

PARTICLE FILTERING AND INFERENCE FOR LIMIT ORDER BOOKS IN HIGH FREQUENCY FINANCE

Pinzhang Wang Lin Li Simon J. Godsill

University of Cambridge, Department of Engineering
Trumpington Street, Cambridge CB2 1PZ, United Kingdom

ABSTRACT

This paper investigates the on-line analysis of high-frequency financial order book data using Bayesian modelling techniques. Order book data involves evolving queues of orders at different prices, and here we propose that the order book shape is proportional to a gamma or inverse-gamma density function. Inference for these models is implemented on-line using particle filters and evaluated on a high-frequency EURUSD foreign exchange limit order book. The two possible order book shapes are tested using particle filter marginal likelihood estimates and in addition, heat maps are constructed based on the inference results to reveal the imbalance of order distributions between the two sides of an order book, thereby offering valuable insights into the movements of future prices.

Index Terms— order book imbalance, particle filter, limit order book

1. INTRODUCTION

A limit order book in financial markets comprises multiple queues of orders to buy or sell a security (for example shares or foreign exchange rates) at different (discrete) price levels. Figure 1 illustrates a single time snapshot of such a system. Orders to sell join a queue at higher prices (the ‘Ask’ side) and orders to buy sit below these (the ‘Bid’ side). Halfway between the bid and ask sides lies the ‘mid-price’ m_t , which is the price usually quoted for reporting purposes, and it evolves over time as adjacent bid or ask order queues become depleted to zero size through execution (buy/sell at the order’s price) or through time-expiry or cancellation of orders. Hence the balance and shape of the order book can be expected to influence future behaviours of market participants and future price fluctuations [1]. Although individual orders arrive at random timings, in this paper we model the evolution of discrete time ‘snapshots’ of the order queue sizes at regular time intervals (say, 0.5s). In each order book time snapshot, limit orders are placed at multiples of the smallest price change d allowed in the market. The volume of orders Q_j declines away from the mid price due to rapid depletion of orders close to the mid-price and low prospect of order-matching at price levels far from m_t . Therefore, in the histogram of volumes Q_j (the histogram bars in Figure 1 and Figure 6), the shape of an order book snapshot can be observed to resemble a Gamma or inverse gamma density function [2], and we have overlaid our model-based sequential fits to these data on the plots.

To model an order book and its imbalance across the spread has been the focus of several research efforts. Cont et al. in [3] and [4] construct stochastic models for order books by modelling the arrival and cancellation of orders as price-level-specific Poisson processes, and subsequently using queueing theory to calculate the probabilities

for the next price movement. Jiang in [2] proposed that the distribution of orders can be represented by Gamma density curves, and a Kalman filter is then used to track the shape and scale parameters α and β .

This paper presents an alternative tracking technique, particle filtering, which can be applied very generally to systems that are sufficiently well defined [5]. The advantage of particle filtering is that it can be applied to non-linear models and non-Gaussian noise [6] (see e.g. [7] for a high-frequency finance example), as is the case for our proposed models here. Hence we can model the queue likelihoods directly and use non-Gaussian transition densities for their parameters over time, in contrast with the relatively crude Gaussian and indirect modelling approaches of [2], in which the filtering is applied to Maximum Likelihood sequences of fitted α/β values. We are also able to evaluate alternative hypotheses for the order book shape model, in this case an inverse-gamma density. We contribute to current research by providing novel state-space modelling approaches to high-frequency order book snapshot data, and algorithms for their estimation. Section 2 presents the proposed state-space model for the limit order book system. The particle filtering algorithms and the derivation of the model probabilities are given in section 3. Section 4 evaluates the model performance and discusses how this novel tracking of order book shape might be used for price inference and prediction.

2. MODELLING

First we introduce the two distributions that form the basis of our order book shape models. The Gamma (Ga) density is:

$$\text{Ga}(x; \alpha, \beta) = \frac{x^{\alpha-1} \exp\left(-\frac{x}{\beta}\right)}{\beta^\alpha \Gamma(\alpha)} \quad \text{for } x > 0 \text{ and } \alpha, \beta > 0, \quad (1)$$

where $\Gamma(\alpha)$ is the gamma function and α and β are shape and scale parameters, respectively. The distribution of $\frac{1}{\mathbf{X}}$, where \mathbf{X} is gamma-distributed, is the inverse gamma (IG) distribution:

$$\text{IG}(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} \exp\left(-\frac{\beta}{x}\right) \quad \text{for } x > 0. \quad (2)$$

Prices in the market are on a discrete scale nd , where $n > 0$ is an integer and the minimum price quantisation is defined as d . We separate the bid and ask side of the data and model them independently of each other. If we consider J possible discrete price levels each side of the mid price then we propose to model the queue size Q_i at price level $i \in \{1, \dots, J\}$ as a multinomial distribution, i.e. we treat the price of each order as a random draw from a discrete distribution $\{p_i\}_{i=1,2,\dots,J}$. The data are then a vector of queue sizes $y_t = [Q_1^t, Q_2^t, \dots, Q_J^t]^T$, one for each of the bid and ask side.

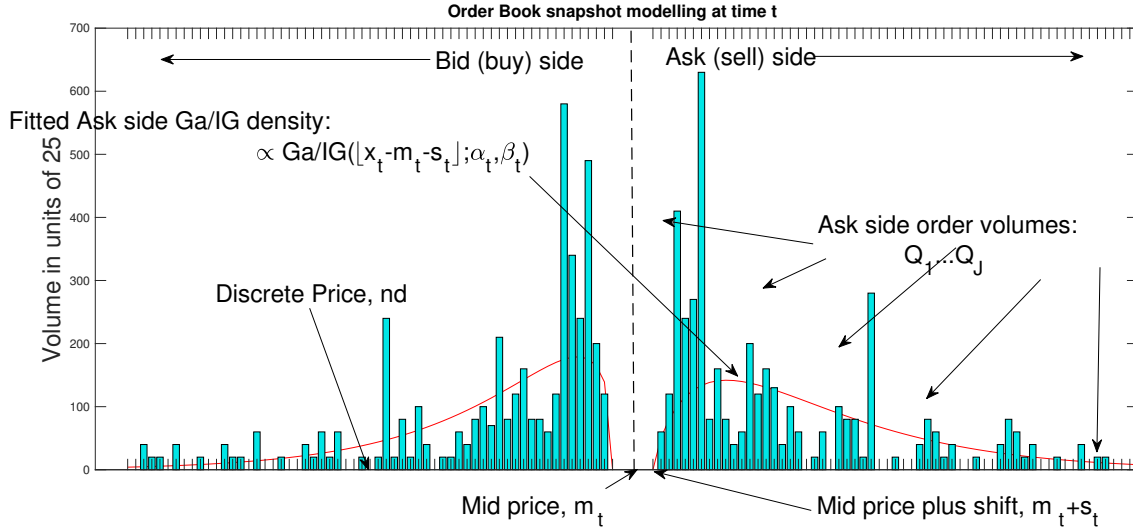


Fig. 1: Illustration of the data and the model, single time t

In order to specify the probabilities we propose to use the Gamma or inverse gamma density function, shifted away from the current (known) mid price m_t by a random (unknown) term s_t . The price corresponding to queue size Q_i is then $d[(m_t + s_t + (i-1)d)/d]$ for the ask side, or $d[(m_t - s_t - (i-1)d)/d]$ for the bid side. This shift parameter s_t has been found in our experience to give a much improved fit to the real data than the non-shifted versions proposed in [2]. Then, the probabilities of the order lying at price level i is determined at each time t as

$$p_i^t = \frac{1}{A} \int_{(i-1)d}^{id} \text{Ga/IG}(x; \alpha_t, \beta_t) dx, \quad A = \sum_{i=1}^J p_i^t \quad (3)$$

and the multinomial ('Multi-Ga/Multi-IG') likelihood is:

$$g(y_t | \mathbf{x}_t) = \frac{(\sum_{i=1}^J Q_i^t)!}{Q_1^t! \dots Q_J^t!} \prod_{i=1}^J (p_i^t)^{Q_i^t} \quad (4)$$

The state of the dynamic order book system on one side of the book can then be defined as:

$$\mathbf{x}_t = [\alpha_t, \beta_t, s_t]^T. \quad (5)$$

with one such state vector independently estimated with its own dynamics and parameters for the data on each of the bid and ask sides. In order to complete the state-space model, we specify the transition density $f(\mathbf{x}_t | \mathbf{x}_{t-1})$. To maintain the positivity of α_t and β_t , their transition densities are chosen as the Gamma distribution (a similar formulation follows for β_t , which is treated as independent of α_t a priori):

$$f(\alpha_t | \alpha_{t-1}) = \text{Ga} \left(\alpha_t; \frac{\alpha_{t-1}}{\sigma_\alpha^2 / \alpha_{t-1}}, \frac{\sigma_\alpha^2}{\alpha_{t-1}} \right) \quad (6)$$

where the degrees of freedom are chosen such that the mean is preserved from the previous α , $E[\alpha_t] = \alpha_{t-1}$, and σ_α is a user-defined standard deviation. The *shift* term s_t will be sampled from a normal distribution with mean equal to s_{t-1} , again mutually independent of α_t and β_t , $f(s_t | s_{t-1}) = \mathcal{N}(s_t; s_{t-1}, \sigma_s^2)$.

We will investigate two models in this paper: Model M_1 has the multinomial gamma (Multi-Ga) likelihood, while M_2 has inverse gamma (Multi-IG) likelihood.

3. INFERENCE METHODOLOGY

We perform statistical filtering to estimate sequentially the posterior density of the hidden state \mathbf{x}_t , $p(\mathbf{x}_t | y_{1:t})$. This is implemented using a particle filter with proposal equal to the transition density for x_t ('bootstrap' filter) with stratified resampling and importance weights proportional to the multinomial likelihood $g(y_t | x_t)$ [8].

Summarising briefly the steps involved:

- At time $t-1$ we have N particles $\mathbf{x}_{t-1}^{(i)} \sim p(\mathbf{x}_{t-1} | y_{1:t-1})$.
- Then for each particle, propagate forward from the prior dynamics $\mathbf{x}_t^{(i)} \sim f(\mathbf{x}_t | \mathbf{x}_{t-1}^{(i)})$ and compute an incremental weight $w_t^{(i)} \propto g(y_t | \mathbf{x}_t^{(i)})$. The incremental weight is stored for calculation of approximate marginal likelihoods, see below.
- The weighted particles then form an approximate updated representation of the new filtering density $p(\mathbf{x}_t | y_{1:t})$.
- Finally, renormalise the weights to sum to 1 and if necessary resample using stratified resampling, resetting the weights to uniform in this case and return to the first step with $t-1 \leftarrow t$.

Model fit is assessed using marginal likelihood $l(M_i) = p(y_{1:t} | M_i)$, $i = 1, 2$, which is computed approximately from the sequence of particle weights as [9]:

$$l(M_i) = l_{1:T}(y_{1:T} | M_i) = l_1(y_1 | M_i) \prod_{t=2}^T l_{t|1:t-1}(y_t | y_{1:t-1}, M_i) \\ \approx l_1(y_1 | M_i) \prod_{t=2}^T \frac{1}{N} \sum_{i=1}^N g(y_t | \mathbf{x}_t^{(i)}).$$

Note that $g(y_t | \mathbf{x}_t^{(i)})$ is the unnormalised weight for each particle i at time t .

	Time point 700-800, Hour 3		Time points 700-800, Hour 9	
	Ask-side	Bid-side	Ask-Side	Bid-Side
M_1	-1.7271 (-1.4192)	-1.5604 (-1.2558)	-1.7550 (-1.4247)	-1.7144 (-1.3721)
M_2	-1.5368 (-1.2637)	-1.4176 (-1.1354)	-1.4735 (-1.2085)	-1.4107 (-1.1430)

Table 1: Model Probabilities $l(M_i)$ ($l_{21:101}(M_i)$) [$\log \times 10^5$]

4. RESULTS

The data set¹, is a EURUSD limit order book that spans 24 hours starting with Hour 0 at ET 5pm, 2 Sep, 2015 (i.e. ‘Hour 3’ = ET 8pm). ‘Snapshots’ of unexecuted order prices and quantities are taken to deliver histogram data y_t .

Results in Table 1 are running 1000 particles on $t = 700 - 800$, Hour 3. The initial mean of \mathbf{x}_0 is $[1.5, 5 \times 10^{-5}, 0]^T$ and the standard deviations σ_α , σ_β and σ_s are 0.1, 3×10^{-5} and 1×10^{-5} , respectively. Both α_0 and β_0 follow a Gamma distribution and s_0 follows a normal distribution. To minimise any bias introduced by random initialisation, the model likelihood for the last 81 time points, $l_{21:101}$, is also shown in brackets in Table 1. Results show that the model M_2 with the IG distribution generates a higher $l(M_i)$ overall, which illustrates that the IG distribution, with heavier tails, accounts better for limit order book here. Similar experiments are also run on other periods to confirm this conclusion. Marginal likelihoods for $t = 700 - 800$ in Hour 9 are also recorded in Table 1. Moreover, we plot the mean of filtered \mathbf{x}_t with their 95% error bars in Figure 4, illustrating the tracking of \mathbf{x}_t over $t = 700 - 800$. Tracking results are also summarised by plotting image heat-maps of the order volumes at different price levels for each time point (red indicates highest volume, blue zero) in order to show the market liquidity, in Figure 2a. The probability densities of the Multi-Ga/Multi-IG can also be scaled by total volume of orders on each side at any t to compare with the original order book as shown in Figure 3b and Figure 3a. We map the shape of the scaled Multi-Ga/Multi-IG to color following the same rule to construct similar heat maps in Figure 2b and 2c for model M_1 and model M_2 . We use the mean of the particle filtered \mathbf{x}_t at time point t to construct the heat map. The empirical difference in modelling between the Ga distribution and the IG distribution can be seen clearly in the heat maps. The order book imbalance can be revealed by the heat map. For instance, on $t = 740 - 750$, there is a high liquidity region on the bid side but a lower liquidity region on the ask side.

We can overlay the estimated Multi-Ga/Multi-IG distribution for the bid and ask sides on an order book snapshot and compare them with original histograms of orders. Both Figure 3a and Figure 3b show a high concentration of orders on the bid side, which reflects an imbalance. But the IG distribution has a heavier tail, accounting better for the orders residing distant from the mid price. From Figure 3a, the height of Multi-IG on the bid side is higher than that of the ask side, which is not immediately clear from the unprocessed data.

Imbalance changes may have non-trivial signaling power for future price changes, which in this case is an intuitive price increase that follows afterwards. For instance, at $t = 735 - 746$ in Figure 2c, an imbalance can be observed while the mid price goes down rather than up. This hints at a more complex correlation between the order book and the mid price, and further statistical analysis is required to strengthen such correlation analysis. The detailed shape and imbalance analysis provided by our models allows potentially for many possible variables to be used in prediction.

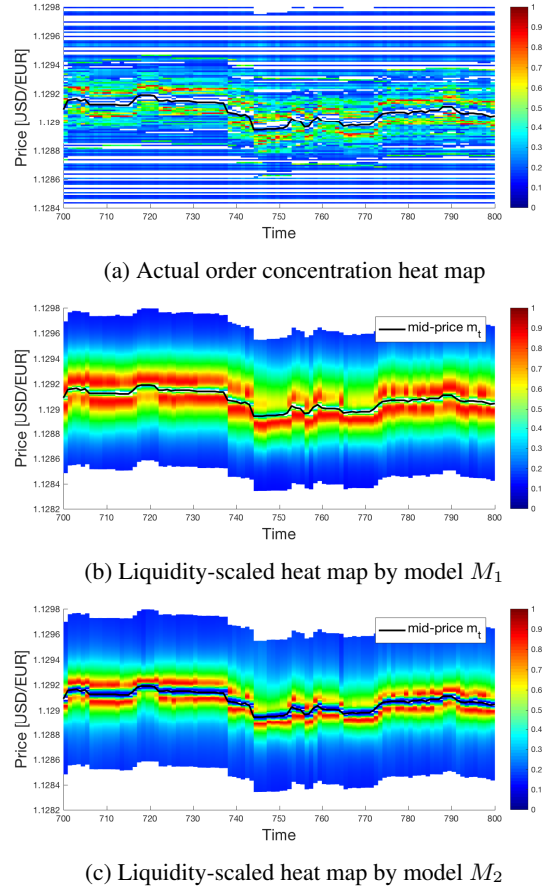




Fig. 2: Probability heat maps on $t = 700 - 800$, Hour 3

¹Thanks to Cambridge Capital Partners for providing data



(a) The order book snapshot at $t = 746$ overlaid by scaled Multi-IG distribution



(b) The order book snapshot at $t = 746$ overlaid by scaled Multi-Ga distribution

Fig. 3: The fitting of an order book snapshot at $t = 746$




Fig. 4: Plots of filtered \mathbf{x}_t for the ask side of M_2 on $t = 700 - 800$, Hour 3

There are also several limitations in using limit order book imbalance for price inference: limit orders only reflect the market demand from ‘patient’ traders [10], whose orders should eventually be executed against, or converted to, market orders, leading to order depletion and price changes. However, it is questionable that limit orders are representative of the total market demand, which also include a large amount of market orders. Moreover, certain limit orders are placed to only ‘probe’ the market, or to discover the existence of ‘iceberg orders’, which are big orders (that tend to drive price changes) only partially reviewed in an order book[11]. Some of such noises may be filtered out through fine-tuning the variance parameters, but difficult to account for thoroughly.

5. CONCLUSIONS AND DISCUSSIONS

Jiang and Ng [2] firstly used the Gamma distribution to track the shape of the limit order book by attempting to capture the dynamics of the shape parameters and the scale parameters on the bid/ask side. But the filtered results depend heavily on the value of the process noise variance Q and the measurement noise variance R . As shown in Figure 5, the filtered parameters with a large value of dynamic noise fluctuate too randomly and can lead to very poor overall density fits to the histogram data. Smoother curves are obtained with a small value of dynamic noise variance, but nevertheless if we look at the plot of Gamma distribution curve fitting with smoothed param-




Fig. 5: Smoothed α and β for the gamma distribution on $t = 700 - 800$, Hour 3 (—: α and β of gamma distribution by MLE, - - -: α and β smoothed by small Q , . . . : α and β smoothed by large Q)




Fig. 6: Gamma distribution with smoothed parameters overlaid on a snapshot of limit order book at $t = 746$, Hour 3 (—: smoothed by Kalman filter, \circ : Multi-Ga of M_1 (filtered), \triangle : Multi-IG of M_2 (filtered))

ters at a time point in Figure 6, the curve doesn’t fit limit order book as well as the models we proposed in our paper.

In contrast with the significant price changes that are usually caused by major macro-economic events, the smaller-scale and more volatile intraday price movements can be better understood through the changes in order book shapes [3]. This paper has shown that the newly proposed model with IG density offers a better fit of the underlying shape of limit order book. Meanwhile, the probability heat maps generated by the particle-filtered means of α , β and s are able to reveal visually the underlying imbalance in shape between the bid and ask sides of an order book. Our future research will motivate future endeavour in improving order book models and subsequently quantifying the correlation between the filtered states and future price movements.

6. REFERENCES

- [1] Christine Parlour and Duane J. Seppi, “Limit order markets. a survey,” 12 2008.
- [2] J. Jiang and W. L. Ng, “Capturing order book dynamics with Kalman filters,” in *The 6th International Conference on Networked Computing and Advanced Information Management*, Aug 2010, pp. 428–433.
- [3] Rama Cont, Sasha Stoikov, and Rishi Talreja, “A stochastic model for order book dynamics,” *Operations research*, vol. 58, no. 3, pp. 549–563, 2010.
- [4] Rama Cont and Adrien De Larrard, “Price dynamics in a Markovian limit order market,” *SIAM Journal on Financial Mathematics*, vol. 4, no. 1, pp. 1–25, 2013.
- [5] Randal Douc and Olivier Cappé, “Comparison of resampling schemes for particle filtering,” in *Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on*. IEEE, 2005, pp. 64–69.
- [6] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, “Novel approach to nonlinear/non-gaussian bayesian state estimation,” *IEE Proceedings F - Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, April 1993.
- [7] Hugh L Christensen, James Murphy, and Simon J Godsill, “Forecasting high-frequency futures returns using online Langevin dynamics,” *Selected Topics in Signal Processing, IEEE Journal of*, vol. 6, no. 4, pp. 366–380, 2012.
- [8] Olivier Cappé, Simon J Godsill, and Eric Moulines, “An overview of existing methods and recent advances in sequential Monte Carlo,” *Proceedings of the IEEE*, vol. 95, no. 5, pp. 899–924, 2007.
- [9] Arnaud Doucet and Adam M. Johansen, “A tutorial on particle filtering and smoothing: fifteen years later,” 2011.
- [10] Frank De Jong and Barbara Rindi, *The microstructure of financial markets*, Cambridge University Press, 2009.
- [11] Hugh L Christensen and Robert Woodmansey, “Prediction of hidden liquidity in the limit order book of globex futures,” vol. 8, pp. 68–95, 06 2013.