Dynamic Order Dispersion and Volatility Persistence in a Simple Limit Order Book Model*

Andrew Hawley[†]
Blake LeBaron[‡]
Mark Paddrik[§]
Nathan Palmer[¶]

May 23, 2024

Abstract

This preliminary paper extends the dynamics of a basic stylized limit order book model introduced in Chiarella & Iori (2002). The original model is capable of generating some key market microstructure features, but it cannot recreate longer range persistence in volatility. We explore a very simple and intuitive addition to the stylized, near zero intelligence behavior of traders that is capable of delivering persistent volatility. We also show that this strategy depends critically on certain key features in the dynamics of supply and demand for liquidity and depth in the limit order book. We believe this is fundamental to understanding both the dynamics of volatility in financial time series, along with variations in liquidity in financial markets. We contribute a parsimonious agent-based model to the literature that may be used as a test bed or sandbox for developing agents with more complex behavior.

Keywords: Agent-based finance, Limit order books, Liquidity, Volatility

*The views expressed in this paper are solely those of the authors and do not necessarily reflect the position of the Office of Financial Research (OFR), the U.S. Department of the U.S. Treasury, or the Federal Reserve Board of Governors.

[†]Board of Governors, Federal Reserve System
[‡]Brandeis International Business School
[§]Office of Financial Research
[¶]Board of Governors, Federal Reserve System

1 Introduction

Agent-based financial markets are often constructed using many different forms of adaptive trader behavior. This can lead to complex models with features that are difficult to connect to any single design feature. Therefore, it is not surprising that very stripped down models of trader behavior have always provided an important benchmark for agent-based models.[1] Beyond getting market simulations closer to their basic core, these models are also important in how they highlight institutional design questions for markets. The goal of this paper is to explore changes in trading mechanisms and their impact on an environment that is not closely linked to any one type of trader, or learning mechanism.

Agent-based models also employ many stylized mechanisms for price setting. We will follow the early paper of Chiarella & Iori (2002) for our market design.[2] In this important early agent-based model the authors put relatively simple, stylized, agents into a basic limit order book. The model presents a useful benchmark for limit order book dynamics that we also take as our foundation. Limit order books are a useful mechanism for modeling trade in agent-based environments for several reasons. First, they are close to many institutions used in actual financial markets. The details may vary, but the basic structures of order placement, crossing, and cancellations are close to many market institutions currently in use. Second, they are a simple trading mechanism that nails down a lot of institutional detail for the modeler. Order book dynamics can only be implemented in a few ways, presenting a useful modeling constraint. Finally, as we will argue in the paper, they form a metaphor for other financial and economic trading that may appear far from a traditional limit order book.

In this paper we explore some features of the basic model. We show that standard time series properties such as uncorrelated returns, and excess kurtosis are relatively generic to the model. We modify the strategies slightly in a simple fashion to generate more realistic features in terms of long range persistence in volatility. The simplicity of the model allows us to analyze the mechanisms going into volatility in a very direct fashion.[3] We find that it requires both simple forms of adaptation to local volatility, but also important features related to liquidity dynamics on the order book. We are able to dramatically alter price/return dynamics by modifying a few key market structures in the model. This gives a new perspective on the dynamics of volatility in markets, and its connection to changing patterns of market liquidity and overall market stability and efficient function.

In Section 2, we summarize the literature of mechanisms which are believed to cause volatility persistence. In Section 3, we describe the base limit order book model and how we introduce

[1] Gode & Sunder (1993) is the classic example of this, influenced by early work of Becker (1962).

[2] Other examples of zero intelligence and limit order books can be found in Farmer, Patelli & Zovko (2005) and Ladley & Schenk-Hoppe (2009), and an early example using purely mechanical order flow is Cohen, Maier, Schwartz & Whitcomb (1983).

[3] Several earlier papers have also explored mechanisms for generating long range persistence in volatility. Among the conjectured mechanisms are order splitting, Toth, Palit, Lillo & Farmer (2015) and evolution and herding in strategies, LeBaron & Yamamoto (2008).

volatility persistence. In Section 4 we experiment with several different mechanisms for influencing volatility persistence, and report on our findings. In section 5 we explore several different ways to eliminate volatility persistence and leptokurtic return distributions from our simulated time series. In section 6 we look at the joint dynamics of the order book and corresponding time series giving us a very realistic picture of changing liquidity in the market. Conclusions follow in Section 7. The appendix presents some detailed parameter sweeps.

2 Background

This paper follows in a long tradition of trying to understand the nature of liquidity in financial markets. The finance world, both academic and practitioner, has struggled for a long time to find a precise definition of a concept for which everyone has an intuitive understanding. The model used here is built from a standard limit order book which is a common trading institution. It serves as a useful platform for understanding endogenous patterns in liquidity which our model is able to generate. Bouchaud (2018), Bouchaud, Farmer & Lillo (2009), and Gould, Porter, Williams, McDonald, Fenn & Howison (2013) are three complete surveys on the structure of limit order books. They emphasize both the wide range of theoretical models in use, and the extensive array of well known empirical features and the puzzles high frequency financial data has presented.

Among these facts are some of the basic empirical features displayed by prices and returns at relatively high frequencies (less than monthly frequency). These include near random walk behavior in prices, return distributions which are generally leptokurtic (fat tailed), and highly persistent volatility.[4] High frequency data displays two other time series features. Trading volume is also highly persistent. Beyond this, there is a strong feature in signed order flow which can be represented simply as $(1, -1)$ at each time period t depending on whether an order is a buy or sell respectively. This series is also highly persistent, and is probably well represented by a long memory process.

There are many other stylized facts to limit order books that hold with varying degrees of regularity across markets.[5] Some pertain to the shape of the book, and another related set of features deals with price impact, or the impact of new orders on the current price. We deal with price impact and its relation to liquidity only through the critical cross correlations it implies. If liquidity generally refers to the ease of trading both in terms of time and low price impact, then it should also be indirectly revealed through bid/ask spreads and price volatility. The positive correlation between these has been well documented in papers such as Bollerslev & Domowitz (1993), Brock & Kleidon (1990), Flemming (1997), and Mike & Farmer (2008).[6] There is also evidence that volatility moves in critical ways with the nature of order flows (market versus limit) and depth on the book. A thin book tends to lead to increases in volatility which then eventually draws in more liquidity providing limit orders, calming eventual price fluctuations. Two examples of this evidence are in Ahn, Bae

[4]These facts are documented in many papers, but Cont (2001) is a good summary.

[5]See Cont (2011) for a full set of high frequency facts.

[6]All of these papers also reveal an interesting daily seasonal pattern which will not be explored here.

& Chan (2001) and Handa & Schwartz (1996).[7] There is also evidence connecting volatility to the state of the order book. Naes & Skjeltorp (2006) report that the slope of the orders on the book (moving away from best bid and ask) is connected to price volatility, trading volume, and the dynamic relationship between volume and volatility. These results are important in that they suggest a very complex relationship between the limit order book, volatility, and trading volume.

The persistence of volatility, or conditional variances, in stock returns has been known since Mandelbrot (1963). These features were an important part of the explosion of ARCH/GARCH models in the 1980's and beyond. The volatility modeling area of econometrics is enormous. The pioneering ARCH/GARCH papers are Engle (1982) and Bollerslev (1986). There are now many surveys of the literature including Andersen, Bollerslev, Christoffersen & Diebold (2013). In the past decade volatility attention has shifted to models built from realized volatility estimates using higher frequency time series to estimate a lower frequency volatility series.[8] This paper is primarily concerned with the mechanisms for generating persistent volatility. Since the introduction of subordinated processes in Clark (1973), and the mechanism of an outside business time clock, much of the literature has assumed some sort of exogenous event, or information timing system must be driving most of the changes in volatility, and concurrently, trading volume. Our paper presents a relatively simple, yet completely endogenous mechanism for volatility persistence based on adaptive agent behavior.[9]

Agent-based financial markets have long been shown to be a reliable generator of many of of these complex time series features. Over time they have become more sophisticated while empirical testing and validation have also made progress. The recent surveys by Dieci & He (2018) and Lux & Zwinkels (2018) contain many examples.[10] Many of these models are analyzed in great detail mathematically, and their replication of the empirical facts of asset returns is usually very robust across parameter sets, and model classes. It is essentially an accepted fact that fat tailed returns, and volatility persistence are a generic output for agent-based financial markets. It is interesting that few of the models make much progress in understanding how to turn off these features.[11] Their complex structure can sometimes mask exactly what is going on in terms of agent behavior.

In this paper we use an agent-based model, but it will follow in the spirit of models such as Gode

[7] Connections between trading volume and volatility have been documented for quite some time going back to at least Clark (1973) and Karpov (1987), and also Andersen (1994), Gillemot, Farmer & Lillo (2007), Ane & Geman (2000), and Tauchen & Pitts (1983).

[8] Two early examples are in Andersen, Bollerslev, Diebold & Labys (2003) and Barndorff-Nielsen & Shephard (2002).

[9] Two studies of the business clock hypothesis are Gillemot et al. (2007) and Ane & Geman (2000) that come to different conclusions about its success. In our model there is no business clock, so the mechanism of a simple subordinated clock related to information flow is ruled out.

[10] A few early examples of agent-based models matching various empirical features are Chiarella & He (2001), Gaunersdorfer & Hommes (2007), Lux (1997), and LeBaron (2001). A recent example which fits a very large number of empirical features is Staccioli & Napoletano (2021).

[11] An exception on this is LeBaron (2001) which shows that forcing agents to use very long times series when evaluating their adaptive behavior moves the market to an equilibrium where asset returns are essentially Gaussian white noise.

& Sunder (1993). These authors originated the idea of zero intelligence (ZI) agents. At the core of this approach is the principle that robust empirical features in socio-economic systems may be driven more by the structure of the institutions themselves than by specific details related to agent adaptive behavior. This idea is very much at the core of what we are trying to do here in that we build markets with our agents following relatively simple rules of thumb. In the area of limit order books there is some precedent for this approach. Farmer et al. (2005) is an example that views the order arrival process as a mechanical random flow. This paper, and some of the Physics inspired approaches are interesting for their ability to bring in more detailed mathematical structure to the limit order book, but some of them abstract away from any representation of an individual agent. Two more agent-based inspired examples are Chiarella & Iori (2002) and Ladley & Schenk-Hoppe (2009). We follow the first of these two papers in terms of utilizing many of their structures for agent design and market dynamics. Both show that with only a small amount of structure in agent trader/forecasters, a rich set of features emerge from the limit order book.

There have been other agent-based models that have extended some of this earlier work on limit order books. LeBaron & Yamamoto (2008) is an example which also builds off of Chiarella & Iori (2002), but adds an evolutionary/adaptive dynamic to the simple agent behavior. Chiarella, Iori & Perello (2009) extend their earlier model to incorporate a well defined share demand from a set of constant absolute risk aversion preferences (CARA).[12]

An important extension of this model is Yamamoto (2011) which adds a more sophisticated order placement system responding to depth on either side of the limit order book. With this addition it is able to generate a wide range of empirical time series features.[13] Chiarella, He & Wei (2015) introduces a model with a greater level of sophistication. Here, learning and adaptive agents try to utilize different aspects of the market state coming from the limit order book itself. This is a direct approach to understanding how agents might process the overwhelming information they are faced with in a limit order book trading situation. It is sort of a middle ground between the zero intelligence approach and more formal equilibrium models of an order book as in Goettler, Parlour & Rajan (2005). There is also an empirical approach to low intelligence agent modeling. Yang, Paddrick, Hayes, Todd, Kirilenko, Beling & Scherer (2012) show how agent behavior can be classified using market micro-structure data. This approach is an interesting balance between pure zero intelligence modeling, and more sophisticated agents.[14] It is also a bottom up empirical methodology that is consistent with the agent-based modeling philosophy.

Recently, agent-based approaches have also been used to better understand periods of intense market stress. Particularly those when there are few macro economic events driving the market dynamics. An example of this is the flash crash of May 2010. This pure liquidity event is very much in the spirit of our model where liquidity changes are completely endogenous. Several papers have been exploring

[12] Recent work by Isaac & Ramaswamy (2023 forthcoming) performs extensive replication and robustness experiments with the Chiarella & Iori (2002) model.

[13] Further agent-based models building off this foundation include Biondo (2020).

[14] It is similar in spirit to work on zero intelligence plus, or ZIP agents introduced in Cliff & Bruten (1997).

this model in settings with different styles of trading strategies brought together in agent-based models.[15]

Our approach blends many of these lines of research. We take the limit-order book structure, along with a low intelligence agent-based approach. Our calibrated market with relatively simple strategies generates near Gaussian white noise in the returns series. Modifying our agents' order shading strategy by a small amount completely changes our price/return dynamics, generating the rich set of features previously discussed. Also, it creates realistic, but completely endogenous pulsations in market liquidity. Lillo (2007) is an early approach to thinking about order shading, and detailed results there are consistent with our simple behavioral rules. Menkhoff, Osler & Schmeling (2010) and Zovko & Farmer (2002) report empirical evidence supportive of our strategies. Finally, Zhan & Friedman (2007) analyzes the impact of different levels of shading or markup in a simulated agent-based experimental setting.

3 Model description

3.1 Limit order book

The limit order book is defined by a price range, (p_{min}, p_{max}), and a price tick size, p_Δ. The order book accepts limit orders at the highest frequency time tick interval, (t), and orders are queued in time-priority at each price. Each order consists of an integer quantity (v), an owner identifier n, and an end-time t_{end} after which the order will no longer be executable. The book maintains a measure of the best bid, best ask, bid-ask spread, and midpoint of the price.

This is illustrated in a very simple book in Figure 1. On the ask side of the book (offers to sell) in red, the prices range from 100 to 102 in the left-most column. The next column represents the orders which have arrived the earliest, and will be the first to be executed at their respective price if a market order comes through. For example, at price 100 (and abusing notation slightly) the first order is the tuple $(1, t_{end}, i)$ – the quantity is 1, the end-time for the price is t_{end}, and the agent who is making this offer has the ID i.

The same holds for the bid side of the book (offers to buy), denoted in blue. Here prices range from 92-95, and the first order which would be executed has quantity 3, also ends at t_{end}, and is owned by the agent j. The best bid and best ask prices are 95 and 100, respectively; the bid-ask spread is 5, and the midpoint price (observed by all agents) is 97.5. If another limit order were to arrive at the best bid price, it would be added to the right-most side of the price queue at 95, while the order at the leftmost side of the queue is the first to be executed. Figure 1 displays a list of two additional bids both for quantities of 1. The best ask also contains two additional orders for quantities of 2 and 1.

[15] Several examples of this literature are Paddrik, Jr., Todd, Yang, Scherer & Beling (2012), Vuorenmaa & Wang (2014), Leal, Napoletano, Roventini & Fagiolo (2016), and Paulin, Calinescu & Wooldrich (2019).

Figure 1: Orderbook Example

For calculation of various statistics, a time period is denoted s, and Δ ticks of t produces one time period s. In each tick, N_{trade} agents are randomly selected to be given the opportunity to trade on the book.

3.2 Agent behavior

The basic model of agent behavior is a simplification from Chiarella & Iori (2002).[16] Agents follow one of two possible strategies to forecast the price over the next τ price ticks:

1. fundamental / mean reverting, and
2. noise trading.

The population is made up of σ_F fraction of fundamental traders, and $1 - \sigma_F$ fraction of noise traders. The price of the asset at time t is given by p_t. Returns over period h are defined as log differences,

$$r_{t+h} = \log(p_{t+h}) - \log(p_t). \tag{1}$$

3.2.1 Fundamental forecasts

The fundamental/mean reverting forecast relies on shared knowledge of a fundamental price, p_F, that all agents following the fundamental strategy believe. The fundamental return forecast for agent i is calculated as a one-tick-ahead forecast, given the current midpoint price p_t. Each forecast model runs a regression to predict the log deviation of the one period ahead price from the fundamental as in,

$$\log(p_{t+1}/p_F) - \bar{p} = \rho(\log(p_t/p_F) - \bar{p}). \tag{2}$$

[16] We simplify their model by eliminating the trend following strategies. This is conistent with results in Isaac & Ramaswamy (2023 forthcoming) who show that these strategies have only minor impact on the generated price dynamics.

Assuming stationarity around the fundamental, gives a long run mean of $\bar{p} = 0$, and we can derive,

$$\begin{aligned}
\log(p_{t+1}/p_F) &= \rho \log(p_t/p_F) \\
\log(p_{t+1}/p_F) - \log(p_t) &= \rho \log(p_t/p_F) - \log(p_t) \\
\log(p_{t+1}/p_F) - \log(p_t) &= (\rho - 1)\log(p_t) - \rho \log(p_F) \\
\log(p_{t+1}/p_t) &= (\rho - 1)\log(p_t) - \rho \log(p_F) + \log(p_F)
\end{aligned}$$

$$E(r_{t+1}) = (\rho - 1)\log(p_t/p_F). \tag{3}$$

If returns are normally distributed, then the expectation of the future price is given by,

$$E(p_{t+1}) = p_t e^{E(r_{t+1}) + 0.5 * \sigma_r^2}. \tag{4}$$

Assuming returns are approximately independent, the h period ahead forecast would be,

$$E(p_{t+h}) = p_t e^{h(E(r_{t+h}) + 0.5 * \sigma_r^2)}. \tag{5}$$

Since the regression as stated is just a regression of $\log(p_{t+1}/p_F)$ on $\log(p_t/p_F)$, ρ is the estimated first order autocorrelation of $\log(p_t/p_F)$ using a rolling band of the previous $10,000$ time steps. It is required to stay in the range $[\rho_L, \rho_H]$. Forecasters will be attempting to learn this parameter from the observed simulation time series. The parameter provides important information on the speed of price convergence back to the fundamental.

Finally, we introduce some small heterogeneity in the fundamental beliefs. Fundamental traders are divided into three groups, and each follows a slightly different value for p_F. The three values are defined in proportion to the central value as in $(0.95p_F, p_F, 1.05p_F)$. We will show that this adds slightly more realism, but it is not required for our main results.

3.2.2 Noise forecasts

The noise trading forecast simply forecasts a random noise return,

$$r_{t+1} \sim N(0, \sigma_e^2). \tag{6}$$

with a multi-horizon price forecast defined as for the fundamental agent. This noise shock is an aggregate shock and is not specific to any agent.

3.3 Agent trading behavior

Once agents forecast the return, it is projected out over one period and the corresponding price forecast is calculated as,

$$p^*_{t+1,i} = E(p_{t+1}) = p_t E(e^{r_{t+1}}) \tag{7}$$
$$= p_t e^{E(r_{t+1})+0.5\sigma_r^2}. \tag{8}$$

This can be extended out h periods as,

$$p^*_{t+h} = p_t e^{h(E(r_{t+h})+0.5\sigma_r^2)}. \tag{9}$$

Agents are making two critical, but false, assumptions here. They are assuming returns are both independent over time, and normally distributed.

Agents use these price forecasts to decide whether to buy or sell. Individual agents enter the market randomly and decide whether to buy or sell a single share of the asset. If $p^*_{t+h,i} > p_t$, the agent will be a buyer, and seeks to place a bid order. This market simplifies most strategic bidding behavior by assuming agents shade their bids below their current price forecast using a random strategy,

$$p^B_{t,i} = p^*_{t+h,i}(1 - k_i), \tag{10}$$

where $k_i \sim U(0, k_{max})$ independently, with $k_{max} \leq 1$.

If this bid crosses (is greater than) the current best (lowest) ask, then the orders are matched, and the buy and sell transaction is completed at the limit order price for that ask. If the bid falls below all the asks on the book, then a limit order for one share is generated and entered on the bid side of the book at p^B_t.

The random variable k_i determines the fraction they bid below their personal current valuation. The k_i value is fixed for each agent for the duration of the simulation. Some agents will always enter the market "aggressively" with respect to their own price expectations, while other "patient" agents will always enter the market far away from their price forecast, and always in their own favor. This is illustrated in figure 2.

When $p^*_{t+\tau,i} < p_t$ agents enter as sellers, and generate an ask price as,

$$p^A_{t,i} = p^*_{t+\tau,i}(1 + k_i). \tag{11}$$

They again enter this on the book unless it falls below the best bid in which case a transaction is executed.

Agents also execute market orders. With a fixed probability a given agent will decide to cross the book and execute their buy or sell order immediately at the best bid or ask. This probability, PMO,

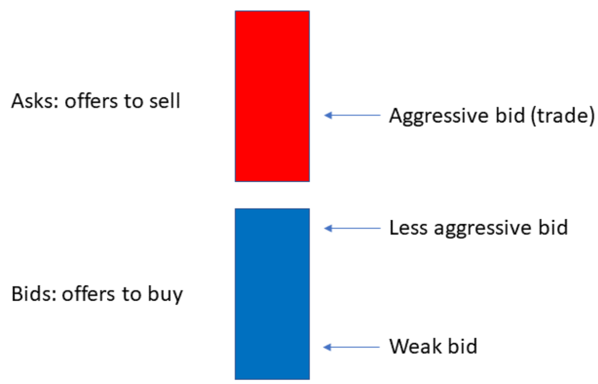

Figure 2: Aggressive vs Non-Aggressive Bidding

will be calibrated to align the fraction of executed orders with reasonable values from equity markets. The presence of market orders keeps market trades flowing even at times when the book might become thin with wide spreads. This could lead to a somewhat disfunctional market generating very few trades, and very little data.

After τ ticks, the order is automatically cancelled. We will see that this is a critical parameter for the model. In this way the basic market endogenously generates movements in liquidity and order placements in this highly stylized limit order setting. There is a small amount of agent intelligence directed at price forecasting, but none is directed at strategic order placement. Agents can be assumed to not see the book except for the prices of the best bid and ask. This is the basic structure in Chiarella & Iori (2002).

3.4 Agent behavior, extended: volatility persistence

Our key change to the original model is to modify the order placement strategies for the agents. We continue to use a mostly random strategy for buyers (sellers) trying to bid (ask) below (above) their forecast valuation. However, we assume that agents are aware of recent market volatility and adjust their strategies accordingly. Bidders now use,

$$\begin{aligned} p_t^B &= p_{t+h,i}^*(1-\widetilde{k}) \\ \widetilde{k} &= \frac{\sigma_S}{\sigma_L} k_i \\ \sigma_L &= \hat{\sigma}(\{r_t\}_0^t) \\ \sigma_S &= \hat{\sigma}(\{r_t\}_{t-v_{lag}}^t), \end{aligned} \tag{12}$$

where $\{r_t\}_{t_0}^t$ is the time series of realized price returns from t_0 to t, $\hat{\sigma}(x)$ is the sample standard deviation of the time series x, and v_{lag} is the lag length over which to calculate the short-term volatility. The same setup is repeated appropriately for all ask prices.

The σ_S and σ_L are measures of short term and long term volatility, respectively. σ_L is measured over the entire experimental period, while σ_S is measured over a much shorter rolling time horizon. Agents perceiving more near term volatility will move their orders farther from the current market. This behavior would be optimal for certain agents who are interested in maximizing trading profits for a constant probability of order execution. For higher volatility, they can take the risk of moving their orders farther away from the current market since they see an increased probability of being hit at a given distance from the last traded price.[17] They also may sense greater uncertainty, and higher probability of meeting better informed agents in the market when volatility is high. This also would push their orders off the current price in the same direction, driven by their fears of the presence of informed traders.[18]

4 Simulation Experiments

4.1 Parameter choices

As with any agent-based model we are faced with many parameter choices. The simplicity of this model keeps this set relatively small, but it still can be intimidating. We take a different approach from fitting our parameters to some chosen representative data set. There are several reasons for this. First, our aim is to recreate features which are common to many data sets in a robust fashion. We do not want to concentrate on what might be very specific characteristics which are not generally shared. Second, we want to gain a lot of understanding for how our parameters impact the model. We view each one as important, and want to understand their impact. Because of this we perform a very different type of parameter calibration exercise.

Although we are just choosing what we feel are the best parameters, and often running one simulation

[17] This is essentially an inventory strategy argument.

[18] In this case it is related to asymmetric information, adverse selection arguments. For theoretical background and references on both these arguments see Hasbrouck (2007). For a detailed model of an optimal order placement strategy see Lillo (2007). Also, see Menkhoff et al. (2010) and Zovko & Farmer (2002) for empirical evidence that traders do react to volaility in this way when placing orders.

with this set, we try to be very careful in our choices of parameter values. We divide our set of parameters into 4 subgroups which depend on our methods and intentions. These groups are all presented in table 1. The first set is a group we have calibrated empirically to general data features. In each of these cases the appendix presents a detailed sweep across different values. The first of these, σ_F, represents the fraction of fundamental traders (the remainder are noise traders). We will set this to 0.1. Larger values give us very strong mean reversion to the fundamental, and an unreasonably large negative return autocorrelation. Our second parameter is PMO, or the probability of a market order getting executed. We find that the value of 0.05 generates reasonable empirical patterns in prices and returns. It also gives us a rough fraction of order execution of 5 percent. Roughly 95 percent of our orders will get canceled. This fraction is not too far off actual values in limit order based markets. Our third parameter, kMax, is the max amount of order shading. This is a parameter which we have little intuition for what it should be, and there is no empirical evidence we can go to for locking it down. We simply sweep this and pull our value, 0.1, as one that gives us reasonable price/return dynamics. Our final critically calibrated parameter is Δ. For us this is the number of short trading periods that make up a trading day. We calibrate this parameter to make our daily returns appear close in terms of standard deviation and kurtosis to actual data. We get an annualized standard deviation of about 0.1 which is low for comparison to real data, but we think this is reasonable since our market is not subject to any new information, and fundamentals are constant.[19]

Our second set of parameters are critical to the dynamics of the model. Our base case simulations concentrate on certain values, but in a later section we will demonstrate that moving these around can have a strong impact on our results. The first of these, τ, is the number of short periods that an order is allowed to remain on the book until being canceled. We generally set this value to 25, or half a day, but increasing this value can change our volatility dynamics significantly. It is swept in the appendix. The next two parameters are booleans controlling the type of order shading (static or dynamic), and whether local volatility estimates are trimmed. The former is critical to our results, and is our basic extension to most simple limit order book models. The second, "volatility trim", is also important, and will be discussed further, as will be the amount of trimming, given by "Volatility trim level". Positive and negative returns are trimmed to a value given by $2\sigma_L$ or two times the long range return standard deviation. Finally, the level of noise, σ_e, is also swept in the appendix. The level of noise appears to have a very dramatic impact. Too little, and interesting behavior and trading seems to disappear. Too large, and the noise dominates the market.

Our next set of parameters are all self referential, and calibrated to be consistent with the market simulations. One can think of these as being nearly endogenous to the market. The first two are concerned with the parameter ρ which represents the speed of mean reversion that agents perceive. This is a critical behavioral parameter for the model. As discussed in the last section, agents will be estimating ρ to the data they are observing from the generated time series in real time. To maintain

[19] One could probably make the case that our standard deviations should be even lower, but we have no examples of market data with no new information arriving to compare with.

market stability we will bound ρ to fall into the set $[\rho_L, \rho_H]$, and these values are given in the table. Agents make price forecasts several short time steps into the future, h. This value is set to the mean length of time an order spends on the book before eventually being executed. Canceled orders are not considered in this estimate. We have found a rough connection of orders executing at about $(1/3)\tau$ short periods. This would correspond to agents forecasting out about 1/6 of a day into the future, so even our fundamental traders are very short term in nature. Although this parameter could be dynamically updated with real time market data we chose the simpler method of setting it to a value from our market, and then making sure that in future runs this value is self-consistent with simulation output.

Our final set of parameters are the number of traders entering each subperiod, N_{trade}, and the lag used in our short term volatility estimates, v_{Lag} (in short time units). Both of these parameters do not have a great impact on our results, and we have not added large scale sweeps for them. v_{Lag} does impact the level of persistence we see in our absolute return autocorrelations. Also, as we will see, the model does not generate as much volatility long memory as real markets do. We will perform some small scale simulations with heterogeneity across values of v_{Lag}. The last parameter, p_F, is the fundamental value which is set to 1000. Moving this around has no impact on the results unless we were to move it very close to zero.

Parameter	Value
σ_F	0.1
PMO	0.05
kMax	0.1
Δ	50
τ	25
Dynamic shading	**True**
Volatility trim	**False**
Volatility trim level	$2\sigma(r_t)$
σ_e	0.0001
ρ_L	0.98
ρ_H	0.999
Forecast horizon (h)	$(1/3)\tau$
N_{trade}	30
v_{Lag}	150
p_F	1000

Table 1: Model parameters

4.2 Fixed order shading

In this section we present results for simulated daily returns using our benchmark parameter set. We will be using the fixed order shading strategy, with constant values of k_i for each agent. The four panel plot in figure 3 displays several key features from the model. First, the upper left panel plots the returns series estimated as log first differences of the simulated prices. Returns show little persistence and no obvious visual patterns. A histogram of these returns is shown in the upper right hand panel superimposed with a standard Gaussian distribution. The returns look very close to Gaussian with an estimated kurtosis of 3.4 which is close to 3, or the target value for a Gaussian.

The lower left panel plots the corresponding price series. This displays a somewhat complex time series. In the long term it is stationary and not going too far away from 1000 which is the fundamental target in the simulations. However, it does display some very persistent deviations from the fundamental which is evident in the picture. This persistent behavior is confirmed with the estimated autocorrelation of 0.988. Prices are not a random walk, but are very persistent.

Finally, the lower right panel searches for structure in the returns, reporting the autocorrelations in returns and absolute returns. It reveals little autocorrelation in either. There is a small amount of negative autocorrelation in returns, but the estimated value at the first lag is only about -0.02. There is almost no autocorrelation in the absolute returns series.

Our results in this figure show that the model is able to take basic random behaviors and reasonable limit order book dynamics, and convert this to a set of results which replicate some features of financial markets, but miss others in a critical way. The counterfactual behavior can be seen in the right column of the figure. Daily returns are not Gaussian, and exhibit excess kurtosis. Also, while returns themselves show little autocorrelation, volatility is typically persistent, which usually appears through positive autocorrelation in absolute returns.

As mentioned previously our simplest version of the model involves some heterogeneity across fundamental traders. There are three groups of traders who persistently believe in fundamental values of $(0.95p_F, p_F, 1.05p_F)$. In figure 4 we repeat our previous graphs for the case of a single fundamental type holding beliefs of p_F only. The figure recreates the previous picture almost exactly in 2 out of the 4 panels. The only key difference is in the lower panels. In the left panel asset prices now mean revert to the fundamental much more quickly which is visually apparent. The quantitative near random walk behavior which was shown in the previous plots is gone with a price autocorrelation of 0.68. A closely related result is shown in the lower right panel where we see a strong negative autocorrelation in the returns series. We view both of these as unrealistic, and will therefore stay with the slight amount of heterogeneity in agents' fundamental views.

We will now explore small modifications to this model that generate more realistic dynamics. In particular, returns will exhibit typical fat tailed distributions and persistent volatility.

4.3 Dynamic order shading: Base case model

We now take our exact specifications from the previous section and make one change. We include dynamic volatility shading, namely the addition of $\frac{\sigma_S}{\sigma_L}$ in the agent's bid (ask) price determination as previously described,

$$p_t^b = p^*_{t+\tau,i}(1 - \widetilde{k}).$$

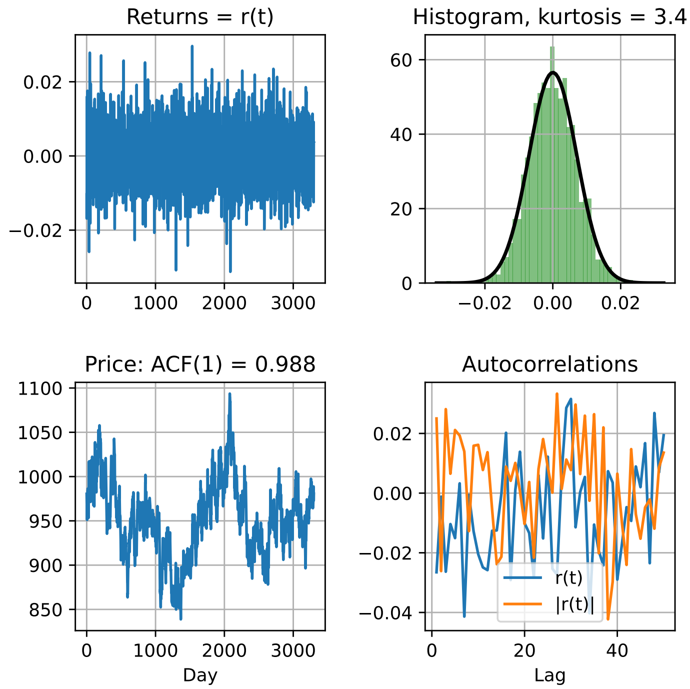

Figure 3: Fixed order shading

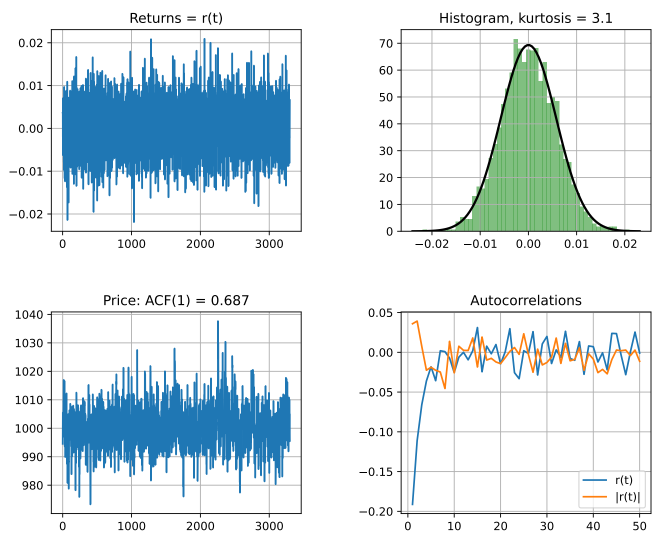

Figure 4: Fixed order shading: Homogeneous fundamentals

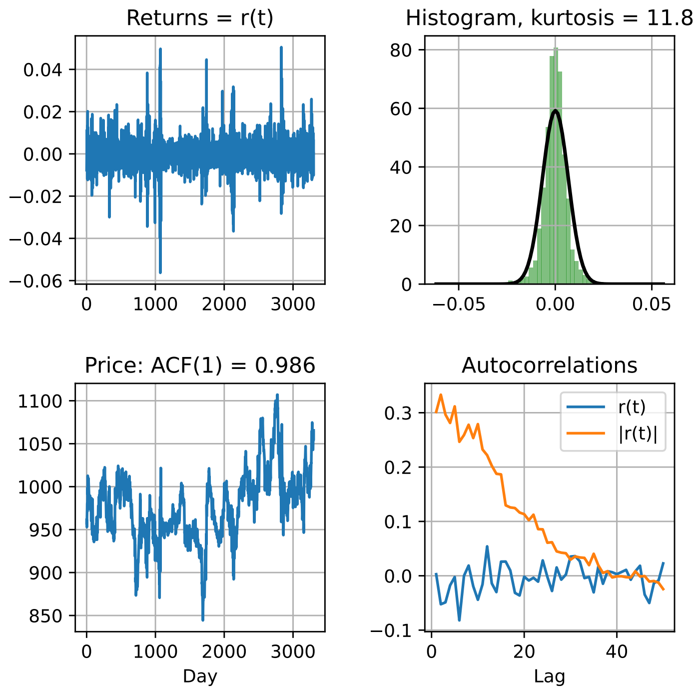

Figure 5: Dynamic order shading: Base case model

Figure 5 presents a similar four panel plot exploring the prices and returns in this situation. The changes are dramatic and revealing. The return series in the upper left panel shows more pronounced extreme values along with pockets of high and low volatility. The return histogram in the upper right panel shows strong leptokurtosis (fat tails), along with a consistent kurtosis estimate of 11.8. In the lower left, prices again display persistent deviations from fundamentals with a slightly larger autocorrelation of 0.986. Finally, the return autocorrelations remain close to zero, but the absolute return correlations are positive, and very persistent. All these features are consistent with actual financial time series.

The next figure, figure 6, repeats all the previous plots for daily returns from the S&P500 ETF (SPY). This a security which is designed to trade like an individual stock, yet also track the overall market. We are sampling this daily from June 2008 through May 2022.[20] For three of the panels the results are similar to the previous figure. One exception is the price series in the lower left panel which is nonstationary, and not converging to any constant value. This is obvious for stock prices, and is consistent with steady growth in underlying fundamentals. One could also take issue with the lower right panel where the long term persistence of volatility is even stronger than in the simulation. This is characteristic of a possible long memory process in volatility which exhibits hyperbolic as opposed to exponential decay in the autocorrelations. We will look at this in more

[20]We use a shorter sample here to line up with the daily length of 3500 in our simulations. For our detailed empirical tests we extend this to start in 1993, and also extend the simulations to 7000 days.

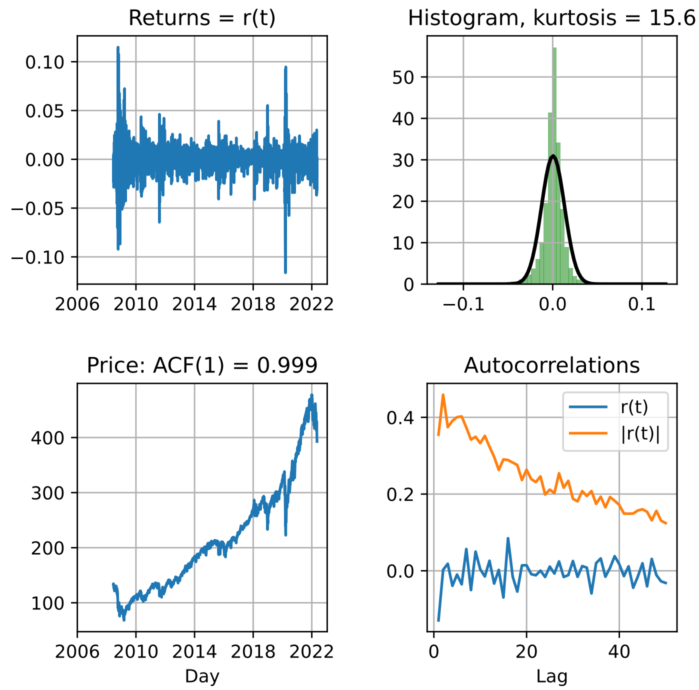

Figure 6: SPY Returns 2008-2022

detail in the next section.

4.4 Stylized Facts

This section looks at some of the empirical details for how close this model comes to the features just shown graphically. We examine three primary price/return features: deviations from fundamentals and random walk behavior, excess kurtosis, and volatility autocorrelation.

4.4.1 Random walk

One of the fundamental concepts of efficient market theory is the martingale property for asset returns arguing that previous returns are uncorrelated with future returns. Empirical studies have not always agreed on this feature. Many older studies have concluded that asset returns exhibit no autocorrelation or slight positive serial autocorrelation, Cont (2001). However, more recent empirical studies of major world indices have found weak, but statistically significant negative autocorrelations, Baltussen, van Bekkum & Da (2019).

We have already demonstrated that our model generates near uncorrelated return series, so we will look into this a little more deeply here. In figure 7 we examine the first order autocorrelation at various intraday time horizons.[21] Since we have calibrated to 50 short periods equalling one day, the

[21]For all results in this section we are using longer sample sizes. Our SPY series now starts in February 1993, and

figure shows returns at various shorter time horizons. One period here would correspond to roughly 6.5 hours/50, or about 7.8 minutes. The SPY comparison series is measured at 10 minute intervals. Autocorrelations at micro structure levels often turn negative, and the figure makes it clear that this is a robust feature for both our simulation and the data. Most of this negative correlation comes from trading moving back and forth between the best bid and ask (bid/ask bounce). Eventually, the impact of this institutional feature dissipates at longer horizons, and the autocorrelations approach zero. Our simulations show a slightly larger negative autocorrelation which is probably due to wider spreads and lower trading volume. However, our alignment with this data here is very good.

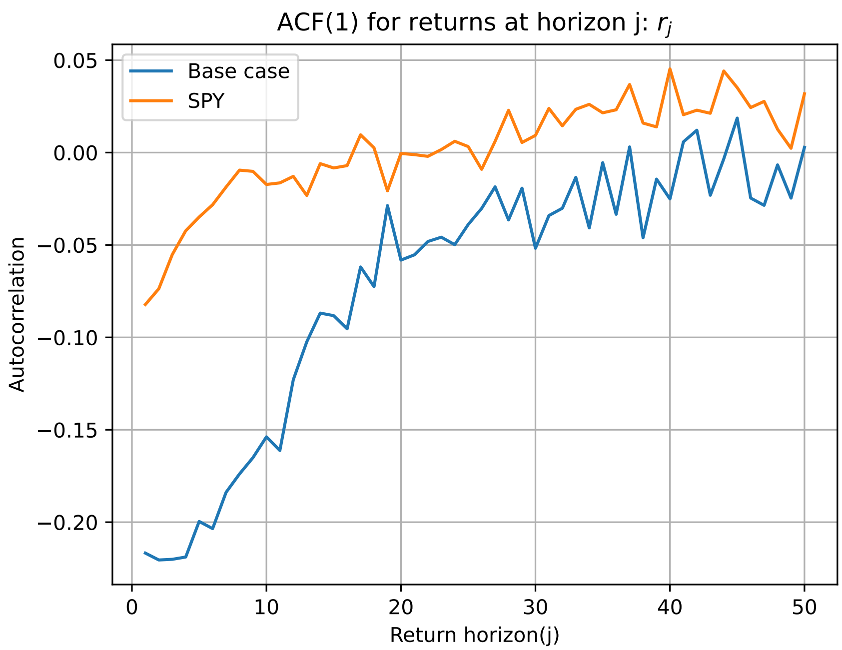

Figure 7: Return autocorrelation (1 lag) at varying horizons (j=50, 1 Day)

4.4.2 Heavy Tails and Aggregate Gaussianity

Asset returns calculated over higher frequencies (daily or intraday) tend to follow a non-Gaussian distribution. The distribution of daily returns exhibit heavy tails, especially on the left side of the distribution, Mandelbrot (1963). There is no general consensus on the exact form of the tails, but a kurtosis value > 3 (Gaussian) is often observed. We estimate it with,

$$\text{Kurt}[r_t] = E[(\frac{r_t - \mu}{\sigma})^4] = \frac{1/n \sum_{t=1}^{n} (r_t - \bar{r})^4}{[(1/n) \sum_{t=1}^{n} (r_t - \bar{r})^2]^2}.$$

Empirically, as returns are calculated at less granular time scales, the distribution tends to converge

the simulations are extended out to 7500 days to align with the sample length of the SPY series, 7373.

to an approximately normal distribution (Cont. 1997, Chakrabortia et al 2010). This stylized fact is referred to as "Aggregational Gaussianity". To test our results we utilize the Kolmogorov-Smirnov (K-S) test to compare return distributions at different time scales to a normal distribution in table 2. The K-S Test Statistic D is given by

$$D = \sup_x |F_n(x) - F(x)|,$$

where $F_n(x)$ is the CDF of returns from the simulation, and $F(x)$ is the CDF of a normal distribution. The values indicate that our simulation results, where returns are calculated at a daily, 2-day, or weekly frequency, exhibit statistically significant differences from the normal distribution. However, as the scale becomes less granular, returns appear to approach a Gaussian distribution. The table also displays a similar pattern for the SPY returns.

K-S p-Value		
Returns Frequency	Sim p-Value	SPY p-Value
Daily	<0.0001	<0.0001
2-Day	<0.0001	<0.0001
Weekly	0.0161	<0.0001
Bi-Weekly	0.1062	0.0370
Monthly	0.2938	0.2128

Table 2: Distributional aggregation

4.4.3 Volatility properties

In addition to simple volatility clustering discussed earlier, it has been well documented that the function of absolute returns decays slowly as a function of the time lag, roughly as a power law. The exponent of the power law varies considerably, but studies typically claim that the range is between $[0.2, 0.5]$.[22] Visually, the simulated returns from our model display this feature. Using high frequency data down sampled to daily frequency we can get a more detailed picture of this persistence by examining autocorrelations of realized volatility measures. Realized volatility is an important method for getting more detailed information on the changes and persistence of volatility.[23]

We first define a realized volatility estimate for day t as

$$RV_t = \sum_{h=1}^{H} r_{t,h}^2 \tag{13}$$

[22]See Cont (2001), Liu, Cizeau, Meyer, Peng & Stanley (1997), Chakraborti, Toke, Patriarca & Abergel (2011).

[23]The use of high frequency data to get more accurate volatility time series at lower frequencies goes back at least to French, Schwert & Stambaugh (1987). The modern implementation, and detailed theoretical developments start with Andersen et al. (2003) and Barndorff-Nielsen & Shephard (2002). A good comparison application to stock returns is Andersen, Bollerslev, Diebold & Ebens (2001). There have been many excellent surveys on the topic including, Andersen et al. (2013). An excellent textbook treatment is contained in Taylor (2005). There is also a zero-intelligence agent approach to testing various realized volatility estimators in Gatheral & Oomen (2010).

where $r_{t,h}$ is the log return over the intraday subperiod, h on day t. Assume there are H of these subperiods in a day. For our SPY data we will use 5 minute time slices. For our simulations we use our high frequency trading period which we have aggregated to daily horizons using $H = 50$. We will also follow most of the literature by looking at the log of the RV process,

$$rv_t = \log(RV_t) \tag{14}$$

We check two of the simplest stylized facts for realized volatility. First, standardized returns are estimated by taking daily returns, and dividing by the square-root of the daily RV_t value. These normalized returns have been shown to follow a near Gaussian distribution. The left panels of figures 8 and 9 display the distributions for the SPY and our base simulations respectively. In both cases we have transformed what was a nonnormal return distribution to something very close to normal. The second fact is that the log of RV_t is also nearly normally distributed itself. The right panels in these two figures display this for the SPY and base model respectively. In both cases it is relatively close to normal, but the simulation data still appears somewhat skewed and yields a kurtosis of 4.1.

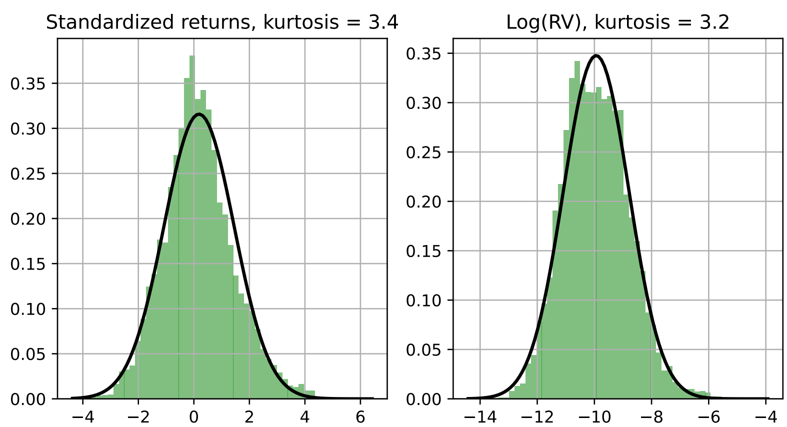

Figure 8: Realized volatility distribution facts (SPY)

Volatility, represented by either absolute returns or realized volatility shows very long scale persistence and slowly decaying autocorrelations. These are consistent with long memory processes, and we will perform some basic tests on our data to see how well it lines up with our comparison SPY series.

True long memory processes exhibit very important scaling features. One is that the variance over

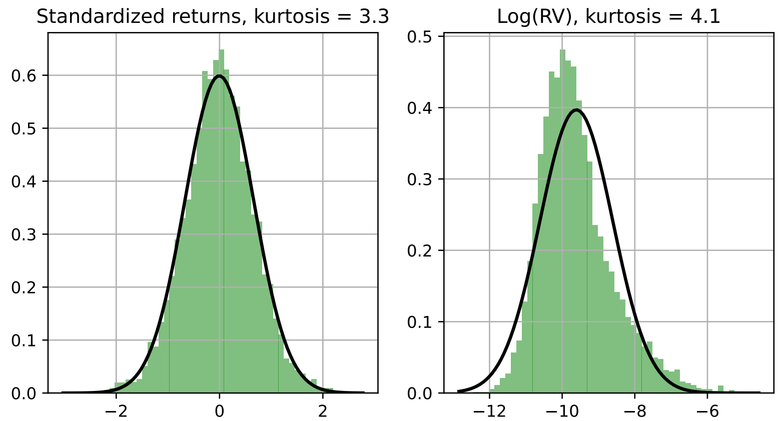

Figure 9: Realized volatility distribution facts (Base case)

nonoverlapping sums will scale as a power law,

$$\mathrm{Var}(\sum_{k=1}^{K} rv_{t+k}) \propto K^{2H} \tag{15}$$

Plotted in log/log space the variances should yield a straight line with slope $\beta = 2H$, where H is the self similarity parameter.[24] Other long memory parameters can be determined from H. The fractional difference parameter, d, is given by

$$d = H - \frac{1}{2} = \frac{1}{2}(\beta - 1). \tag{16}$$

Finally, autocorrelations will decay as a hyperbolic,

$$\rho_k \propto k^{-\gamma} \tag{17}$$

where

$$\gamma = 2 - 2H. \tag{18}$$

An independent process is represented by $H = 0.5$, with variances scaling in proportion to the sample lengths, K, and $d = 0$.

One simple estimate of H is to use scaling plots of the variance sums, where the slopes are given by $2H$. The upper panel in figure 10 shows a variance scaling plot for log realized volatility in our base simulation case. The range for aggregation goes from 1 to 75 days. The estimated slope in this region is 1.75 which gives a values of $H = 0.874, d = 0.37, \gamma = 0.25$. The upper panel also displays a scaling example for a single redraw (with replacement) of the original rv_t series to show that its slope is far away from the actual series, and very close to 1 as it should be.

[24]The core definitions for long memory processes are contained in Beran (1994).

The lower panel plots the autocorrelations for rv_t along with several theoretical candidates. It shows the expected long positive persistence in autocorrelations in the simulation with values remaining positive even at 75 day lags. An AR(1) is fit, and its ACF drops to zero quickly. The long memory autocorrelations use the hyperbolic decay given above along with γ. It is lined up to match the data at the lag of 32, or halfway through the range. It is better for fitting the data, but is low for the early lags, and high for the later ones. The best fit here comes from an ARMA(1,1) fit to the rv_t time series. This corresponds to a model with latent volatility following a highly persistent AR(1) process, but obscured by some amount of noise. It is interesting that this relatively simple model, with a slow, but exponential, decay pattern performs the best here.

Figure 11 repeats the previous figure for the SPY series. Here we see a slope of 1.88 which gives $H = 0.94, d = 0.44, \gamma = 0.12$. The lower panel again shows the ACF for the rv_t series along with the comparison values. In this case both the AR(1) and ARMA(1,1) fall much more quickly than the data. However, the long memory appears to be a very good depiction of the ACF decay in this case.

The base model generates very long persistence in volatility, but still not on the order of that seen in the data. Also, the SPY data appears closer to a true long memory process. There could be many reasons for this. First, our model consists of a single type of agent when it comes to volatility based order shading. There is only a single value for v_{lag}. Possibly using heterogeneous values for this might add more realistic patterns. In the next two figures we do a quick exploration of this concept.

We make our agents' volatility forecasting models heterogeneous. Agents draw a forecasting horizon from a finite set given by $[0.3, 1.5, 3, 6, 12, 24, 48, 96]$ days. They are uniformly distributed across these horizons. The previous experiments are a subset of this, putting all weight on the 3 day horizon.[25] Figure 12 repeats our long memory experiments from figure 10 for the heterogeneous forecasters experiment. The upper panel shows our scaling exponent to be very close to the SPY data, and the lower panel displays an autocorrelation pattern which is now much closer to the fitted long memory model, and does not line up well with the two candidate short memory processes. A more detailed comparison is performed in figure 13 which compares the autocorrelations for several different homogeneous forecast horizons, the heterogeneous model, and also the SPY. It is clear that the heterogeneous simulation lines up well with the SPY data. Also, interesting is the fact that none of the homogeneous experiments come close to looking like the real data. This makes it clear that just extending the forecast horizon beyond the original 3 days is not enough to replicate long memory volatility as seen in the SPY. Even the very long horizon experiment generates a very persistent autocorrelation, but it is nearly constant at a value near 0.25. It clearly does not replicate the data well. Heterogeneity in this one parameter appears enough to match up with the long memory in the data.

Finally, it is well known that discrete changes in time series can also appear as long memory. The

[25]Our choice of horizons is loosely based on powers of 2, but not exactly. We have found these parameters through casaul experimentation, and they are not estimated. Our goal here is to show that this is a plausable mechanims for long memory.

SPY values contain some distinct short periods of relatively high volatility. Examples would be the 2008 financial crisis, and the COVID crisis. These may end up showing up as more extreme persistence in the real data. The simulation contains no discrete jumps like this. Deeper testing on all these mechanisms is something we intend to do in figure research.

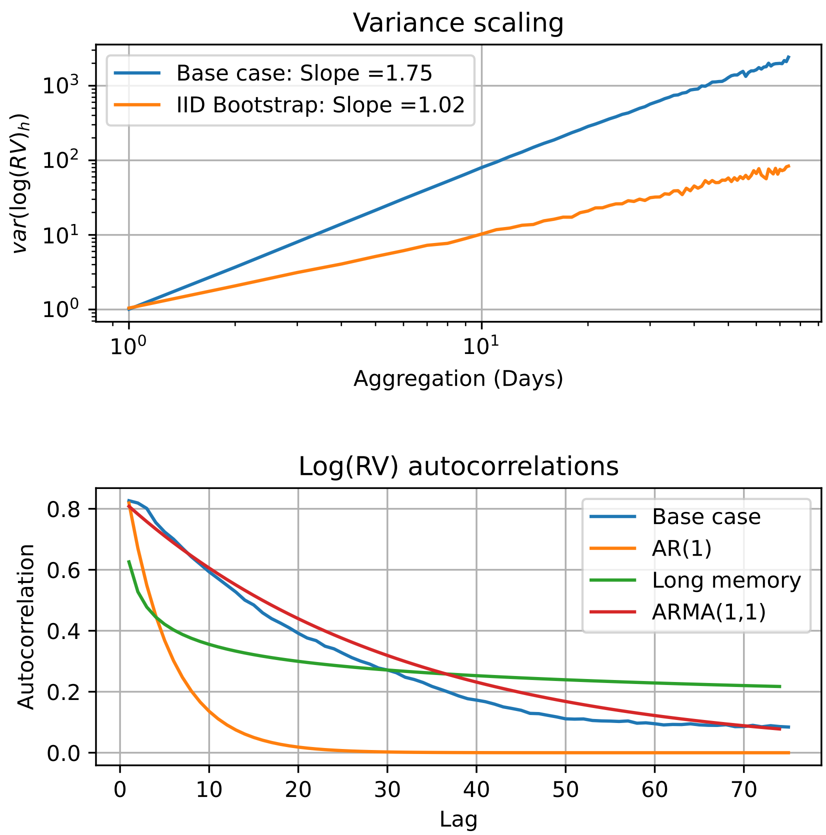

Figure 10: Long memory tests (Base case)

5 Shutting down volatility persistence

We now take the modified model with the volatility adjustment as our benchmark and demonstrate that we can shut down the volatility persistence by changing other market parameters. We explore three main experiments below: extending the time orders stay on the book, trimming agent volatility

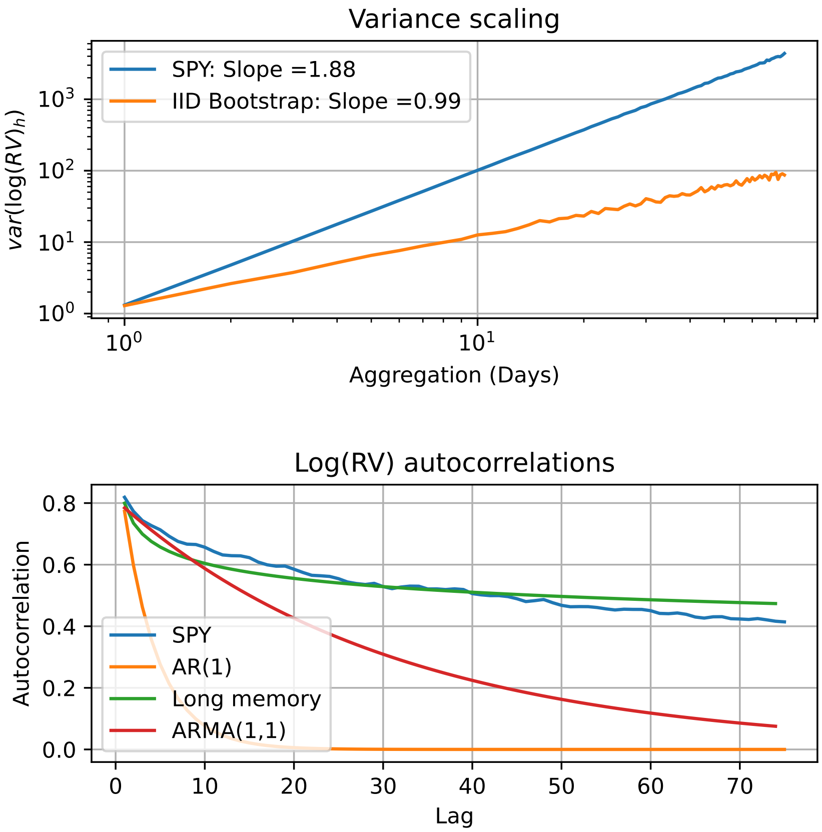

Figure 11: Long memory tests (SPY)

estimates, and increasing the level of noise for the noise traders.[26]

5.1 Reduce rate of order cancelations

Our first experiment demonstrates that increasing market liquidity and depth can eliminate the volatility persistence effect. We do this by changing the order cancellation policy. The model allows orders to stay on the book for a fixed amount of time, related to the forecast horizon of the agents. Agents do not actively cancel their orders as the world evolves. Rather the lifetime for an order is a parameter under our control. Increasing this value allows orders to stay on the book longer and increases liquidity in the market. Increasing this value by a large enough amount will eliminate the volatility persistence effect altogether. This gives a picture of the core dynamic that is necessary between agent order placement strategies and the actual order density on the book. As long as there is some scarceness to liquidity, when agents begin to place orders farther from the market midpoints, the density of the book drops, leading to a reduction in near midpoint depth, and increasing price impact for a given trade size. This has to lead to an increase in market volatility as it processes

[26] Using a more complex, but well calibrated, model, Biondo (2019) explores several similar policy interventions.

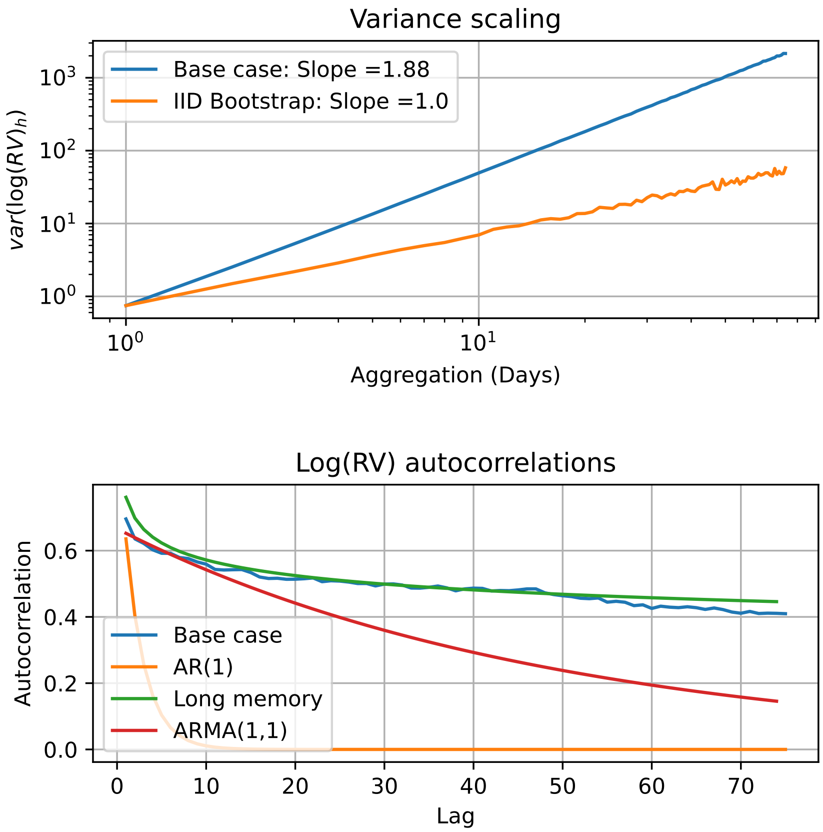

Figure 12: Long memory tests (Base case: Heterogeneous agents)

new incoming trade demands. There will then be a feedback from this increased volatility into the agents' strategies, leading to more order dispersion and a magnification of volatility.

This situation will only occur when there are enough order cancellations. If the book remains fully loaded with orders, then a kind of "wall of liquidity" will protect the market from increases in volatility. We show that we can raise and lower this wall under our control. Figure 14 demonstrates the effects of changing the order cancellation time from 25 ticks (half a day) to 75 ticks (1.5 days). Volatility clustering noticeably disappears from the time series as the "wall of liquidity" cannot be broken down. Figure 29 in the appendix displays results for $\tau \in [10, 25, 50, 250]$. One can observe for values of τ that are too low, the returns become volatile because a dense order book can never develop. Our model relies on the dynamics from both sparse and dense order books. The τ parameter is essential in ensuring these dynamics take place.

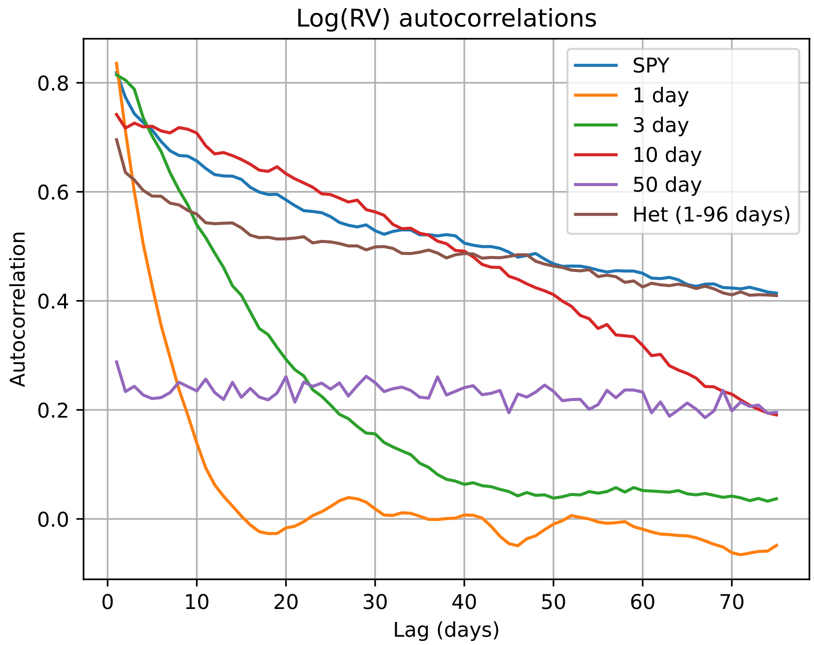

Figure 13: Autocorrelation comparisons

5.2 Trimmed volatility estimates

Our second experiment involves the learning behavior of agents. How are they processing volatility, and could they be slightly confused about it?[27] Given our original base case results from the first section we know that the model generated excess kurtosis in the returns series. Could these return spikes confuse agents to perceive market jumps as increases in the volatility of the overall price diffusion?[28]

To control for the impact of extreme returns we replace the simple variance estimate for short term volatility with one using returns which are trimmed to a multiple of the long-term volatility. We modify equation (12) by changing the value for σ_S as follows:

[27] See Benhabib & Dave (2014) for some early examples of the difficulty of learning in time series subject to large deviations.

[28] Some early econometrics for sorting between jumps and diffusion volatility would be Barndorff-Nielsen & Shephard (2004).

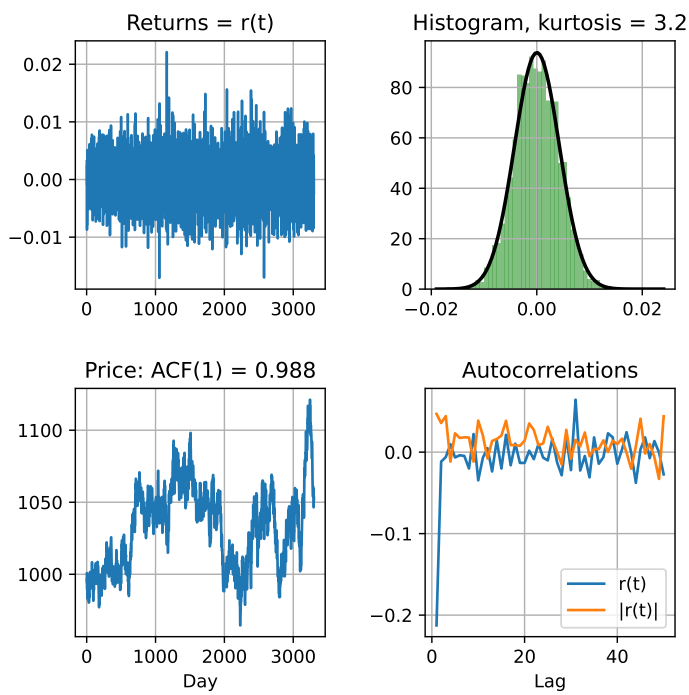

Figure 14: Reduced Order Cancelations

$$
\begin{aligned}
p_t^B &= p^*_{t+\tau,i}(1-\widetilde{k}) \\
\widetilde{k} &= \frac{\sigma_S}{\sigma_L}k \\
\sigma_L &= \hat{\sigma}(\{r_t\}_0^t) \\
\sigma_S &= \hat{\sigma}(\{\hat{r}_t\}_{t-v_{lag}}^t) \text{ where,} \\
\hat{r}_t &= \max\{-2\sigma_L, \min\{2\sigma_L, r_t\}\} \ \ \forall t \in [t-v_{lag}, t].
\end{aligned} \tag{19}
$$

Thus the returns are trimmed to within 2 standard deviations of zero. We find that simply trimming this volatility estimate stops the persistent volatility dynamics. This suggests an interesting connection between return jumps and persistent volatility. Jumps confuse traders, and they filter some of this into volatility estimates and therefore future volatility. Shutting down this channel again removes the volatility persistence from the return series. Figure 15, demonstrates the effects of having agents trim their volatility estimates. Once again, volatility clustering noticeably disappears from the time series while other stylized facts are not as affected.

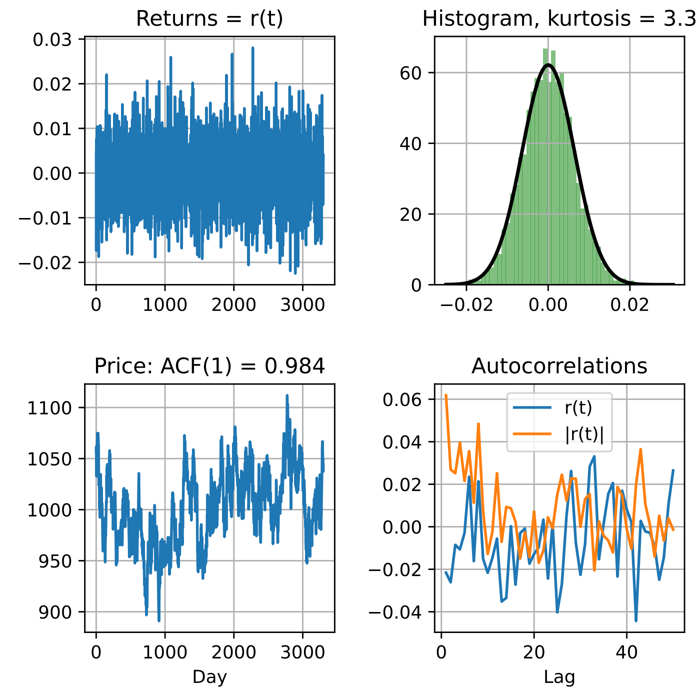

Figure 15: Trimmed Volatility Estimates

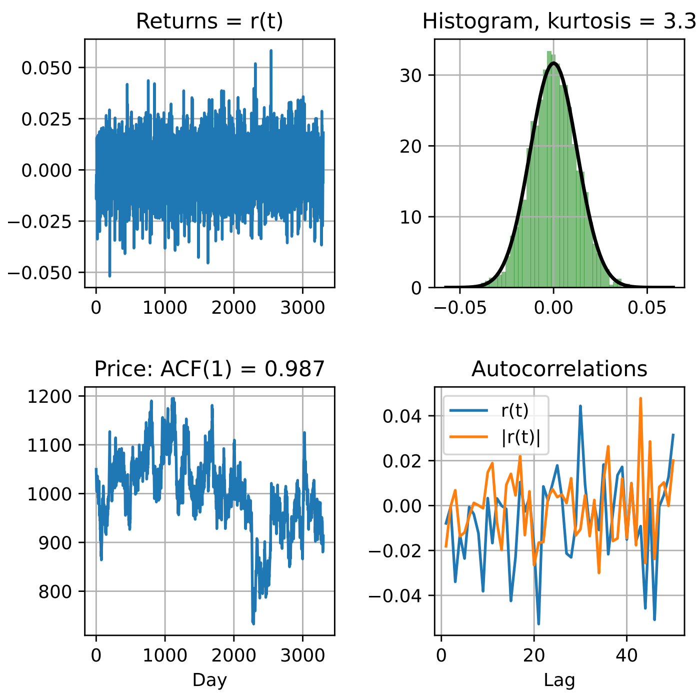

Figure 16: Increased Noise Trading

5.3 Increased noise trading

Finally, we are able to change the volatility characteristics by increasing the noise level. If we set the noise level to $\sigma_e = 0.001$, a factor of 10 over the highest noise level we use in our experiments, we can increase the variation in the noise trader valuations (equation 6). With agents experiencing a very high level of noise-driven activity in the order book, the effects of shading appear to be overwhelmed, and volatility persistence does not take hold. In figure 16 we see the impact of increased noise trading on our results. Most of our interesting return features have again been removed.

6 Endogenous liquidity and connections to the order book

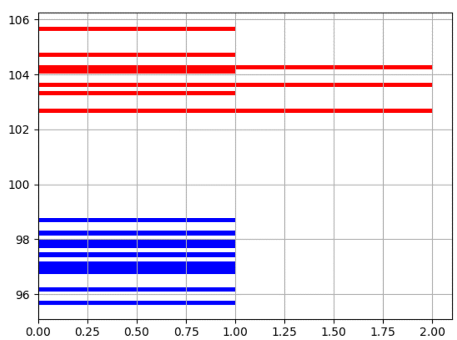

Figure 17: Sparse, Illiquid Limit Order Book

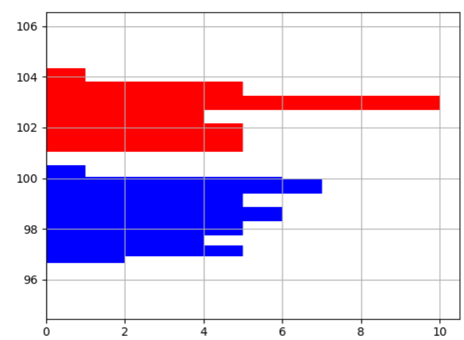

Figure 18: Dense, Liquid Limit Order Book

In this section we will demonstrate how our markets endogenously move through periods of high and low liquidity as shown in some basic features of the limit order book. We say that the limit order book is sparse when there are many gaps at price points between orders. Such a limit order book cannot support a high volume of order flow without large jumps in price. A dense order book is the opposite: every price tick has multiple limit orders stacked. A high volume of order flow in a dense order book will induce a smooth change in price. Figure 17 visualizes a sparse orderbook. One can observe that as market orders are filled from 102 to 104, the price will change suddenly for relatively small volume of orders. Conversely, figure 18 visualizes a dense orderbook; here price will move slower and more smoothly for the same volume that would move quickly through the sparse book.

It has been documented in Naes & Skjeltorp (2006) (NS) that, there are distinct relationships between the shape of the order book and volatility. Order books are typically "V" shaped around the bid-ask spread. The orders near the bid/ask spread are executed much quicker than orders further away. Therefore, buy(sell) orders far below(above) the best bid(ask) tend to accumulate on the order book. Order accumulation allows for the calculation of order book slope based upon the density plots of bids and asks. Ask slopes are typically positive, while bid slopes are typically negative. One of their major findings is that a flat order book corresponds to greater volatility, and flatter books are typically indicative of a less liquid order book. We use one of several different slope estimates used in their paper. Bid slope at time t is calculated in the following manor:

1. Define the minimum bid (minBid) as the lowest bid currently on the book.
2. Define the maximum bid (maxBid) as the largest bid (at the bid/ask spread).
3. Divide a range of n prices from maxBid to minBid, $p_1, \ldots, p_n$. p_i is decreasing over i.
4. Find the cumulative volume on the book from p_1 to p_n, and place in V_n.
   (a) V_1 is the number of orders at p_1. V_2 is the sum of the orders at p_1 and p_2.
   (b) V_n are the total orders in the range $1, \ldots, i$.
   (c) V_i is nondecreasing in i.
5. Normalize V_i by dividing by total volume as in $v_i = V_i/V_n$.
6. Finally, define slope as:

$$\text{slope} = \frac{1}{n-1}\sum_{i=1}^{n-1}\frac{v_{i+1}-v_i}{p_{i+1}/p_i - 1}. \tag{20}$$

This defines the bid slope at a fixed time snapshot. This is negative since p_i is decreasing, and volume is increasing. The ask slope is calculated in exactly the same fashion, but will be positive since the ask prices will be increasing. Total orderbook slope is calculated as:

$$slope = 0.5(|BidSlope| + |AskSlope|)$$

We follow NS by averaging the slope for each day (50 periods in our model).

Figure 19 presents panels comparing various measures of the order book status along with the slope. The top panel shows a plot of our daily returns over roughly 3500 daily periods. This clearly shows our typical periods of low and high volatility. The second panel displays the mean bid/ask spread over each day. It shows a pattern of widening during the periods of high volatility which is consistent with limit order books. Our next panel shows the estimate of slope on the book. Consistent with NS we see the slope being largest during periods of low volatility. This means that these are periods with a steep and dense limit order book keeping trading close to the bid/ask spread. When the book flattens and becomes more sparse, slope falls, and volatility increases. A final measure of market liquidity is shown in the lower panel. This shows the market depth measured as the number of orders close to the bid/ask boundary.[29] This shows a large dense book when return volatility is low, and a much sparser book with lower depth when volatility is high. Figure 20 summarizes the connections from the last figure in a series of cross correlations. The figure shows the strong negative correlations between slope and depth and volatility along with the positive correlation between spreads and volatility.

[29]This sums the orders within σ_L, one return standard deviation, of the (best bid+best ask)/2.

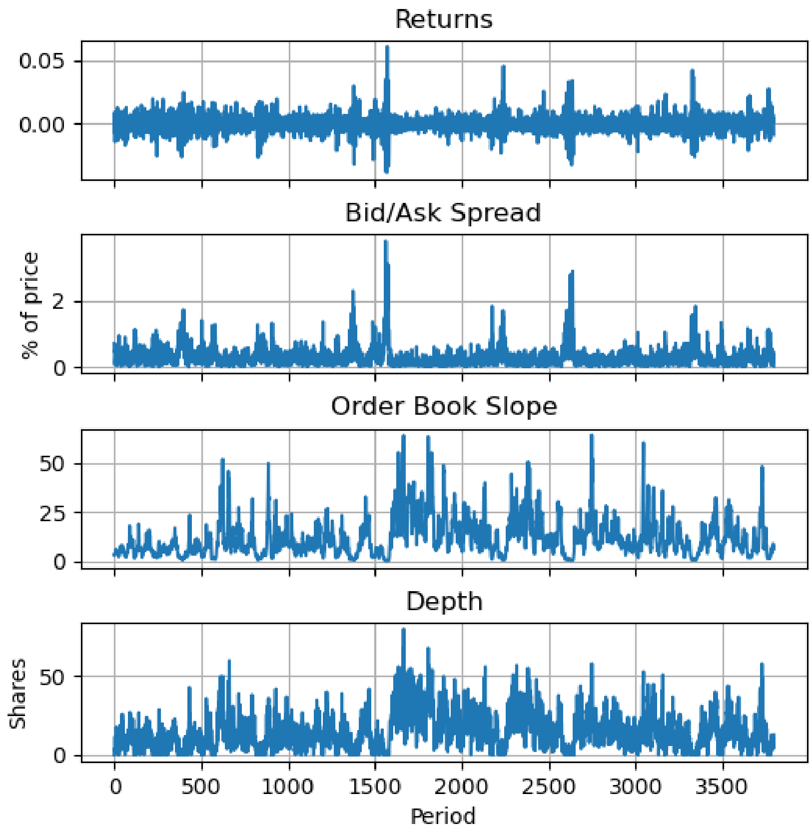

Figure 19: Order Book Dynamics (Base case)

Figure 21 presents some evidence for the correlations involving trading volume. We see that steep slopes are associated with high trading volume. NS find an overall negative relationship between volume and slopes, where flat books are connected to relatively high trading activity. However, they document that this relationship is more complicated. When they measure the slope on the parts of the book closest to the midpoint the sign on this relationship changes and the steeper book corresponds to higher volume as we have. It might be that this inner part of the book is the most important for measuring liquidity.

Our direct relationship between trading volume and volatility is more complicated. There is overall a negative correlation, but there is a large jump up to near zero contemporaneously. So it looks that when volatility jumps up, there is more volume as orders get swept off the book. When higher volatility levels settle into the new order shading strategies, the book starts to thin out, and volume starts to drop off. The overall negative correlation between volume and volatility is inconsistent with a lot of time series evidence from markets. We feel this is a clear example of new information, and real markets reacting to this as an increased demand for liquidity. This demand would reduce liquidity, increase spreads, and it could increase volume as well. We only have pure supply changes in

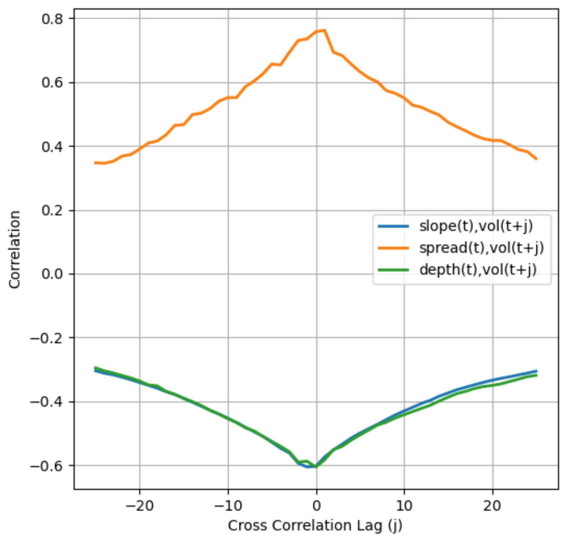

Figure 20: Volatility Cross Correlations (Base case)

liquidity which move positively with trading activity. We would need to introduce new information events to correctly test our model's ability to capture more realistic volume dynamics.

NS report an interesting connection between slopes, volume and volatility which we are able to replicate. They demonstrate that there is a connection between order book slope, and the covariance between volume and volatility. When the order book is steeply sloped the covariance between volume and volatility falls, and when the book is flatter this covariance rises. This is consistent with a flat book indicating a period of heterogeneous beliefs, and new information arriving which could also lead to positive connections between volume and volatility. We see evidence in our simulations for a similar pattern represented by a strong negative correlation between slopes and the covariance shown in figure 21. This is interesting since we do not allow for any changing dynamics of information coming into the market. It suggests this may be a more fundamental property of volume and volatility, and might be related to some of the empirical ambiguity related to trying to estimate the correlation between volatility and trading volume.

Figure 22 repeats figure 19 but now with dynamic order shading turned off. In this case our previous runs showed very simple patterns in price/return time series that were very close to a random walk. We see this again here in returns, but now the lack of time series patterns is repeated in all our order book statistics. In all three cases, spreads, slopes, and depth, there is little evidence for much variability over time as compared with our other runs. The stable return series is represented by a relatively stable limit order book. Market liquidity by all measures appears constant throughout the experiment. Figure 23 displays the cross correlations between our order book features and volatility for this case with static order shading. They show none of the dynamics of the previous case, and

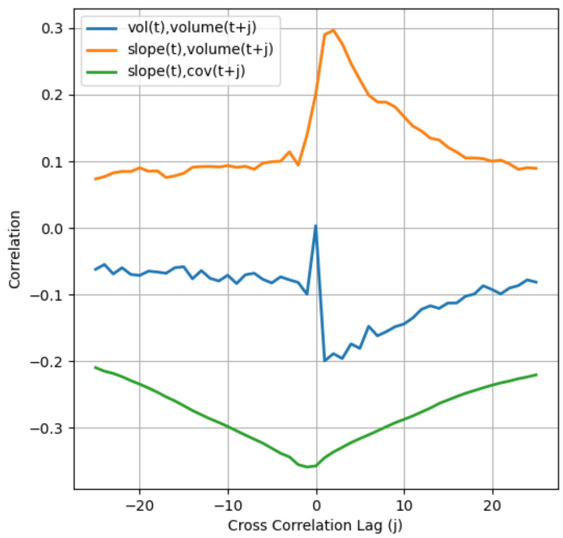

Figure 21: Slope Cross Correlations (Base case)

all are mostly zero. There are some sharp contemporaneous connections which do mimic our usual liquidity patterns, but only as a one day shock. In these cases, high volatility (or maybe just jumps) are connected to large spreads and low order book depth, but the relationship with slope is reversed from our previous figure. Now days with high volatility also appear to have steep slopes. This is somewhat puzzling, as it is counter intuitive. It is also important to note that the magnitudes of these cross correlations are well below those from figure 20.

7 Conclusions

We have developed a simple extension to the models of Chiarella & Iori (2002) and Chiarella et al. (2009). Our key addition is to change the simple order placement strategies of agents to adjust for short range volatility. This modification, accompanied by relatively restrictive order cancellation policies makes sure that market liquidity stays scarce enough to generate persistent volatility patterns. This is consistent with a realistic joint dynamic of volatility and concentration of order book depth around pricing midpoints.

We demonstrate that the model generates all the essential time series stylized facts in asset returns, along with some common features from limit order books. All are consistent with our story of endogenous pulsations in liquidity, and also agents' response to perceived informed trading. We do not claim to be the first agent-based model to accomplish such a task. Many that have proceeded us have done an equivalent, and sometimes better, job of quantitatively hitting the features we are going after. We have two contributions in our model. First, we do this in a very transparent fashion. It is clear where volatility persistence comes from, and the liquidity based story we present

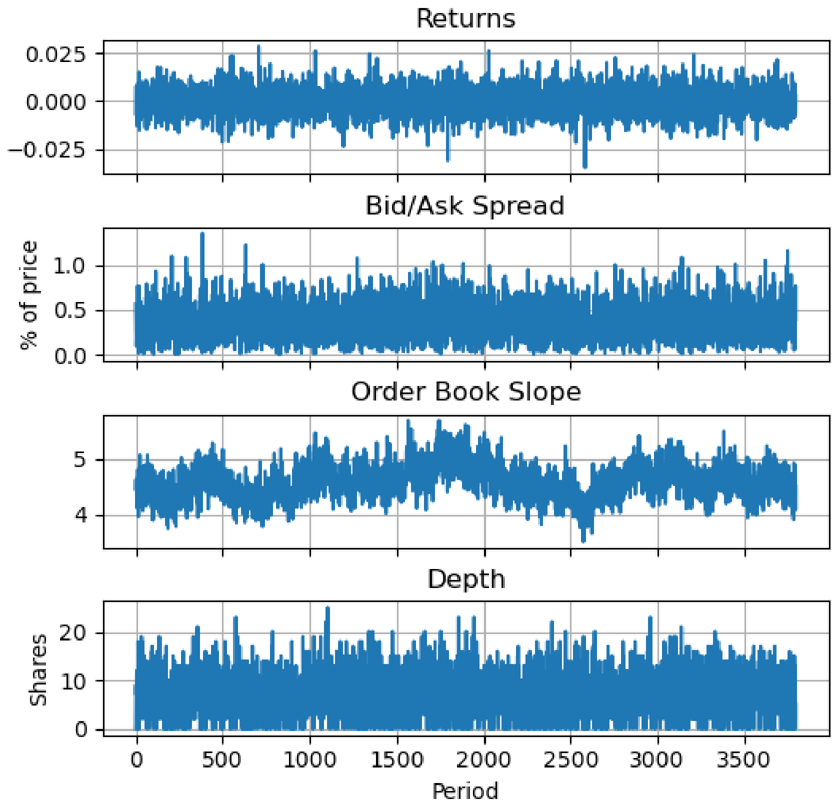

Figure 22: Order Book Dynamics (Static shading)

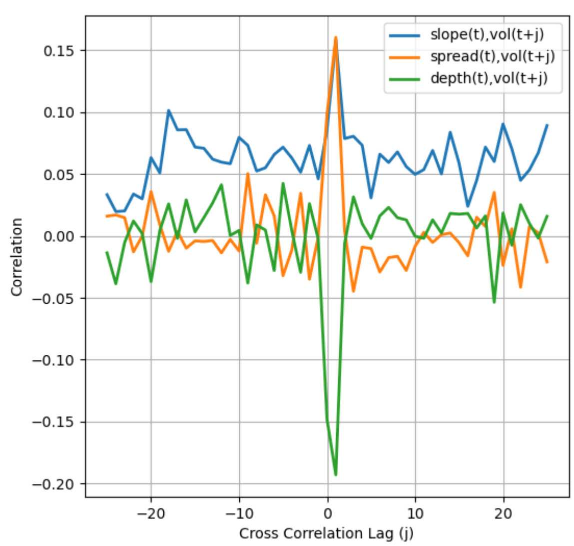

Figure 23: Volatility cross correlations (Static shading)

makes sense in terms of the dynamics of trading. Second, we show that we are able to shut down the interesting dynamics in our model through several routes. The most policy relevant for these would be our manipulation of cancellations so that orders stay on the book for a long time.

Although we hit many empirical features with this simple model there are several critical ones that we do not replicate. One, mentioned previously, is the persistence of order flow. By design our orders are independent, and exhibit no persistence. Also, there is no statistical interaction between limit order strategies and patterns in the order flow. Both of these have been shown to be important in microstructure data.[30] In our model we have decided to avoid trying to model these since they would require us to move to significantly more complex agents, and make our very transparent mechanism less clear.

As a benchmark, we now have a simple limit order book generating reasonable price/return dynamics. This can be very useful when developing more complex empirical and theoretical market models. Beyond limit order books one can view this as a more general mechanism for volatility persistence in other markets. Limit orders represent possible trade points for many agents, or the threshold at which they might make an offer, or take an offer on an asset. The endogenous feedback of current volatility into these threshold prices can become self-fulfilling in terms of pricing and volatility dynamics. We hope to eventually show that this may be a foundational mechanism for an empirical feature that still remains elusive in terms of accepted explanations.

[30] Again, see the surveys of Bouchaud (2018) and Bouchaud et al. (2009) for some of the fundamental theoretical challenges these results present for order book modeling.

References

Ahn, H.-J., Bae, K.-H. & Chan, K. (2001), 'Limit orders, depth, and volatility:Evidence from the stock exchange of Hong Kong', *Journal of Finance* **56**(2), 767–788.

Andersen, T. G. (1994), 'Stochastic autoregressive volatility: A framework for volatility modeling', *Mathematical Finance* **4**, 75–102.

Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2013), Financial risk measurement for financial risk management, *in* G. Constantinedes, M. Harris & R. Stulz, eds, 'Handbook of the Economics of Finance', Vol. 2, Elsevier, pp. 1127–1220.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Ebens, H. (2001), 'The distribution of realized stock return volatilty', *Journal of Financial Economics* **61**(1), 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **96**, 579–625.

Ane, T. & Geman, G. (2000), 'Order flow, transaction clock, and normality of asset returns', *Journal of Finance* **55**(5), 2259–2284.

Baltussen, G., van Bekkum, S. & Da, Z. (2019), 'Indexing and stock market series dependence around the world', *Journal of Financial Economics* **132**, 26–48.

Barndorff-Nielsen, O. E. & Shephard, N. (2002), 'Econometric analysis of realized volatility and its use in estimating stochastic volatility models', *Journal of the Royal Statistical Society B* **64**(2), 253–280.

Barndorff-Nielsen, O. E. & Shephard, N. (2004), 'Power and bipower variation with stochastic volaitlity and jumps', *Journal of Financial Econometrics* **2**, 1–37.

Becker, G. S. (1962), 'Irrational behavior and economic theory', *Journal of Political Economy* **70**, 1–13.

Benhabib, J. & Dave, C. (2014), 'Learning, large deviations and rare events', *Review of Economic Dynamics* **17**(3), 367–382.

Beran, J. (1994), *Statistics for Long-Memory Processes*, Chapman and Hall, New York, NY.

Biondo, A. E. (2019), 'Order book modeling and financial stability', *Journal of Economic Interactions and Coordination* **14**, 469–489.

Biondo, A. E. (2020), 'Information versus imitation in a real-time agent-based model of financial markets', *Journal of Economic Interaction and Coordination* **15**, 613–631.

Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **21**, 307–328.

Bollerslev, T. & Domowitz, I. (1993), 'Trading patterns and prices in the interbank foreign exchange market', *Journal of Finance* **48**, 1421–1444.

Bouchaud, J.-P. (2018), Agent based models for market impact and volatility, *in* C. Hommes & B. LeBaron, eds, 'Handbook of Computational Economics: Heterogeneous Agent Modeling', Vol. 4, North Holland.

Bouchaud, J.-P., Farmer, J. D. & Lillo, F. (2009), How markets slowly digest changes in supply and demand, *in* 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland, pp. 57–160.

Brock, W. A. & Kleidon, A. W. (1990), 'Exogenous demand shocks and trading volume: A model of intraday bids and asks', *Journal of Economic Dynamics and Control* **16**, 451–489.

Chakraborti, A., Toke, I. M., Patriarca, M. & Abergel, F. (2011), 'Econophysics review: I. Empirical facts', *Quantitative Finance* **11**(7), 991–1012.

Chiarella, C., He, X. & Wei, L. (2015), 'Learning, information processing and order submission in limit order markets', *Journal of Economic Dynamics and Control* **61**, 245–268.

Chiarella, C. & He, X.-Z. (2001), 'Asset pricing and wealth dynamics under heterogeneous expectations', *Quantitative Finance* **1**, 509–526.

Chiarella, C. & Iori, G. (2002), 'A simulation analysis of the microstructure of double auction markets', *Quantitative Finance* **2**, 346–353.

Chiarella, C., Iori, G. & Perello, J. (2009), 'The impact of heterogeneous trading rules on the limit order book and order flows', *Journal of Economic Dynamics and Control* **33**, 525–537.

Clark, P. K. (1973), 'A subordinated stochastic process model with finite variance for speculative prices', *Econometrica* **41**, 135–55.

Cliff, D. & Bruten, J. (1997), Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets, Technical Report HPL-97-141, Hewelett-Packard Laboratories, Bristol, UK.

Cohen, K. J., Maier, S. F., Schwartz, R. A. & Whitcomb, D. K. (1983), 'A simulation model of stock exchange trading', *Simulation* **41**, 181–191.

Cont, R. (2001), 'Empirical properties of asset returns: stylized facts and statistical issues', *Quantitative Finance* **1**, 223–236.

Cont, R. (2011), 'Statistical modeling of high-frequency financial data: Facts, models and challenges', *IEEE Signal Processing Magazine* **28**(5), 16–25.

Dieci, R. & He, Z.-Z. (2018), Heterogeneous agent models in finance, *in* C. Hommes & B. LeBaron, eds, 'Handbook of Computational Economics: Heterogeneous Agent Modeling', Vol. 4, North-Holland, chapter 5, pp. 257–328.

Engle, R. F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation', *Econometrica* **50**, 987–1007.

Farmer, J. D., Patelli, P. & Zovko, I. (2005), 'The predictive power of zero intelligence models in financial markets', *Proceedings of the National Academy of Sciences of the United States of America* **102**, 2254–2259.

Flemming, M. J. (1997), 'The round-the-clock market for U.S. treasury securities', *FRBBNY Economic Policy Review* pp. 9–24.

French, K. R., Schwert, G. W. & Stambaugh, R. F. (1987), 'Expected stock returns and volatility', *Journal of Financial Economics* **19**, 3–29.

Gatheral, J. & Oomen, R. C. A. (2010), 'Zero-intelligence realized variance estimation', *Finance and Stochastics* **14**, 249–283.

Gaunersdorfer, A. & Hommes, C. (2007), A nonlinear structural model for volatility clustering, *in* A. Kirman & G. Teyssiere, eds, 'Micro Economic Models for Long Memory in Economics', Springer-Verlag, pp. 265–288.

Gillemot, L., Farmer, J. D. & Lillo, F. (2007), 'There's more to volatility than volume', *Quantitative Finance* **6**(5), 371–384.

Gode, D. K. & Sunder, S. (1993), 'Allocative efficiency of markets with zero intelligence traders', *Journal of Political Economy* **101**, 119–37.

Goettler, R., Parlour, C. A. & Rajan, U. (2005), 'Equilibrium in a dynamic limit order market', *Journal of Finance* **60**, 2149–2192.

Gould, M. D., Porter, M. A., Williams, S., McDonald, M., Fenn, D. J. & Howison, S. D. (2013), 'Limit order books', *Quantitative Finance* **12**(11), 1709–1742.

Handa, P. & Schwartz, R. A. (1996), 'Limit order trading', *Journal of Finance* **51**, 1835–1861.

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

Isaac, A. G. & Ramaswamy, V. (2023 forthcoming), 'Rule-based trading on an order-driven exchange: a reassessment', *Quantitative Finance* .

Karpov, J. M. (1987), 'The relation between price changes and trading volume: A survey', *Journal of Financial and Quantitative Analysis* **22**, 109–126.

Ladley, D. & Schenk-Hoppe, K. R. (2009), 'Do stylized facts of order book markets need strategic behavior', *Journal of Economic Dynamics and Control* **33**, 817–831.

Leal, S. J., Napoletano, M., Roventini, A. & Fagiolo, G. (2016), ‘Rock around the clock: An agent-based model of low and high frequency trading’, *Journal of Evolutionary Economics* **26**(1), 49–76.

LeBaron, B. (2001), ‘Empirical regularities from interacting long and short memory investors in an agent based stock market’, *IEEE Transactions on Evolutionary Computation* **5**, 442–455.

LeBaron, B. & Yamamoto, R. (2008), ‘The impact of imitation on long-memory in an order-driven market’, *Eastern Economic Journal* **34**, 504–517.

Lillo, F. (2007), ‘Limit order placement as an utility maximization problem and the origin of power law distribution of limit order prices’, *The European Physical Journal B* **55**(453-459).

Liu, Y., Cizeau, P., Meyer, M., Peng, C. K. & Stanley, H. E. (1997), ‘Correlations in economic time series’, *Physica A* **245**, 437–440.

Lux, T. (1997), ‘Time variation of second moments from a noise trader/infection model’, *Journal of Economic Dynamics and Control* **22**, 1–38.

Lux, T. & Zwinkels, R. C. J. (2018), Empirical validation of agent-based models, *in* C. Hommes & B. LeBaron, eds, ‘Handbook of Computational Economics: Heterogeneous Agent Modeling’, Vol. 4, North Holland, chapter 8, pp. 437–488.

Mandelbrot, B. B. (1963), ‘The variation of certain speculative prices’, *Journal of Business* **36**, 394–419.

Menkhoff, L., Osler, C. L. & Schmeling, M. (2010), ‘Limit-order submission strategies under asymmetric information’, *Journal of Banking and Finance* **34**, 2665–2677.

Mike, S. & Farmer, J. D. (2008), ‘An empirical behavioral model of liquidity and volatility’, *Journal of Economic Dynamics and Control* **32**, 200–234.

Naes, R. & Skjeltorp, J. A. (2006), ‘Order book characteristics and the volume-volatility relation: Empirical evidence from a limit order market’, *Journal of Financial Markets* **9**(4), 408–432.

Paddrik, M. E., Jr., R. H., Todd, A., Yang, S. Y., Scherer, W. & Beling, P. (2012), An agent-based model of the E-mini S&P 500 and the Flash Crash, *in* ‘2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr)’, pp. 1–8.

Paulin, J., Calinescu, A. & Wooldrich, M. (2019), ‘Understanding flash crash contagion and systemic risk: A micro-macro agent-based approach’, *Journal of Economic Dynamics and Control* **100**, 200–229.

Staccioli, J. & Napoletano, M. (2021), ‘An agent-based model of intra-day financial markets dynamics’, *Journal of Economic Behavior and Organization* **182**, 331–348.

Tauchen, G. E. & Pitts, M. (1983), 'The price variability-volume relationship on speculative markets', *Econometrica* **51**(2), 485–505.

Taylor, S. J. (2005), *Asset price dynamics, volatility, and prediction*, Princeton University Press, Princeton, NJ.

Toth, B., Palit, I., Lillo, F. & Farmer, J. D. (2015), 'Why is order flow so persistent?', *Journal of Economic Dynamics and Control* **51**, 218–239.

Vuorenmaa, T. A. & Wang, L. (2014), An agent-based model of the flash crash of May 6, 2010, with policy implications, Technical report, SSRN.

Yamamoto, R. (2011), 'Order aggressiveness, pre-trade transparency, and long memory in an order-driven market', *Journal of Economic Dynamics and Control* **35**, 1938–1963.

Yang, S. Y., Paddrick, M. E., Hayes, R., Todd, A., Kirilenko, A., Beling, P. & Scherer, W. (2012), Behavior based learning in identifying high frequency trading strategies, *in* '2012 IEEE Conference on Computational Intelligence for Financial Engineering and Economics (CIFEr)', IEEE, pp. 1–8.

Zhan, W. & Friedman, D. (2007), 'Markups in double auction markets', *Journal of Economic Dynamics and Control* **31**, 2984–3005.

Zovko, I. & Farmer, D. (2002), 'The power of patience: a behavioral regularity and limit-order placement', *Quantitative Finance* **2**(5), 387–392.

Appendix A: Parameter Sweeps

In this Appendix, we demonstrate parameter sweeps over the values:

$$\sigma_f \in \{0.05,\ 0.1,\ 0.2,\ 0.4,\ 0.6\},$$
$$\text{PMO} \in \{0,\ 0.025,\ 0.05,\ 0.1\},$$
$$\text{kMax} \in \{0.05,\ 0.1{,}0.015,\ 0.2\},$$
$$\sigma_e \in \{0.00005,\ 0.0001,\ 0.0005,\ 0.001\},$$
$$\tau \in \{10,\ 25,\ 50,\ 250\}$$
$$\Delta \in [10, 25, 40, 50, 60, 75, 100, 125]$$

σ_f indicates the fraction of traders who are fundamental traders. PMO is the probability that a noise trader will submit a market order instead of placing an order on the book. kMax represents the maximum amount of order shading. σ_ϵ is the noise level for noise traders. τ is the time before order cancellation on the book (in small time units), and Δ represents the number of small time units per simulated day. These sweeps demonstrate relative robustness across a wide range of parameters.

These plots are generated by running 5 long run simulations for each combination of parameters and reporting the average statistics for the daily return series. The results are plotted across slight variations of specific parameters.

Increasing σ_f in figure 24 results in quicker reversion to the fundamental price. For large values this can lead to autocorrelation. However, $\sigma_f \in \{0.05, 0.1, 0.2\}$ produces reasonable results. For all values of σ_f we show excess kurtosis, and positive absolute return autocorrelations. The increased level of mean reversion driven by fundamental traders is evident in the decreasing return autocorrelation (negative), and the reductions in return standard deviation.

Increasing the probability of market orders (PMO) in figure 25 increases general trading activity in the market. The sweep shows that with this we see increases in return standard deviations, and decreases in return autocorrelations, and increases in absolute return autocorrelations. Kurtosis starts below 3, but quickly increases to the typical range. In general we are showing all the basic features for a wide range of PMO. We decided to stay with $PMO = 0.05$ which gives us an executed order probability of about 0.05 which is close to what is commonly seen in many markets.

Increasing kMax in figure 26 makes the returns series more 'choppy' due to the wider spread that result from agents placing orders further from their valuation. For large values of kMax kurtosis and standard deviation become extreme, as does the absolute return autocorrelation. It therefore seems reasonable to keep it at the intermediate value of 0.10.

As we increase the noise level, σ_ϵ, in 27 most of the interesting time series results disappear as the market is dominated by noise. Kurtosis moves toward 3, and absolute return autocorrelations approach zero. The annualized std. for returns keeps steadily increasing as more noise is added. Interestingly the return autocorrelations become more negative as more noise is added to the market.

Increasing the value of τ in figure 28 we see how letting orders stay on the book longer can shut down much of the time series dynamics in the model. By the time we reach the value used in the earlier experiments, 75, it is clear that kurtosis and the autocorrelation of absolute returns have both dropped significantly. It is again interesting that further increases in τ will eventually induce some negative autocorrelations in the daily return series. 75 seems like a reasonable threshold for this phase shift in the model dynamics.

Finally, changing the number of small periods per day, Δ is explored in figure 29. As the aggregation level is increased we see a small reduction in kurtosis, and a steady increase in return std. and autocorrelations with the latter approaching zero. Also, as the time period increases, the amount of autocorrelation in the absolute returns falls off to near zero. Our mid range value of 50 seems reasonable.[31]

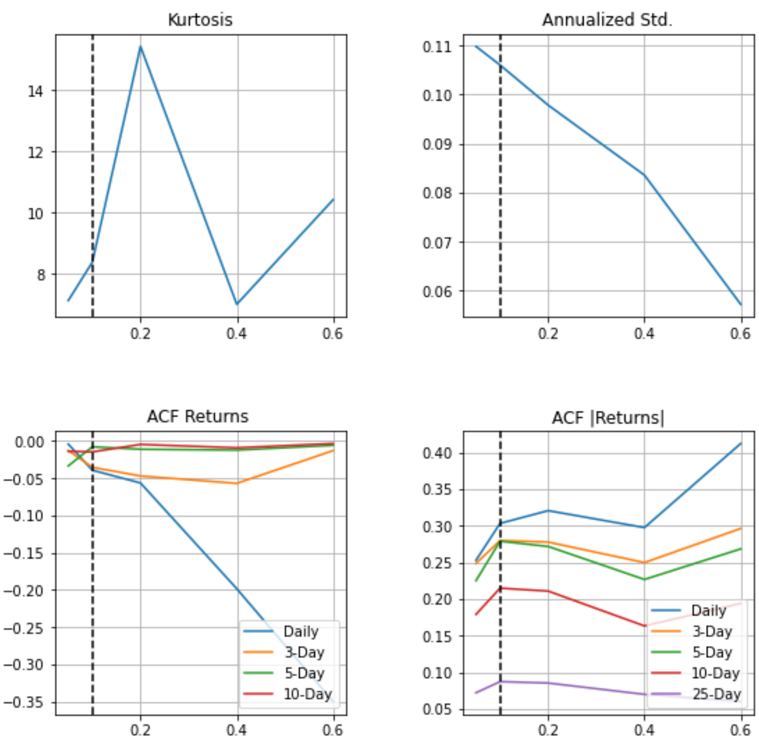

Figure 24: $\tau = 25,\ \sigma_e = 0.0001,\ kMax = 0.1,\ \sigma_F \in [0.05, 0.1, 0.2, 0.4, 0.6]\ PMO = 0.05$

[31]The annualized std. of 0.1 is low relative to most stock return series. Depending on the period this value is between 0.15 and 0.2 for the SPY. We still believe this is fine, since our model assumes no new information arrivals, and a fixed fundamental value.

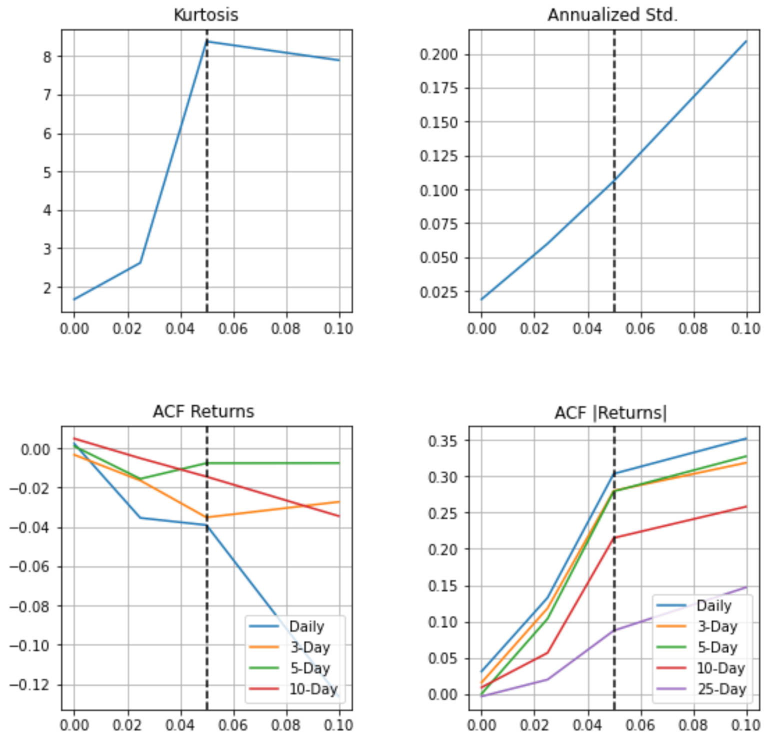

Figure 25: $\tau = 25,\ \sigma_e = 0.0001,\ kMax = 0.1,\ \sigma_F = 0.1,\ PMO \in [0, 0.025, 0.05, 0.1]$

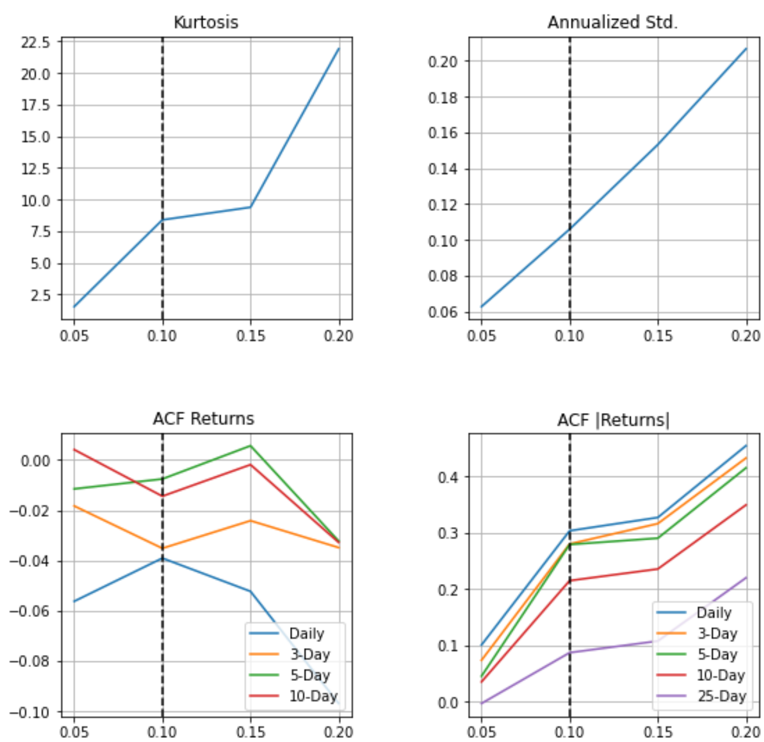

Figure 26: $\tau = 25,\ \sigma_e = 0.0001,\ kMax \in [0.05, 0.1, 0.15, 0.2],\ \sigma_F = 0.1,\ PMO = 0.05$

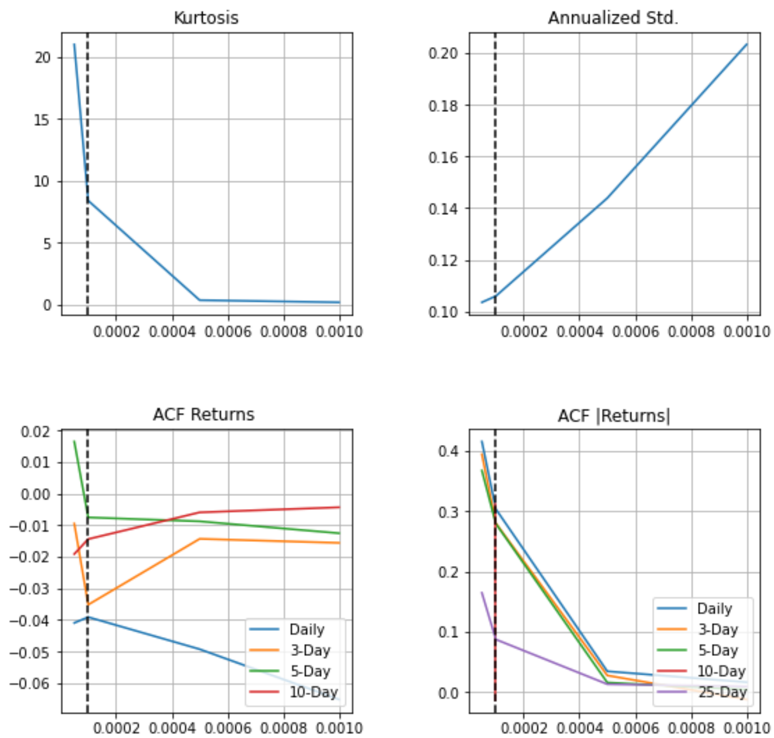

Figure 27: $\tau = 25,\ \sigma_e \in [0.00005, 0.0001, 0.0005, 0.001],\ kMax = 0.1,\ \sigma_F = 0.1,\ PMO = 0.05$

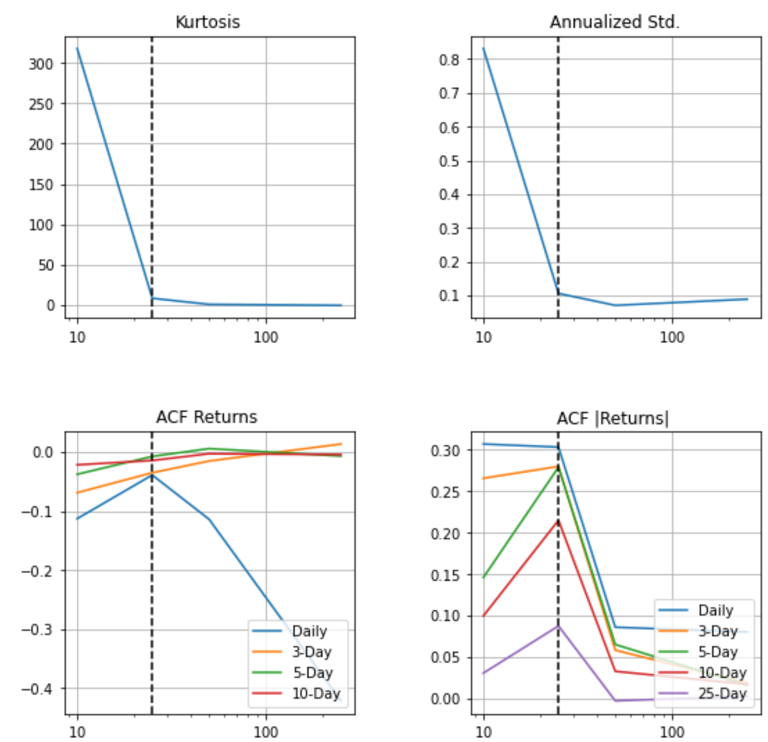

Figure 28: $\tau \in [10, 25, 50, 250],\ \sigma_e = 0.0001,\ kMax = 0.1,\ \sigma_F = 0.1,\ PMO = 0.05$

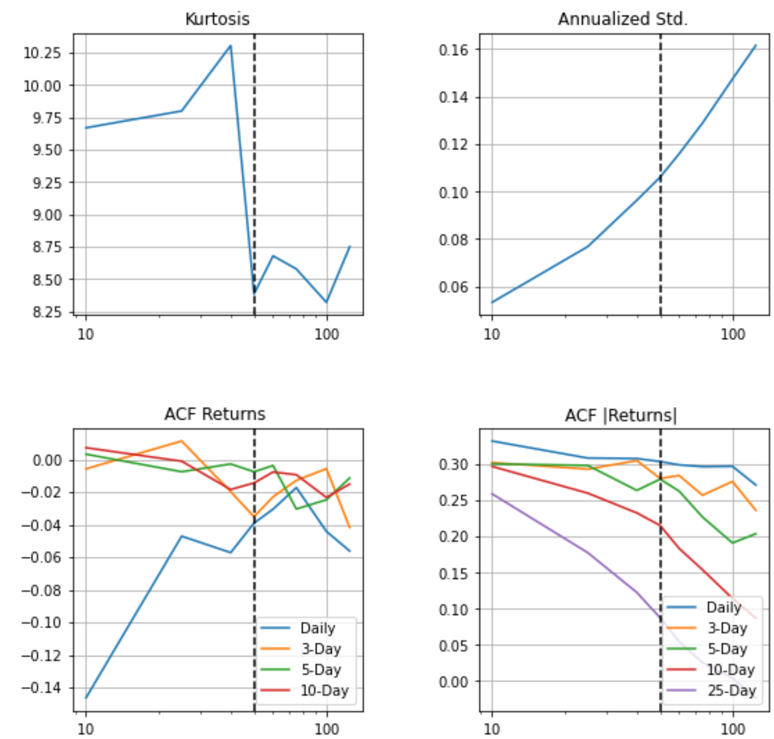

Figure 29: $\tau\ \sigma_e = 0.0001,\ kMax = 0.1,\ \sigma_F = 0.2,\ PMO = 0.05, \Delta \in [10, 25, 40, 50, 60, 75, 100, 125]$