A Stochastic Model for Limit Order Books

Applications to High Frequency Trading

Chung, Dahan, Hocquet, Kim

MS&E 444

Professor Giesecke

June 10, 2009

1. Introduction

This paper outlines our use of Rama Cont et al's stochastic model for limit order book dynamics in applications to high-frequency trading. The model itself uses Poisson processes to model order arrivals and cancellations. The resulting birth-death process allows the use of inverse Laplace transform methods involving continued fractions to calculate various conditional probabilities of interest, such as the probability for the mid price increasing given the current state of the order book. These probabilities can then be used as short-term predictions to develop high-frequency trading and execution strategies.

Section 2 of this paper outlines the model developed by Cont, as well as various diagnostic tests affirming the model's statistical accuracy. Section 3 expands on the inverse Laplace transform method, as well as the results of the computation for a selection of probabilities. Section 4 discusses trading strategies involving the computed probabilities and their tractability.

2. Model Assessment and Diagnostics

Cont's model for limit order book dynamics is as follows. There are two types of order, limit orders and market orders. Limit orders persist until they are cancelled, or filled by a market order.

- All orders are of unit size, corresponding to the average size of orders passed
- Limit buy (sell) orders arrive at a distance of i ticks from the opposite best quote at independent, exponential times with rate $\lambda(i)$
- Market buy (sell) orders arrive at independent, exponential times with rate μ
- Order cancellation rate is proportional to the number of outstanding orders
- Arrival rate of orders at a given price is function of the distance to the bid/ask, and follows a power law: $\lambda(i) = k/i^{\alpha}$, for parameters k, α

Given these assumptions, the order book state can be considered a continuous-time Markov Chain.

Before fully utilizing Cont's model and subsequent results in trading strategies, we conducted tests to determine whether an order book populated according to Cont's dynamics possesses traits common to limit order books as found in the literature. MATLAB was used to simulate several days worth of orders and cancellations according to Cont's model, using the parameters set out in the paper. In particular, we present the following characteristics as discussed in detail by Slanina:

- Long-term shape of the order book
- Distribution of durations
- Negative lag-1 autocorrelation
- Microstructure noise
- Hurst coefficient > 0.5

The shape of the order book is found to be reasonable. The average number of limit orders active at a given distance peaks very close to, but not at the best bid/ask, which is realistic. The distribution past the peak for increasingly large distances from the mid price seems follow a power law with exponent ≈ 2, which is also expected, but the distribution between the mid price and the peaks is unclear.

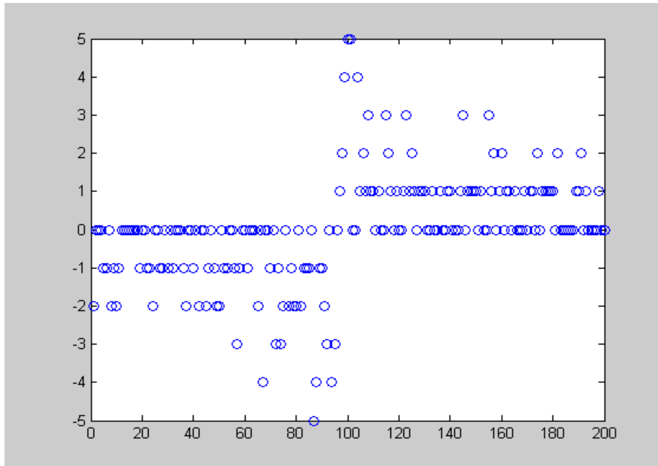

Figure 1. Long term shape of the order book

According to Slanina, empirical studies have shown that incoming limit orders have volumes which are distributed according to a power law, with an exponent of ≈1.4 for market orders and a higher exponent of ≈2 for limit orders. Cont's model is fundamentally different because one of the main assumptions is that each order has the same unit size, the average size observed in the past. The distribution of the distances of limit orders from the mid price also follows a power law, although the value of the exponent reported differs rather widely (≈1.5 to ≈2.5) from one study to another.

Limit orders persist until they are satisfied or cancelled, and the lifetime of limit orders is also distributed according to a power law with exponent ≈2.1 for cancelled orders and ≈1.5 for satisfied orders. Here again, the model is satisfying.

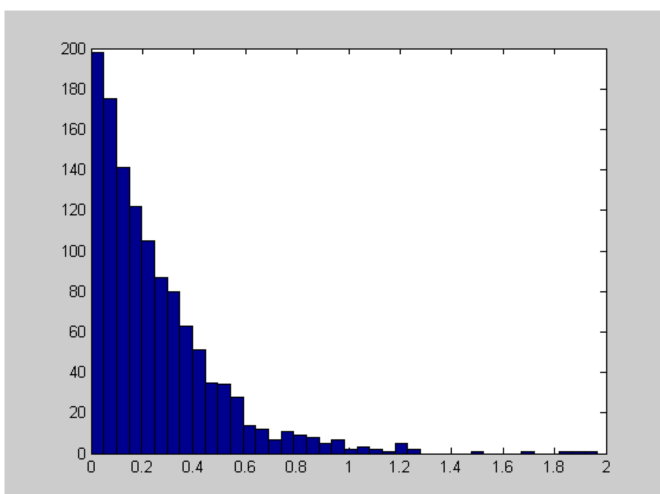

Figure 2. Distribution of durations

The sample autocorrelation function shows a large negative lag-1 autocorrelation, and no other significant values for other lags, as we expect.

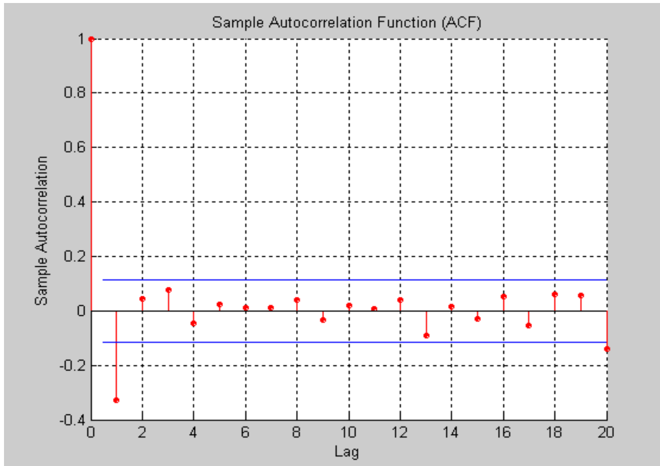

Figure 3. Autocorrelation

The volatility estimation as a function of sampling frequency shows that volatility estimation increases with sampling frequency, which conforms to the observation of microstructure noise.

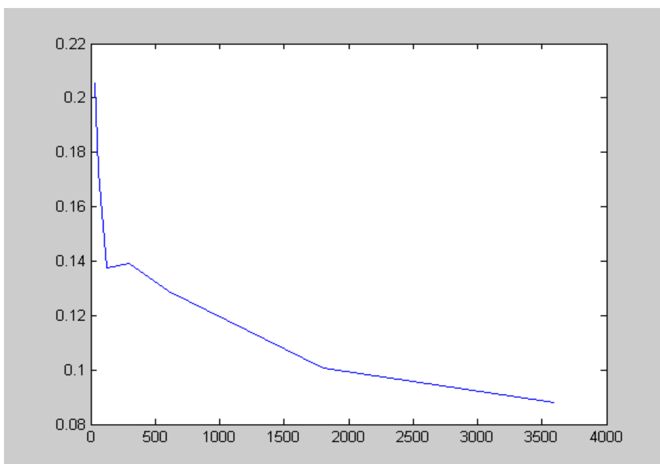

Figure 4. Volatility as a function of sampling frequency

3. Inverse Laplace Method and Results

As mentioned earlier, the motivation for a stochastic model for order book is to predict short-term behavior of various quantities which are useful in trade execution and algorithmic trading. We focus on three probabilities of interest as in Cont (2008): the probability of the mid-price increase, the probability of executing a limit order at the bid before the ask price moves and the probability of executing both a buy and a sell order at the best quotes before the mid-price moves. These quantities can be quickly obtained by computing an inverse Laplace transform (LT) of each respective conditional probability distribution. We confirm these quantities using Monte Carlo simulations, whose results agree with the Laplace transform computations and show that these quantities are well captured by the model. In practice, the Laplace transform method allows efficient computation of the probabilities of interest, bypassing the need for Monte Carlo simulation.

A. Birth death processes

Consider a birth-death process with birth rate λ and death rate μ_i in state $i \geq 1$, and let σ_b denotes the first passage time of this process to 0 given it starts from state b. The first passage time can be written as

$$\sigma_b = \sigma_{b,b-1} + \sigma_{b-1,b-2} + \cdots + \sigma_{1,0}$$

where $\sigma_{i,i-1}$ denotes the first passage time of the process from state i to $i-1$. Let f_i be the LT of the pdf of the first passage time from state i to $i-1$. We obtain the recursion as

$$\widehat{f}_i(s) = \frac{\mu_i}{\lambda + \mu_i + s} + \frac{\lambda \widehat{f}_{i+1}(s)\,\widehat{f}_i(s)}{\lambda + \mu_i + s}$$

From which we obtain

$$\widehat{f}_i(s) = \frac{\mu_i}{\lambda + \mu_i + s - \lambda \widehat{f}_{i+1}(s)}$$

Iterating on the above equation produces the continued fraction (CF),

$$\widehat{f}_i(s) = -\frac{1}{\lambda} \Phi_{k=i}^{\infty} \frac{-\lambda \mu_k}{\lambda + \mu_k + s}$$

Since each first passage time is independent, the LT of the first passage time can be written as

$$\widehat{f}_b(s) = \left(-\frac{1}{\lambda}\right)^b \left(\prod_{i=1}^{b} \Phi_{k=i}^{\infty} \frac{-\lambda \mu_k}{\lambda + \mu_k + s}\right)$$

B. Probabilities of interest

The first move in the mid-price occurs at the first passage time of the bid or ask queue to zero. If the spread (S) is greater than one, the mid-price moves at the time a limit order arrives inside the spread. Throughout this section, let $\sigma_b(\sigma_a)$ be the first passage time of orders in the best bid (ask) to zero. Furthermore, let $\varepsilon_b(\varepsilon_a)$ be the first passage time of remaining orders at time t in the best bid (ask) to zero. For a simple notation, λ denotes the limit orders arrival rate and μ denotes the market orders arrival rate and θ denotes cancellation rate of limit orders.

Given the initial configuration of the book, the mid-price increasing probability can be written as

$$P(\sigma_a - \sigma_b < 0)$$

σ_b and σ_a are birth death process with birth rate λ and death rate $\mu + i\theta$ in state $i \geq 1$. Using the result from the previous section, we can represent this probability as

$$\widehat{F}_{a,b}(s) = \frac{1}{s}\hat{f}_a(s)\hat{f}_b(-s)$$
$$P(\sigma_a - \sigma_b < 0) = F_{a,b}(0)$$

where $\widehat{f}_b(s) = \left(-\frac{1}{\lambda}\right)^b \left(\prod_{i=1}^{b} \Phi_{k=i}^{\infty} \frac{-\lambda(\mu + k\theta)}{\lambda + \mu + k\theta + s}\right)$

This is the special case when S = 1 and the case with S > 1 can be easily extended from this result.

Similarly, the probability that a limit order placed at the best bid is executed before the mid-price moves is given by

$$P(\varepsilon_b - \sigma_a < 0)$$

Since ε_b is a pure death process, the LT of the pdf of $\varepsilon_{i,i-1}$ has an exponential distribution.

$$\hat{g}_{i,i-1}(s) = \frac{\mu + \theta(i-1)}{\mu + \theta(i-1) + s}$$

Note that we assume that our order is not cancelled and the cancellation rate is $\theta(i-1)$. Therefore, we can represent the probability of executing a limit order before the price movement (S = 1) as

$$\widehat{F}_{a,b}(s) = \frac{1}{s}\hat{g}_b(s)\hat{f}_a(-s)$$
$$P(\varepsilon_b - \sigma_a < 0) = F_{a,b}(0)$$

where $\hat{g}_b(s) = \prod_{i=1}^{b} \frac{\mu + \theta(i-1)}{\mu + \theta(i-1) + s}$

We now compute the probability that two orders, one placed at the best bid price and one placed at the best ask price, are both executed before the mid-price moves, given that the orders are not cancelled. Using the result from previous analysis, this probability can be written as

$$P(\max(\varepsilon_b, \varepsilon_a) < \min(\sigma_b, \sigma_a)) = P(\varepsilon_a < \sigma_b, \varepsilon_b < \varepsilon_a) + P(\varepsilon_b < \sigma_a, \varepsilon_a < \varepsilon_b)$$

Let us define $h_{a,b} = P(\varepsilon_b < \sigma_a, \varepsilon_a < \varepsilon_b) = P(\varepsilon_b < \varepsilon_a < \sigma_b)$. Conditioning on the value of ε_b,

$$h_{a,b} = \int_0^\infty P(\varepsilon_b < \varepsilon_a < \sigma_b \mid \varepsilon_b = t) g_b(t) dt$$

where $g_b(t)$ is the inverse LT of $\hat{g}_b(s)$. Focusing on the first integrand in the above equation, it can be shown that

$$P(\varepsilon_b < \varepsilon_a < \sigma_b \mid \varepsilon_b = t) = \sum_{i=1}^{\infty} \sum_{j=1}^{a} P(\varepsilon_j < \sigma_i) P(X_b(t) = i \mid \varepsilon_b = t) P(W_a = j)$$

where $X_b(t)$ represents the number of orders at the bid and $W_a(t)$ represents the number of orders remaining at the ask at time t of initial $X_a(0)$. We already derived $P(\varepsilon_j < \sigma_i)$ from the probability of executing a limit order before the price movement. In addition, we can regard $P(X_b(t) = i \mid \varepsilon_b = t)$ as the number of orders at time t from initially empty state and it can be written as

$$P(X_b(t) = i \mid \varepsilon_b = t) = \frac{e^{-\Lambda(t)} \Lambda(t)^i}{i!}, \quad \Lambda(t) = \frac{\lambda(1 - e^{-\theta t})}{\theta}$$

Finally, $P(W_a = j)$ is the probability that a pure death process with death rate $\mu + (k-1)\theta$ in state $k \geq 1$ is in state j at time t, given it begins in state a. Therefore it can be written as

$$P(W_a = j) = \left(e^{Q_a t}\right)_{a,j}$$

where Q_a is the infinitesimal generator of this pure death process. Summarizing the above analysis, the probability of executing both a buy and a sell order at the best quotes before the mid-price moves is

$$P(\max(\varepsilon_b, \varepsilon_a) < \min(\sigma_b, \sigma_a)) = h_{a,b} + h_{b,a}$$

where $h_{a,b} = \sum_{i=1}^{\infty} \sum_{j=1}^{a} P(\varepsilon_j < \sigma_i) \int_0^\infty P(X_b(t) = i \mid \varepsilon_b = t) P(W_a = j) g_b(t) dt$

C. Numerical method

In order to calculate the probabilities in the previous section, the inverse LT of a CF must be implemented. Due to Euler (1773), there is relatively simple recursion formula for approximating CF.

$$
\begin{aligned}
&\widehat{F}(s) = \Phi_{n=1}^{\infty} \frac{a_n}{b_n} \approx \frac{P_n}{Q_n} \\
&P_0 = 0, \quad P_1 = a_1, \quad P_n = b_n P_{n-1} + a_n P_{n-2} \\
&Q_0 = 1, \quad Q_1 = b_1, \quad Q_n = b_n Q_{n-1} + a_n Q_{n-2}
\end{aligned} \tag{1}
$$

After this rational approximation, we apply a Fourier series method based on Hoog (1982) to calculate the inverse LT. It can be shown that the Bromwich integral in the inverse LT can be represented as

$$f(t) = \frac{1}{2\pi i} \oint \exp(st) \hat{F}(s) ds = \frac{\hat{F}(\gamma)}{2} + \sum_{k=1}^{\infty} \hat{F}(\gamma + \frac{ik\pi}{T}) \exp(\frac{ik\pi t}{T}) = \sum_{k=0}^{\infty} a_k Z^k$$

where $s = \gamma + i\omega = \gamma + ik\pi / T$. Instead of applying the epsilon algorithm to the partial sum of the above series, we use the quotient-difference algorithm which makes the rational approximation available in the form of a CF. This enables the above series to be evaluated at any particular time value by recursion. In particular, given the above power series, we want to calculate the CF

$$v(Z) = \frac{d_0}{1+} \; \frac{d_1 Z}{1+} \; \frac{d_1 Z}{1+} \; \cdots$$

In practice, this means determining $v_{2M}(Z)$ where

$$
\begin{aligned}
&u_{2M}(Z) = \sum_{k=0}^{2M} a_k Z^k, \quad v_{2M}(Z) = \frac{d_0}{1+} \; \frac{d_1 Z}{1+} \; \frac{d_1 Z}{1+} \; \cdots \frac{d_{2M} Z}{1} \\
&u_{2M}(Z) - v_{2M}(Z) = O(Z^{2M+1})
\end{aligned}
$$

The coefficient d_k can be calculated using the quotient-difference algorithm. Once we obtain d_k , the same rational approximation as in Eq (1) can be done again to get

$$f(t) = \frac{e^{\gamma t}}{T} \mathrm{Re}\left\{ \frac{P_{2M}}{Q_{2M}} \right\}$$

D. Simulation results

Since the numerical method in the previous section is based on the one-sided LT and the probabilities of interest are two-sided LT, we shift the distributions by multiplying by a complex exponential and perform the inverse LT.

Tables 1, 2, 3 and 4 compare Monte Carlo simulation results (10000 iterations) and the LT method for the probability of a mid-price increase when S=1. Tables 9 and 10 show the same probability when S=2. The row represents the number of orders at the best bid and the column represents the number of orders at the best ask. We can confirm that the LT method successfully replicates the Monte Carlo simulation and the order book dynamics are well captured by the model.

Tables 5 and 6 compare Monte Carlo simulation results (10000 iterations) and the LT method for the probability of executing a bid order before mid-price movement when S=1. Table 11 and 12 show the same probability when S=2.

Tables 7 and 8 compare Monte Carlo simulation results (10000 iterations) and the LT method for the probability of making the spread. We only demonstrate the probability when S=1. As noted in Cont (2008), for a fixed value of a, as b is increased, the probability of making the spread is not monotone.

	1	2	3	4	5
1	0.5	0.3368	0.2615	0.2188	0.1912
2	0.6637	0.5003	0.4085	0.3504	0.3105
3	0.7392	0.5922	0.5003	0.438	0.393
4	0.7819	0.6503	0.5627	0.5003	0.4537
5	0.8096	0.6903	0.6078	0.547	0.5004

Table 1. probability of mid-price increase, S=1 (Laplace)

	1	2	3	4	5
1	0.4992	0.3594	0.2929	0.2455	0.2189
2	0.6505	0.5009	0.4171	0.3753	0.3353
3	0.7163	0.5807	0.5018	0.4481	0.4051
4	0.7482	0.6325	0.5611	0.5046	0.4534
5	0.781	0.6666	0.5977	0.5395	0.4996

Table 2. probability of mid-price increase, S=1 (Monte Carlo)

	1	2	3	4	5
1	0.5291	0.4927	0.4868	0.4735	0.454
2	0.5627	0.5321	0.5168	0.4938	0.4705
3	0.5735	0.5531	0.5336	0.4914	0.5182
4	0.5743	0.5507	0.5675	0.5442	0.531
5	0.5744	0.5351	0.5674	0.5593	0.4531

Table 3. probability of mid-price increase after 2 movements, S=1 (Monte Carlo)

	1	2	3	4	5
1	0.4663	0.3729	0.3183	0.2994	0.2834
2	0.559	0.4671	0.4056	0.3636	0.3326
3	0.6098	0.528	0.4532	0.4158	0.3878
4	0.6502	0.5511	0.5046	0.485	0.4304
5	0.6645	0.5755	0.5279	0.4969	0.4885

Table 4. probability of mid-price increase after 10 movements, S=1 (Monte Carlo)

	1	2	3	4	5
1	0.5081	0.7038	0.7992	0.8531	0.8866
2	0.3665	0.5595	0.6726	0.7448	0.7939
3	0.2998	0.4756	0.5886	0.6661	0.7218
4	0.2602	0.4203	0.5288	0.6066	0.6648
5	0.2332	0.3807	0.4838	0.5601	0.6187

Table 5. probability of execution of bid order, S=1 (Laplace)

	1	2	3	4	5
1	0.6159	0.7829	0.855	0.8995	0.922
2	0.4702	0.6622	0.7563	0.8086	0.8486
3	0.3966	0.5799	0.6779	0.744	0.7851
4	0.3593	0.5184	0.6161	0.6869	0.7433
5	0.3198	0.4724	0.5738	0.645	0.6965

Table 6. probability of execution of bid order, S=1 (Monte Carlo)

	1	2	3	4	5
1	0.2756	0.3194	0.3207	0.3115	0.2998
2	0.3194	0.3994	0.4201	0.4211	0.4146
3	0.3207	0.4201	0.4561	0.4676	0.4683
4	0.3115	0.4211	0.4676	0.4877	0.4949
5	0.2998	0.4146	0.4683	0.4949	0.5076

Table 7. probability of making the spread (Laplace)

	1	2	3	4	5
1	0.2301	0.2606	0.2632	0.2524	0.2399
2	0.2601	0.3235	0.3363	0.3349	0.3273
3	0.2544	0.3322	0.359	0.3685	0.3624
4	0.2582	0.3314	0.3635	0.3803	0.3834
5	0.2413	0.3203	0.3647	0.3805	0.3931

Table 8. probability of making the spread (Monte Carlo)

	1	2	3	4	5
1	0.4986	0.4041	0.3786	0.3703	0.367
2	0.5946	0.4996	0.4706	0.4596	0.4554
3	0.62	0.5287	0.4997	0.4885	0.4837
4	0.6281	0.5392	0.5097	0.5005	0.495
5	0.6276	0.5427	0.5173	0.505	0.5

Table 9. probability of mid-price increase, S=2 (Laplace)

	1	2	3	4	5
1	0.5106	0.43	0.4196	0.421	0.423
2	0.5643	0.4999	0.488	0.4737	0.4749
3	0.5861	0.5206	0.4976	0.4954	0.4888
4	0.5899	0.52	0.5083	0.4925	0.4879
5	0.584	0.5223	0.5128	0.5139	0.4899

Table 10. probability of mid-price increase, S=2 (Monte Carlo)

	1	2	3	4	5
1	0.1502	0.1816	0.1909	0.1942	0.1956
2	0.0386	0.0522	0.0573	0.0595	0.0605
3	0.0131	0.019	0.0218	0.0231	0.0237
4	0.0053	0.0081	0.0096	0.0104	0.0108
5	0.0025	0.0039	0.0047	0.0052	0.0055

Table 11. probability of execution of bid order, S=2 (Laplace)

	1	2	3	4	5
1	0.1695	0.1905	0.1983	0.1897	0.1945
2	0.0486	0.0602	0.057	0.0622	0.0602
3	0.0162	0.0206	0.0231	0.0236	0.025
4	0.0058	0.0093	0.0098	0.0131	0.0119
5	0.0041	0.0047	0.0057	0.0052	0.0055

Table 12. probability of execution of bid order, S=2 (Monte Carlo)

4. Application to Trading Strategies

With the inverse Laplace method, we have at our disposal interesting conditional probabilities that we can apply to create high-frequency trading strategies. We treat two strategies, betting on an increase in mid-price and also making the spread. For now, we are still considering an order book defined by the parameters estimated from the stock “Skyperfect Communications” used by Cont et al. in their article. Note that executing a strategy can have a direct influence on the shape of the order book, for example our actions can result in a change of the mid-price. Thus it is impossible to back-test our strategies on real data, because we would neglect our market impact. Thus we instead test our strategies on trading data simulated with our trades integrated “real time” so to speak.

A. Betting on an increase of the mid-price

This strategy is quite simple, and the symmetric strategy can also be used. If the volume at the best ask is low (in this model we consider that it is equal to 1), and the spread is equal to one tick, we can send a market order. Then the spread becomes equal to 2. If the volume at the best bid is high (more than 3) and the volume at the new best ask is low (for instance 1), then we have a good chance that the mid-price increases (around 0.62).

As for the exit strategy, we do not want to lose too much money (two ticks or more). Thus, if the volume at the best bid becomes equal to 1, we sell the stock at the market price, and we lose one tick. But if the best bid price increases by more than two ticks, we make a profit if we sell the stock.

That is the strategy that we choose to use first. Here is the P&L, in ticks, as a function of the number of trades (over 10 simulated days):

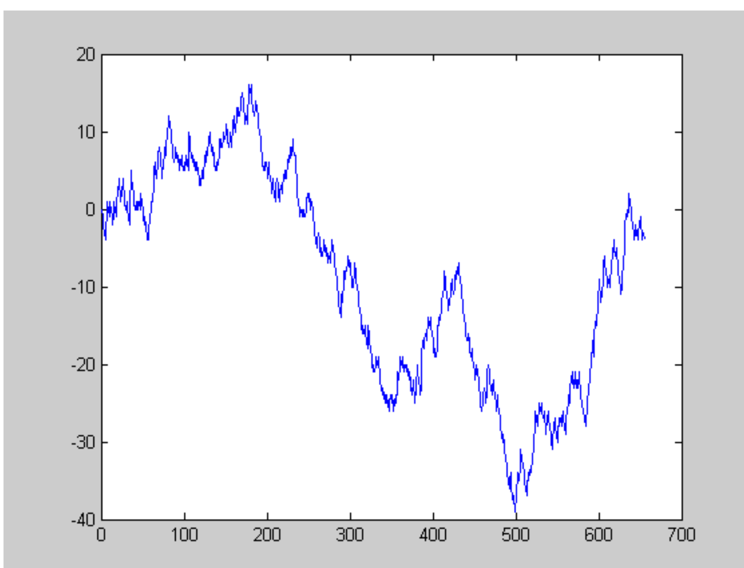

Unfortunately, this strategy does not seem to be good enough. Without transaction costs, the average loss is 0.006 ticks by trade, and the 95% confidence interval [-0.05; 0.05] shows us that this strategy does not generate a reliable profit (except in a bull market). The distribution of profits per trade, shown below, demonstrates that our earnings on successful trades are invalidated by too frequent losses of one tick.

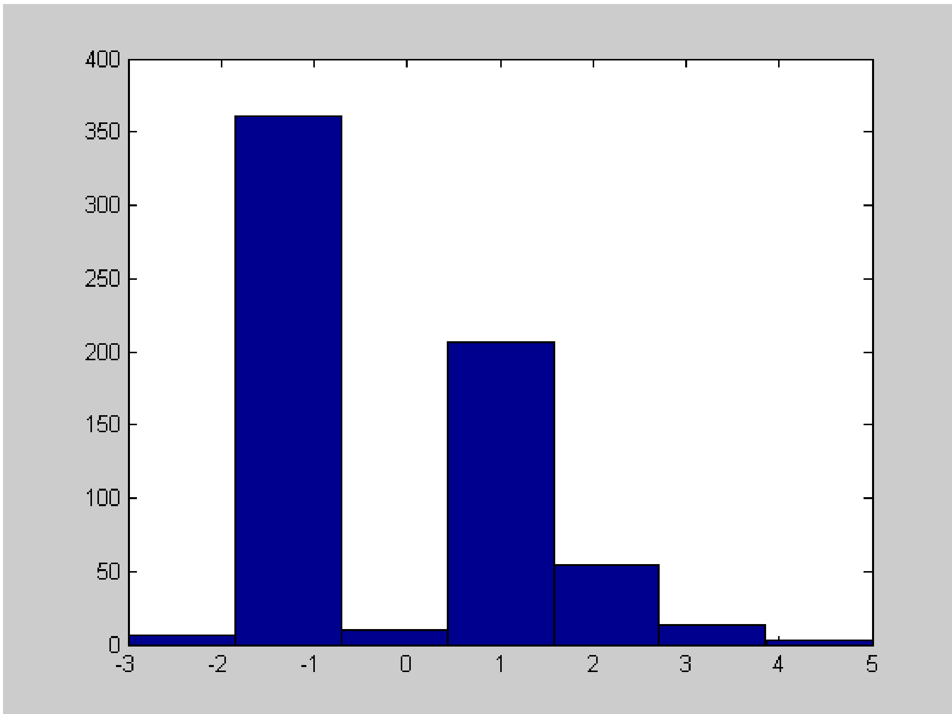

Variants derived by changing the initial strategy (entering the position) according to the probabilities we computed in the previous part, as well as the exit strategy (allowing losses of 2 ticks) did not improve results.

Despite the poor performance detailed above, this does not imply that this strategy as a whole is ineffective. These tests were run based on the Skyperfect Communications parameters, and it is clear that these parameters, if changed, would have a huge influence on the relevant probabilities. For instance, if we decrease the rate of market orders (μ) or cancellations (θ) compared to the rate of limit orders, we decrease the probability that the volume at the best bid decreases before a higher limit order arrives, and thus we decrease the probability to lose 1 tick.

Thus if we expand this strategy to the whole universe of stocks, we can guess that there exist stocks for which a profit can be generated, even when taking transaction costs into account.

B. Making the spread

Another strategy is to try to make the spread, i.e. placing an order at the best bid and another one at the best ask and hoping they will be both executed before the mid-price moves. A successful set of trades yields the spread in profit, in this case 1.

Computing this probability using both inverse Laplace transform and Monte-Carlo methods yields a very low rate of success, except when the volume at the best bid and the best ask is very high, e.g. more than 5 times the average order size. Indeed, in this case, the probability to succeed is slightly greater than ½, due to an increased sense of "price inertia".

If we do not succeed, our potential loss can be high depending on our exit strategy. For example, if one of our orders is executed but then the mid-price moves against us, an immediate exit will cost us two ticks (cancelling our second order and replacing it by a market order). We can avoid this by cancelling our second order if the volume at the opposite best quote becomes equal to 1. Then we lose only 1 tick but this also decreases our probability to succeed.

Once again we find that our profit expectation is not high enough. However, with other stocks, we can expect to find some securities whose parameters create a sufficiently large probability to make the spread. Here, we see that the ratio between θ or μ and the rates or limit orders are still very important.

5. Challenges and Further Work

We have detailed above the process used to validate Cont's stochastic order book model and use inverse Laplace transform methods to generate trading strategies. The single major challenge we faced in pursuing this project was the application to real data, specifically to the American stock markets. The original intent was to use trading data provided by Evnine and Associates to estimate the model parameters, and compared the results to those of Cont, which were derived from a security on the Tokyo Stock Exchange.

However, we were unable to acquire a complete data set, as Cont's framework requires both trade data as well as level 2 order book data. Unfortunately, we had level 2 data for dates for which we did not have trade data, and vice versa. Thus we were not able to use the provided datasets for parameter estimation.

That having been said, the next step in this research is clear – to obtain a suitable data set, perform a suitable parameter estimation procedure, and then compute the aforementioned probabilities for the new set of parameters. This having been done, we can not only compare the parameters of securities on the NYSE/NASDAQ versus those on the TSE, but hopefully will also be able to find stocks for which our computed probabilities can lead to profitable trading strategies.