

# Simulation of a Limit Order Driven Market

Julian Lorenz\* and Jörg Osterrieder †

June 3, 2008

## Abstract

We present an order flow model framework for limit order driven markets. Different from previous models we explicitly model a reference price process that “sweeps” the limit order book as it fluctuates up and down. Our framework allows us to use any stochastic process to model this reference price and very general specifications of the limit order flow. We believe that this framework can fruitfully combine order flow models with well-studied models for stock price processes and provides a step towards developing realistic, yet tractable models for complex limit order driven markets. We use public order data from SWX as an example to estimate the model parameters.

Most modern stock exchanges use automated auction systems to facilitate trading. Examples of purely automated auction markets include Xetra (Frankfurt/Germany), SWX (Swiss Stock Exchange), Hong-Kong, and ECNs such as Island. Such kind of markets operate as purely order-driven electronic stock markets without market makers, and liquidity is provided only by public orders. Contrary, in a quote-driven market, designated market makers supply liquidity continuously, quoting bid and ask prices. Hybrid trading structures, for instance NYSE and NASDAQ, may combine both features (see e.g. Domowitz [1993] for a taxonomy of trading systems).

---

\*ETH Zürich, Universitätstrasse 6, 8092 Zürich, Switzerland; jlorenz@inf.ethz.ch. Supported by UBS AG.

†ETH Zürich, Universitätstrasse 6, 8092 Zürich, Switzerland; joerg.osterrieder@math.ethz.ch.

## Modeling Limit Order Book Markets

A tractable model of a limit order driven market would be of tremendous value. It would not only help to better understand the dynamics of limit order driven markets from a theoretical point of view, but would be very valuable in practice as well. It would make it possible to evaluate and fine-tune intraday trading and order placement strategies, instead of backtesting solely using historical order flow. Clearly, trading strategies that use information from the order book are superior to such that only use transaction information (stock price), since every order that is submitted may be interpreted as new information.

However, the inherent complexity of such kind of a market makes it very difficult to develop a model that is at the same time realistic, analytically or at least computationally tractable, and can be reasonably well fitted to actual market activity. While very accurate models for stock price processes have been developed, modeling order book dynamics is still in its infancy with some notable exceptions.

Smith et al. [2003] model order driven markets under the assumption of random order flow. The general idea is that the aggregated effect of all market participants are flows of buy and sell orders with certain distributions for limit price, size etc. Smith et al. assume that limit orders (of fixed order size) arrive with the same frequency at every possible price. Using analytical methods and numerical simulations they estimate quantities of interest such as the expected order book shape.

A slightly different approach is used by Chiarella and Iori [2002]: they simulate an order driven market not by modeling aggregate order flows, but by setting up a number of heterogeneous agents who trade with each other via an order matching mechanism. The agents set their limit prices individually according to some fixed trading strategies, ranging over fundamentalist, chartist and “noise-induced” (uninformed) components. Chiarella and Iori argue that such kind of agent based models exhibit similar features as real markets. However, from practical point of view it is rather hard to fit the model to historical data.

## Reference Price Process

We propose a modeling approach that is similar to Smith et al. [2003], but differs in a very important aspect. We also start by modeling limit orders as random order flow, governed by some kind of distribution. But instead of modeling the flow of market orders (respectively, marketable limit orders) that generate trades with the limit orders waiting in the order book, we model a *reference price process* that “sweeps” the limit order book as it fluctuates up and down. This process is to be understood as a virtual reference, not as the actual stock price process. Every time the reference price process sweeps over bid or ask limit orders stored in the book, these orders are executed at their limit price. Consequently, the resulting sequence of transaction prices is very closely linked to that virtual reference price process, though not exactly identical. Instead of a continuous stock price path this model generates a discrete sequence of transactions with varying volume – a realistic feature indeed.

In this paper, we shall only give a brief informal description and analysis of the model. For a mathematically rigorous presentation the reader is referred to Osterrieder [2007]. In this paper we mainly focus on numerical results for public order book data from SWX.

Let us describe the model more explicitly: The reference price process follows a continuous time random walk, which we may choose to make the dynamics as realistic as we wish. In the most general setting it may follow any semimartingale dynamics, but for concreteness here we shall consider processes of the form

$$dX_t/X_t = \mu_t dt + \sigma_t dW_t \quad (1)$$

where  $W$  is a standard Brownian motion.  $\mu_t$  and  $\sigma_t$  are – possibly time-dependent – drift and volatility coefficients.

Limit orders are submitted relative to this price and arrive in discrete time. This limit order flow is modeled by a Poisson process, which describes the number of orders that arrive in a certain subset of a three-dimensional space spanned by time, relative limit order price and limit order size (see Osterrieder [2007] for details). The idea behind this model is that there is a large number of potential buyers and sellers who act independently from each other, submitting limit orders every now and then. Other than Smith

et al. [2003], this model allows for the more realistic case of non-uniform order sizes.

Buy orders are only allowed to be placed below the reference price and sell orders above. We assume that the flow of limit order submission does not influence the reference price (1). Whenever the reference price reaches a limit order that is stored in the book, this limit order is completely executed and removed from the book. This means that we do not consider time priority or the possibility of partial execution. We model order cancellation by introducing a cancellation rate, which governs the probability of any given order being canceled in a certain period of time.

This framework is very versatile. While the reference price process is rather straightforward, the order arrival process and the cancellation process make the model very complex in general. The specification of the order arrival process in the three-dimensional space of time, relative limit price and order size allows for very general dependencies between these variables. In order to obtain a practical model, further assumptions will have to be made, which we will discuss below.

The key advantage of our model is that the stock price is not a process that is a *resulting* process of other model input parameters (as in the models of Smith et al. [2003] or Chiarella and Iori [2002], for instance), but is modeled explicitly. Instead of requiring a great effort to calibrate the model such that the resulting stock price process is realistic and fitted to the observable market volatility, in our model the stock price process is realistic from the beginning: we can either use a simple Geometric Brownian motion with constant volatility, or more complex models with stochastic volatility etc. in order to capture realistic features. We believe that this approach can fruitfully combine order flow models with well-studied models for stock price processes. One can think of the reference price process as a continuous stream of incoming market orders, which manifest themselves once they hit either the bid or the ask price.

## Theoretical Properties

The model outlined above allows for some interesting analytical analysis. As shown by Osterrieder [2007], the order book can be described as the difference of two doubly stochastic Poisson processes at every point in time.

Equivalently, a characterization as the difference of two infinite sums of Bernoulli distributed random variables can be given. From there various quantities and properties of interest can be derived, for instance the relation between the volatility of the stock price and the volume of orders in the order book, the distributions of bid and ask prices, or the time-to-execution of a limit order in the book.

A very interesting application is the relation between the distribution of the relative limit price of arriving orders and the distribution of traded volume. Bouchaud et al. [2002] observed that the density of the distribution of the distance of arriving orders relative to the current price follows a power-law rule with parameter 1.6 (also see Zovko and Farmer [2002]). Plerou et al. [2001] investigated the distribution of traded volume. For  $V_{\Delta t}$ , the traded volume in an infinitesimal small time interval, it was found as

$$P(V_{\Delta t} > v) \sim 1/v^\lambda, \quad v \in \mathbb{R}_+,$$

where  $\lambda = 1.7 \pm 0.1$ .

Interestingly, in the order book model outlined above, it can be shown under fairly general conditions on the reference price process that indeed a value of  $\lambda = \frac{5}{3} \approx 1.67$  follows from 1.6 as the parameter in the order arrival density. The precise formula is

$$\lambda = \beta/(1 - \gamma),$$

where  $\beta$  is a parameter which describes a property of the stock price process. Use  $\beta = 1$  in the equation above and you will be able to recover 1.6 and 1.67 respectively. The reader is referred to Osterrieder [2007] for more theoretical results in this model.

## Data

Our dataset contains the history of trades and public orders<sup>1</sup> of stocks traded on SWX between May 2005 and January 2007.

The Swiss Exchange (SWX) is one of the world's most technologically advanced securities exchanges. It is a purely order-driven market, and was the first stock exchange to launch a fully computerized, integrated trading

---

<sup>1</sup>Provision of public historical SWX order book data by UBS AG is gratefully acknowledged.

system in August 1996. With a total exchange turnover of 1.8 trillion CHF in 2007, and a market capitalization of 1.4 trillion CHF, SWX is among the top ten international stock exchanges. Trading takes place continuously during the trading day with automatic matching of orders based on the usual rules of price and time priority. There are no market makers or floor traders with special obligations or with privileged access to trading. Liquidity is provided exclusively through limit orders, and the whole central computerized order book is public and available in real time.

Besides the public order book and public transactions, SWS allows off-order book trading under certain circumstances. We do not consider this kind of order flow and we do not take account of hidden orders or transactions during the opening and closing auctions.

For each stock, our dataset reports the transaction data (time stamp, price, and volume in number of shares), the order flow entering the book (time stamp, price, volume in number of shares, and side, i.e. bid or ask), and order cancellations (time stamp, price, and canceled volume in number of shares). Each order has a unique ID, by which it can be traced throughout the trading day. Using this information, the state of the order book can be reconstructed for any given point of time in the sample period. For the presentation in this paper, we choose UBS (UBSN/CH0024899483). Exhibit 1 shows a typical order book.

## Model Implementation

Recall that in our framework we need to specify the parameters of the stock price process (1), the parameters of the order arrival process as well as the cancellation rates. As mentioned above, our model allows for very complex specifications with very general dependencies between relative limit price, order size and time.

For fitting the reference price process to historical stock price data a lot of work has been done before. As argued above, here we see one of the key advantages of our model. In the simplest implementation, we model the reference price as a Geometric Brownian motion with constant volatility. While constant volatility is certainly a very simplified assumption, for very short term intraday time horizons we can expect reasonable results. Improved models could use stochastic volatility models or also mean-reversion

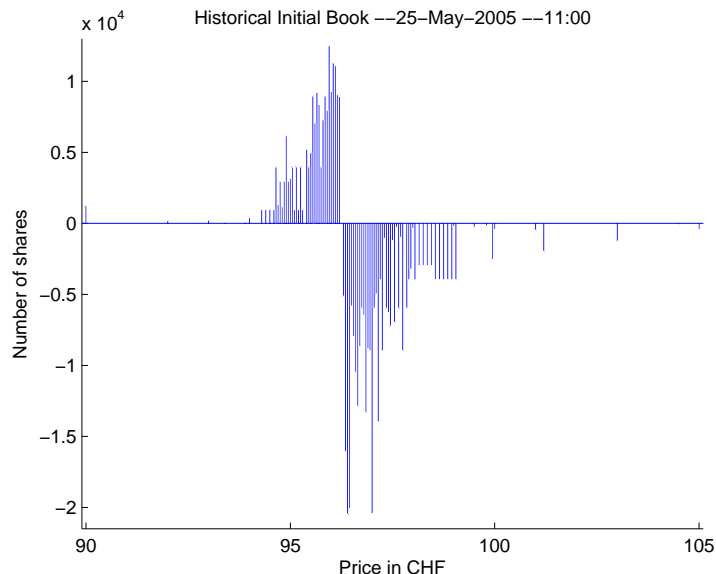


Figure 1: Typical order book of UBS (UBSN/CH0024899483).

or momentum based price processes.

The order arrival process can be empirically estimated by observing arrival times, relative prices and sizes of all incoming orders. In the following, we will present the results for the sell (ask) order book of UBS as an example. The results for the buy (bid) side of the order book are similar.

We first analyse the frequency of arriving limit orders. Not surprisingly, the rate varies significantly across the day. Similar to the well-known U-shape intraday pattern of traded volume, we find strong U-shape seasonality for the order arrival frequency as well. Exhibit 2 shows this intraday pattern for UBS.

Our analysis shows that the inter-arrival times of limit orders for UBS are Weibull( $k, \lambda$ ) distributed with shape parameter  $k \approx 0.7$ , see Exhibit 3. The Weibull distribution is a generalization of the exponential distribution, which is recovered as a special case with the shape parameter  $k = 1$ . Compared to exponentially distributed inter-arrival times, Weibull distributed inter-arrival times with shape parameter  $k < 1$  lead to order arrivals that are much more clustered in time. Exhibit 4 illustrates this behavior.

We next analyse the distribution of order sizes of arriving limit orders.

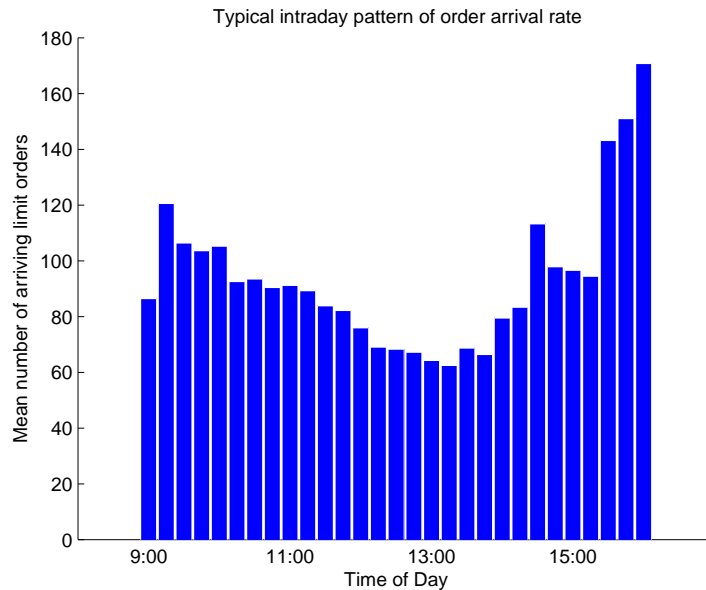


Figure 2: Typical U-shape intraday pattern of the order arrival rate for UBS.

As it is well known, orders have a tendency to cluster in size. For instance, orders sizes of multiples of 100 or 1000 are much more frequent than odd-sized orders. Because of this irregularity, it is hard to fit the empirical distribution of order sizes to any “standard” probability distribution. For our order book simulations, we found it more appropriate to sample order sizes from the empirical distribution in Exhibit 5 directly. Interestingly, we found that order sizes also feature a U-shaped intraday pattern (see Exhibit 6), similar to the intraday pattern of the order arrival rate in Exhibit 2.

For the (unconditional) prices of arriving limit orders, as Bouchaud et al. [2002], Zovko and Farmer [2002] we find that the limit prices relative to the last transaction price follow a power-law like distribution (see Exhibit 7). Our empirical analysis showed that there seems to be a slight tendency of smaller orders being placed closer to the best bid/ask and larger orders deeper in the book.

For the estimation of the cancellation rates of limit orders we must take into account the “censoring effect”: Simply considering the times between order submission and order deletion (given a deletion occurred) for all limit



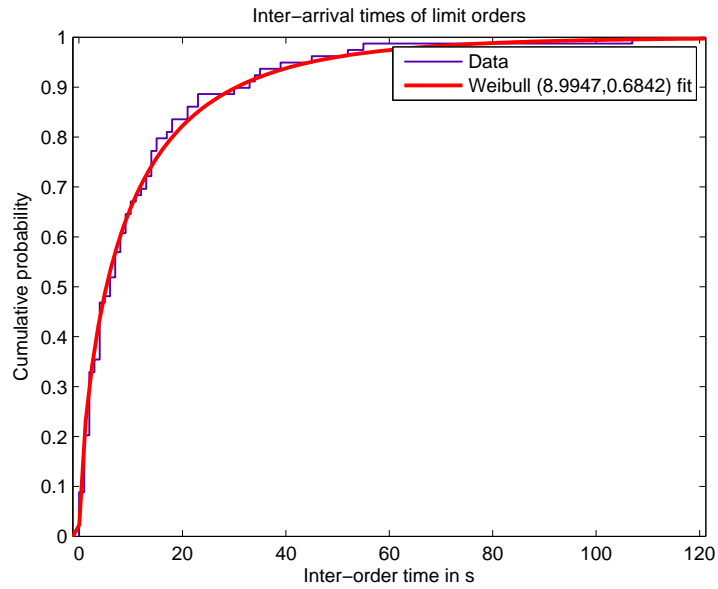


Figure 3: Inter-arrival times of limit orders follow a Weibull( $k, \lambda$ ) distribution with shape parameter  $k < 1$ .




Figure 4: Upper sequence has exponentially distributed inter-arrival times, and lower sequence Weibull distributed inter-arrival times with shape parameter  $k = 0.5$  with same expectation. Weibull-distributed inter-arrival times with  $k < 1$  result in more clustered arrival times.

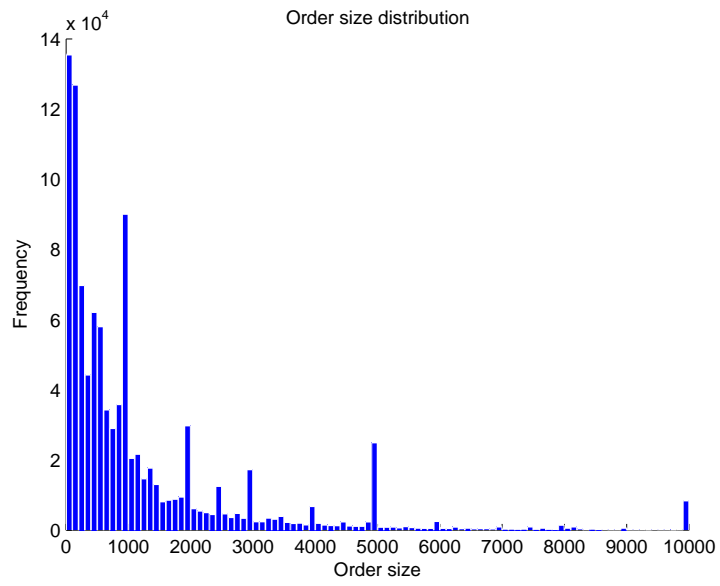


Figure 5: Histogram plot of limit order sizes for UBS.

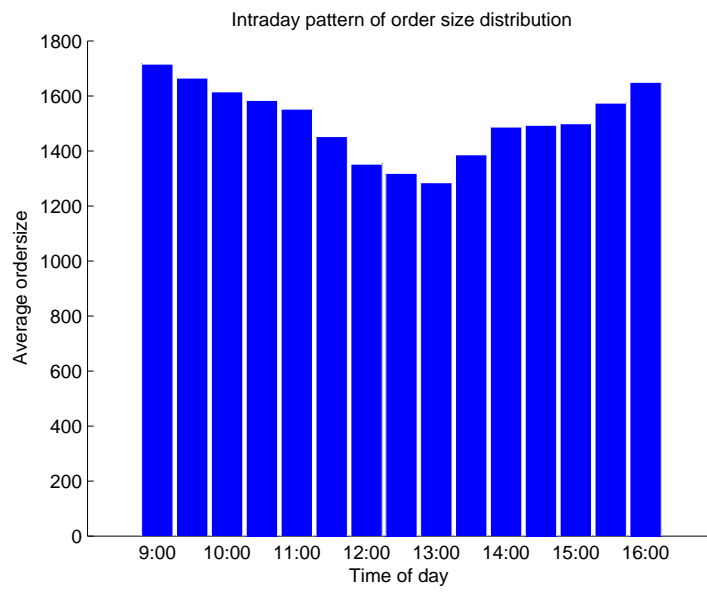


Figure 6: Typical U-shape intraday pattern of order sizes for UBS.

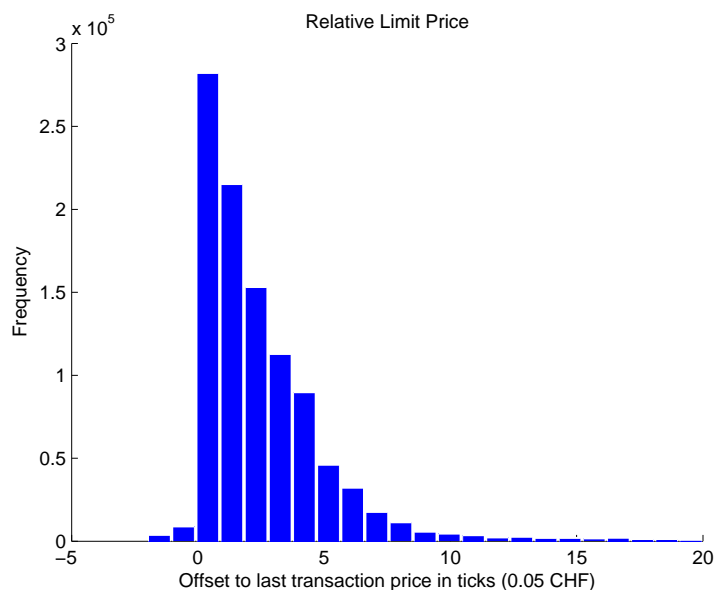


Figure 7: Limit prices relative to the last transaction price for UBS limit orders.

orders neglects the effect that we only see those orders being deleted which were not executed. In terms of *survival analysis*, this corresponds to censored data points (Type III, or *random censoring*). The well-known *Kaplan-Meier estimator* is then the standard technique to estimate times-to-cancellation and cancellation rates.

Exhibit 8 shows the distribution of the time-to-cancellation for arriving UBS limit orders, as obtained by the Kaplan-Meier estimator. A big fraction of arriving limit orders are deleted from the book very quickly, if they are not executed. Furthermore, multiples of 30 seconds occur more frequent. Because of this anomaly, it is again difficult to fit the data to any standard distribution and we found sampling from the empirical distribution to yield better results. Our analysis showed that there seems to be a slight dependency between order size and time to cancellation, with larger orders and especially orders of certain “chunk” sizes (e.g. 1000, 5000, ... shares) being canceled much less frequent than smaller orders. Similarly, orders placed deep into the book with limit prices far away from the best bid/ask also seem to remain slightly longer in the book before they are canceled.

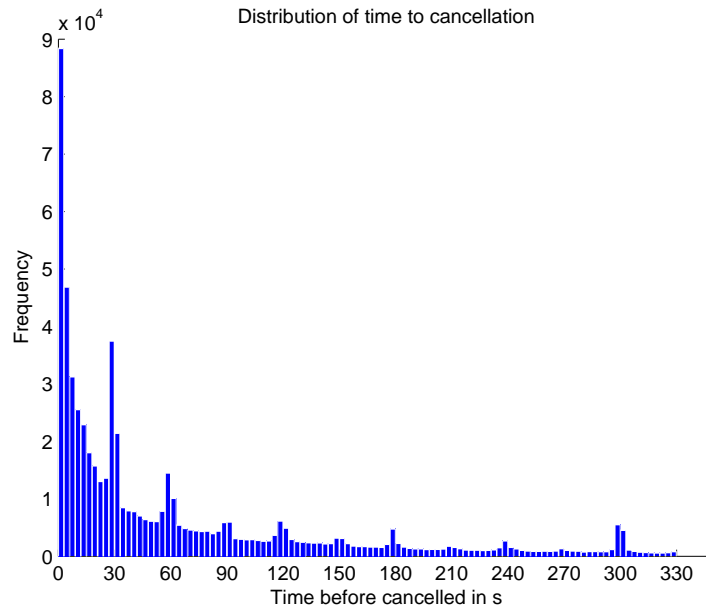


Figure 8: Distribution of the time-to-cancellation for arriving UBS limit orders using the Kaplan-Meier estimator.

An important observation is that the cancellation rate of limit orders stored in the book is much lower than the cancellation rate of newly arriving limit orders. That is especially important if we want to simulate the evolution of the order book *starting from some certain initial state* (for instance, the current state of some real market order book). The reason is a classical “survivorship bias”. Many of the newly arriving orders are from impatient market participants who delete their orders very quickly if they are not filled (see Exhibit 8). On the other hand, orders that are stored in the book are typically orders from more patient market participants and tend to remain in the book longer.

### Order Book Shape and Impact Function

Exhibits 9 and 10 show a historical order book evolution of UBS (ask side) between 10am and 11am on 08/03/2005 in comparison with one sample path (60 minutes) of our order book simulation starting from the same state of the order book at 10am. The Exhibits show one random sample path, and hence

the stock price process in the simulation and in reality (a-posteriori) are of course different – but the general behavior of the order flows are similar. If we want to optimize intraday trading or order placement strategies, we simulate for a collection of sample paths. Furthermore, we can use different market parameters to assess performance in various scenarios.

Exhibit 11 compares the historical and simulated shape of the orderbook for UBS (ask side), showing that our simulation indeed gives reasonable results. The order book shape corresponds to a *temporary impact function*, which is the price concession to buy a certain number of shares using a market order (immediate execution), exhausting the liquidity available at the first price levels of the order book until we have filled the desired quantity.

Note that if we want to evaluate a trading strategy in our simulation, in the formulation (1) the reference price is not affected by our own strategy's activity. If we want to incorporate permanent price impact of our trading strategy, we can indeed easily add a feedback term to the dynamics (1) of the reference price. More detailed results on no-arbitrage considerations in such a large trader setup can be found in Osterrieder [2007].

## Conclusion

We presented an order flow model framework for limit order driven markets. Limit orders are modeled as random order flow. We allow here for very complex specifications with very general dependencies between relative limit price, order size and time of arriving limit orders. This is supplemented by explicitly modeling a reference price process that “sweeps” the limit order book as it fluctuates up and down. Our framework allows us to use any stochastic process to model this reference price. In particular, there is no need to explicitly model market order flow. We believe that this framework can fruitfully combine order flow models with well-studied models for stock price processes and provides a step towards developing realistic, yet tractable models for complex limit order driven markets.

## References

Bouchaud, J.-P., M. Mezard, and M. Potters (2002). Statistical properties of stock order books: empirical results and models. *Quant. Finance* 4(6), 251–256.




Figure 9: The upper panel shows order arrivals (historical data, ask side only) and the transaction price path for UBS between 10am and 11am on 08/03/2005, and the lower panel one sample path (60 minutes) of our order book simulation. Black dots are arriving orders (size information not shown), and the blue curve the price path.

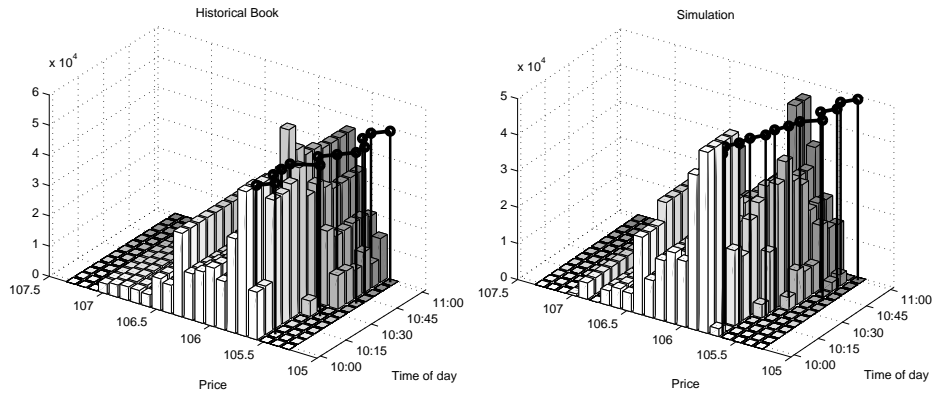


Figure 10: Historical order book evolution (ask side) of UBS (left) between 10am and 11am on 08/03/2005 (snapshots taken every 5 minutes), and one sample path (60 minutes) of our order book simulation (right). The black connected line is the stock price path.

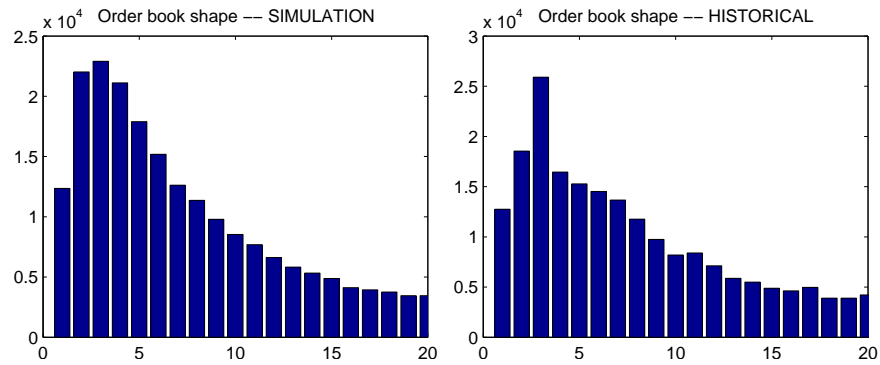


Figure 11: Historical and simulated shape of the orderbook for UBS (ask side).

- Chiarella, C. and G. Iori (2002). A simulation analysis of the microstructure of double auction markets. *Quantitative Finance* 2(5), 346–353.
- Domowitz, I. (1993). A taxonomy of automated trade execution systems. *Journal of International Money and Finance* 12(6), 607–631.
- Osterrieder, J. (2007). *Arbitrage, the limit order book and market microstructure aspects in financial market models*. Ph. D. thesis, ETH Zurich. DISS. ETH Nr. 17121.
- Plerou, V., P. Gopikrishnan, X. Gabaix, L. A. N. Amaral, and H. E. Stanley (2001). Price fluctuations, market activity and trading volume. *Quantitative Finance* 1(2), 262–269.
- Smith, E., J. D. Farmer, L. Gillemot, and S. Krishnamurthy (2003). Statistical theory of the continuous double auction. *Quant. Finance* 3(6), 481–514.
- Zovko, I. and J. D. Farmer (2002). The power of patience: a behavioural regularity in limit-order placement. *Quantitative Finance* 2(5), 387–392.