. This may differ from CI, or their coarser volatility scale may simply hide it. In either case, our finding supports the positive volume–volatility relationship seen in real-world financial markets. Indeed, we find (in the third panel) an almost linear relationship between volume and λ. As discussed in our paper, this reflects a stronger relationship between market participation and volume than is seen in the original CI results. The muted response of volume to λ in CI is counter-intuitive, as is the behavior across variables like the spread and volatility, as we argue in the main text. Finally, the evolution of the number of active bids and asks looks very similar to the CI result, both in trend and level.

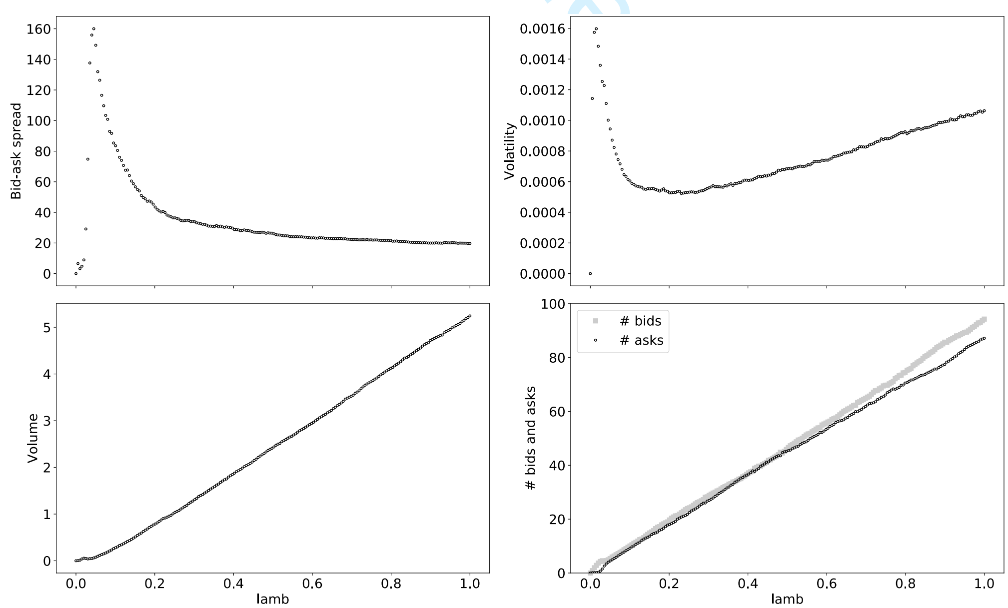

Figure 20: Variations in Market Participation
Bid–ask spread, volatility, trading volume and number of bid and ask orders in the book as a function of market participation (λ). Baseline parameterization except for λ. Compare to Chiarella and Iori (2002, Figure 3).

A.6.2 CI Figure 4

Figure 21 attempts to reproduce the findings presented in Chiarella and Iori (2002, Figure 4). For ready comparison with the CI figure, we sample 200 values for order duration (τ) from a uniform distribution in the interval $[0, 3]$. As before, each plot marker represents data from a full simulation of 1000 periods.

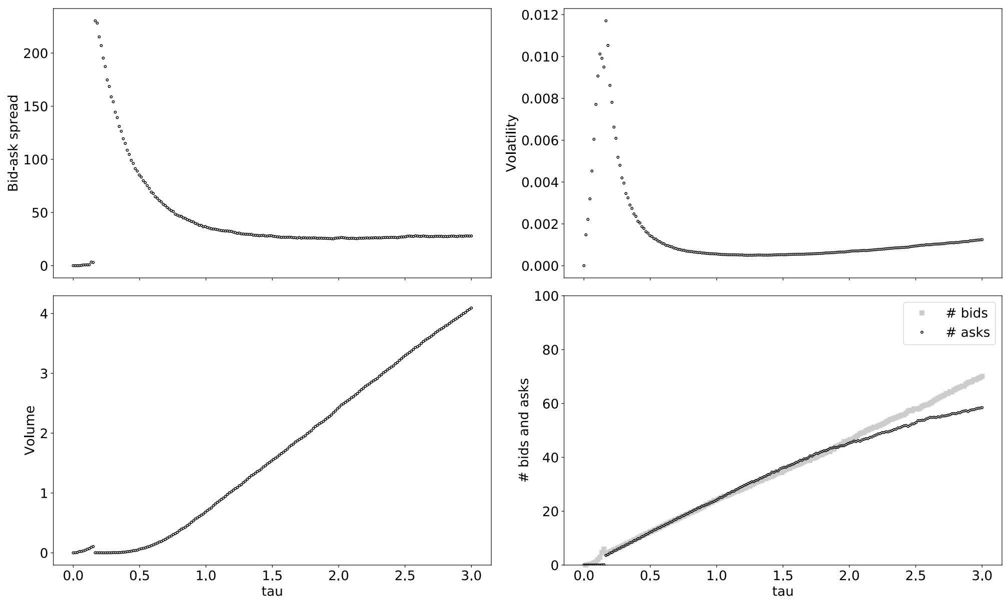

Figure 21: Variations in Order Duration
Bid–ask spread, volatility, trading volume, and number of active bid and ask orders as order duration (τ) varies. Compare to Chiarella and Iori (2002, Figure 4).

Our findings are mostly similar to CI's, but there are some important differences. (Attend to the vertical scale in our Figure 21.) In particular, we cannot reproduce the behavior of periodic trade volume in their third panel. At $\tau \approx 0.15$ we find a minor hint of the nonlinearity they show, but for the most part we find volume increases almost linearly in τ. That trade volume would increase with order duration is our expected outcome. As in Figure 20, we again find a larger bid-ask spread. However, the number of bids and asks as well the evolution of volatility are close to the CI results.

A.6.3 CI Figure 5

This section grapples with the results presented in Chiarella and Iori (2002, Figure 5), which considers the effects of tick-size variation. Our tick-size sampling is uniform on a logarithmic interval (for easier comparison with the CI figure). After uniformly sampling 200 values from the interval $[-2.0, 2.0]$, we calculate the values for the tick size (Δ) as 10^i for each i in the above set of values.

When considering tick-size variation, some of the results we report in the paper are far from the CI results. Ascertaining the reason was not easy, but we finally located a likely suspect. The market tick size forces order prices onto a finite grid. CI do not document the rounding rule they use for this. One possible approach is ordinary *undirected* price rounding. This rounds the order price to the closest tick level, regardless of the direction of adjustment with respect to the order price. A better approach is *directed* price rounding, which rounds bids down and asks up to the nearest tick in the specified direction.

Under directed rounding, order prices are always adjusted towards less aggressive price points. Under undirected rounding, order prices may get adjusted not just more aggressively than the shaded price but even beyond the target price. This becomes increasingly likely as the tick size increases. This is both unrealistic and in contradiction to the spirit in which CI implement price shading. (See section 2.3.2 of the paper.) Directed rounding is therefore preferred. However, we find evidence that CI use undirected rounding.

Figures 22 and 23 display the outcomes with directed and undirected rounding. The key difference between the two approaches lies in the response of volume and volatility at large tick sizes. With directed rounding, the results provide another window on the summary results presented in our paper. As expected, volume falls as the tick size becomes larger, and eventually we get a small increase in volatility due to the jumps between adjacent prices. This reflects the adjustment of order prices to more passive price points. The reduced price aggressiveness implies fewer successful trades, and the decline in volume in the third panel reflects this. However, when trades do occur, higher tick size requirements generate larger

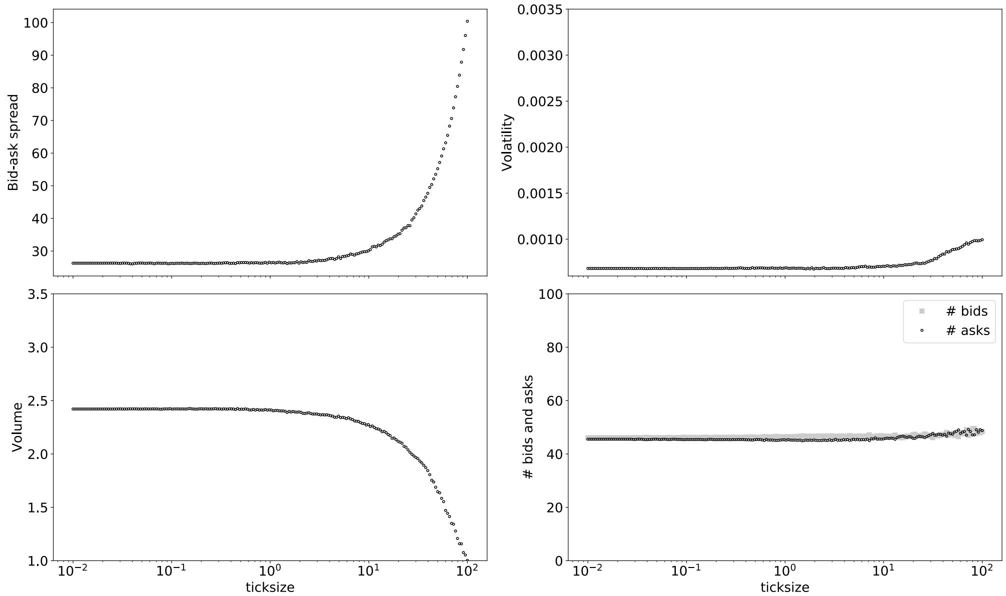

Figure 22: Ticksize Variation with Directed Rounding

Bid–ask spread, volatility, trading volume and number of active bid and ask orders as tick size (Δ) varies. Baseline parameterization, except for tick size, *so* order-price rounding is directed. Compare to CI's Figure 5 and to our Figure 23.

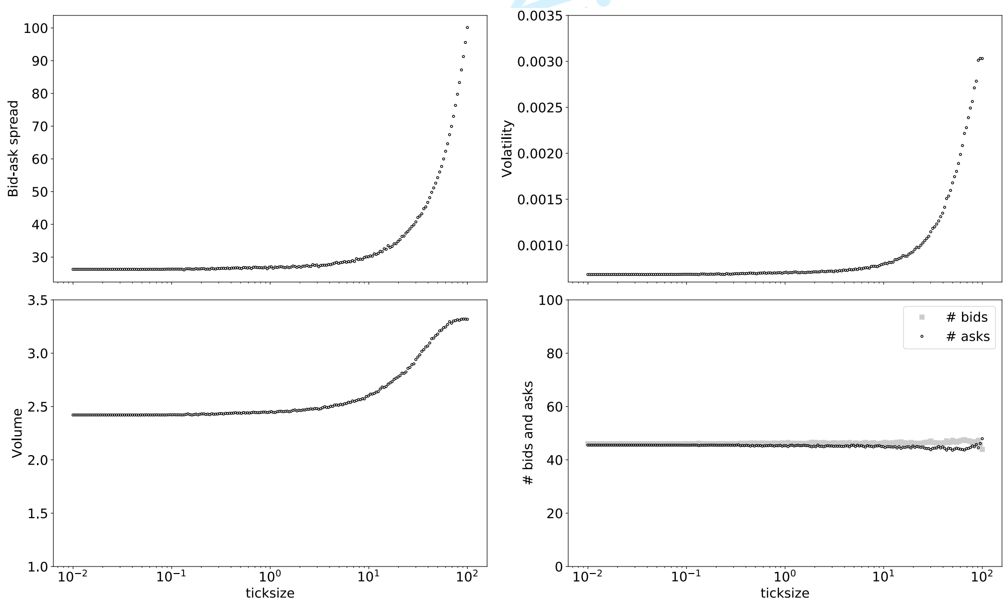

Figure 23: Ticksize Variation with Undirected Rounding

Bid–ask spread, volatility, trading volume, and number of active bid and ask orders as tick size (Δ) varies. Baseline parameterization, except for tick size *and* use of undirected order-price rounding. Compare to CI's Figure 5 and to our Figure 22.

price movements. At large tick sizes, this yields slightly higher measured volatility.

When rounding is undirected, the result is very different. Rounding to the nearest tick causes some orders to become more aggressive, and this is reflected in higher volume as more orders cross the book. The associated volatility is also correspondingly much higher. Note in particular that the results under undirected rounding are qualitatively very similar to CI's results.[2] This suggests that CI adopted undirected rounding.

[2]A final additional source of confusion is the oddly labeled y-axis on the volatility plot in CI's Figure 5 (second panel), which suggests that volatility is *decreasing* in tick size. Since the text characterizes the response of volatility as increasing in Δ. we presume that this axis is mislabeled, and we correspondingly favor their textual description.