MODELLING HIGH-FREQUENCY FINANCIAL DATA

Alet Roux

Submitted in partial fulfilment of the requirements of the degree

Master of Philosophy in Statistical Science

at

Darwin College

UNIVERSITY OF CAMBRIDGE

Supervisors:

Prof. L. C. G. Rogers (external)
Dr. J. R. Norris (internal)

July 2003

Abstract

We are concerned with the number and frequency of transactions on financial markets. Using transactions data on FTSE 100 index futures in 2002 from LIFFE, we illustrate the stylized features of transaction tick data. We find effects pointing to the presence of strong seasonal patterns, long-term serial dependence, clustering as well as strong interdependence between asks and trades and bids and trades, respectively.

We also develop a model for the times and frequencies of transactions. We propose that transactions in individual assets form inhomogeneous Poisson processes, their intensities being generated by a continuous-time Markov chain describing the level of activity in the market. In addition, an iterative maximum likelihood estimation procedure is developed; we use it with success to model transactions in FTSE 100 index futures.

Acknowledgements

During the course of my practical project, I have received tremendous assistance from my supervisors. I am greatly indebted to my external supervisor, prof. L. C. G. Rogers for sharing his insights and experience on suitable models and estimation methodologies. Dr. J. R. Norris, my M.Phil. supervisor, has also, during our brief discussions, asked thought-provoking questions, leading to the inclusion of material on the validity of the Poisson assumption.

I am also grateful to the Cambridge Commonwealth Trust and the Commonwealth Scholarship and Fellowship Programme for financial support during my period of study at Cambridge.

Alet Roux
Cambridge, 3 July 2003

Contents

List of Tables

List of Figures

Chapter 1

Introduction

In this work, we study the number and frequency of transactions on financial markets. In addition to considering the stylized features of such count data, we develop a modelling framework in aid of our understanding of the processes generating activity on the trading floor.

During their opening hours, financial markets operate on a continuous, high frequency basis. Traditionally, however, virtually all available data sets on market activity were based on discrete sampling at lower (often much lower) intervals. The advent of the age of technology has brought about a dramatic fall in the cost of gathering and storing data, as well as making the simultaneous transmission of information to physically dispersed viewers possible. In this sense, the very structure of markets has changed (Goodhart and O'Hara, 1997, p. 74), having important implications for both the availability of and the need to interpret high frequency data. In recent years, data sets containing data sampled at ever increasing frequency have become available; we now have what Engle (2000, p. 2) terms *ultra-high frequency data*, commonly known as *tick data*. Every single transaction is recorded; data sampled at higher frequency, even continuously, cannot provide more information about financial markets than what we have at our disposal at the moment. Chapter 2 is devoted to studying the stylized features of high-frequency transaction data; we use data from transactions in Financial Times Stock Exchange (FTSE) 100 index futures in 2002 on the London International Financial Futures and Options Exchange (LIFFE).

Applied research in high frequency financial data has a rather short history. According to Lin, Knight and Satchell (1999, p. 28), it may be dated from the *First International Conference on High Frequency Data in Finance*, on 29-31 March 1995, sponsored by Olsen and Associates; several papers presented at this conference subsequently appeared in the collection by Lequeux (1999). Perhaps most indicative of the speed at which interest in this area has developed is the fact that, as early as 1997, two issues of the *Journal of Empirical Finance* (Baillie and Dacorogna, 1997) were devoted to this topic; *inter alia* including a comprehensive review (Goodhart and O'Hara, 1997) of a host of literature, containing the availability of databases, statistical properties, problems and difficulties involved with high-frequency data.

Interest in the timing and frequency of transactions has been greatly influenced by the results of several studies indicating the importance of the number of trades as an informational proxy in the empirical study of stock returns. Indeed, Jones, Kaul and Lipson (1994, p. 631) categorically states that

> it is the occurrence of transactions per se, and not their size, that generates volatility; trade size has no information beyond that contained in the frequency of transactions.

In support of this conclusion, Lin et al. (1999, p. 56) also finds that the number of trades and the number of price changes seem to be the best choices for informational variables, volume being decidedly inferior. Dacorogna, Gençay, Müller, Olsen and Pictet (2001, p. 46) warns, and rightly so, against the indiscriminate use of quote data as an informational proxy, as quotes are prone to manipulation by market agents, and therefore quote data provides as much information about human behaviour as it does about financial markets. However, with the ready availability of actual trade data, this problem of mis-information is greatly alleviated.

Against this backdrop, Chapter 3 considers a model of the timing of transactions, using ideas developed by Rogers and Zane (1998). We propose that the level of activity in a financial market can be represented by a continuous-time Markov chain; occurrences of transactions on different assets are then modelled as a multidimensional Cox process with intensity dependent on the state of the Markov chain. We also briefly discuss two related models, developed by Engle and Russell (1998) and Rydberg and Shephard (1999).

Following the methodology of the potential approach of Rogers and Yousaf (2002) to modelling interest and exchange rates, we also suggest a (quasi-)Bayesian iterative approach to maximum likelihood estimation of the parameters of our model. We also illustrate the use of this method on transaction data on trading in FTSE 100 futures.

In Chapter 4, we briefly discuss our results. As could be expected from such a complex modelling framework, some issues remain unsolved, leaving potentially profitable future lines of research.

Chapter 2

Stylised features of transaction tick data

In this chapter, we investigate the stylised features of high frequency transactions data, using information on transactions—where transactions are loosely interpreted, to encompass asks and bids as well as actual trades—on the London International Finance Futures and Options Exchange (LIFFE) in FTSE 100 index futures during 2002. For the most part, we will follow same order as Rydberg and Shephard (2000) in their analysis of IBM shares traded on the NYSE in 1995.

In Section 2.1, we summarise the essential knowledge needed when studying high frequency transactions data, paying particular attention to the origin and nature of our tick data. With our data set in hand, we then consider, in turn, the four outstanding features of transaction data, namely the presence of strong seasonal patterns (Section 2.2), long memory (Section 2.3), clustering (Section 2.4) as well as the strong dependence between bids, asks and trades (Section 2.5). Finally, in Section 2.6, we illustrate how the microstructure of transactions depend on the trading method used by the market.

The aim of this chapter is not to present an exhaustive analysis of tick data: such an undertaking will certainly fill several volumes, if it is to be judged by the volume of the data alone. Instead, we will illustrate the outstanding characteristics of our data set with a broad brush, while attempting to explain them with the tools financial theory affords us.

2.1 Introduction to FTSE 100 index futures

The Financial Times Stock Exchange (FTSE) 100 index was launched in 1984. It tracks the share price movements of the 100 most highly capitalised blue chip companies on the London Stock Exchange, representing approximately 80% of the UK market (FTSE, 2003). It is calculated every 15 seconds throughout the day from 08.00 until the 16.30 closing auction (LIFFE, 2003*a*). The FTSE 100 index has long been a barometer by which professional money managers and private investors measure portfolio performance.

A LIFFE FTSE 100 index futures contract is an agreement to buy or sell a given quantity of the assets underlying the FTSE 100 index at a specified future date, at a pre-agreed price. At any one time, futures with delivery dates at the end of the four nearest quarterly months (March, June, September and December) are listed for trading. Most trades take place in the "front" month; trading in futures with this maturity date is allowed until 10:30:30 on the third Friday, or the last business day before the third Friday, in the delivery month.

Delivery of FTSE 100 futures takes place on the first business day after the expiry date. Rather than requiring physical delivery of the underlying basket of 100 shares, futures are cash settled at the Exchange Delivery Settlement Price (EDSP), which is based on the average values of the FTSE 100 index every 15 seconds from 10:10 to 10:30 on the last trading day (LIFFE, 2003*c*). Of the 81 measured values, the highest 12 and lowest 12 are discarded and the remaining 57 are averaged and rounded to the nearest half index point to calculate the EDSP.

As expected, activity in front month futures increases—exponentially—from the date of first listing until the last trading day. We will not consider the complex relations that exist between trading in futures with different delivery trades here. In fact, as we will see shortly, the total number of transactions does not exhibit any real quarterly seasonal pattern. This is most likely due to the tendency of investors to gradually transfer their positions from the existing front month to the next contract month as the last trading day of a contract month approaches (LIFFE, 2003*c*).

From its establishment in 1982 until November 1998, contracts on LIFFE traded predominantly via open outcry. Each product was traded in a designated area, called a pit, by traders who would register—often shout—bids and offers. Pit observers, wearing distinctive blue jackets, were tasked with observing that all behaviour in the pit conformed with the rules of the market. They also reported prices via microphones to price reporters, located off the trading floor, who would enter the prices, in real time, into terminals. The delay in prices being vocally transmitted from pit observer to price reporter and being entered into a terminal was typically less than a second (MacGregor, 1999, p. 308–309).

During late 1998 and 1999, LIFFE moved from open outcry to automated trading, introducing *LIFFE Connect*, "the world's most efficient, versatile and functionally rich electronic trading platform" (Hunter, 1998). Automated trading is a completely different trading method to open outcry: there is no trading floor, and the role of pit observers is no longer to report prices, but to observe orders as they are matched, and to ensure an orderly market place. Collection of tick data is fully automated, inexpensive and error-free; prices are automatically appended to tick data files generated directly from the trading floor.

Our data set contains, for every transaction, a time stamp, the delivery date of the future, an alphanumeric code indicating the type of transaction, the price (in index points) and the volume of the transaction. Transactions are grouped into the following seven types (the alphanumeric codes which will subsequently appear in tables and figures are given in brackets) (Turan, 2003; LIFFE, 2003*b*; WebFinance, 2003):

1. An *ask* (A) is the declaration by an agent of his/her willingness to sell a specified number of futures, accompanied by the lowest price that the agent is willing to accept.

2. A *bid* (B) is the declaration by an agent of his/her willingness to buy a specified number of futures, accompanied by the highest price that the agent is prepared to pay.

3. A *basis trade* (J) (or cash and carry trade) is to either buy cash bonds and sell futures, or to sell cash bonds and buy futures. Basis trading is usually a consequence of the perception that the futures and bonds are mispriced with respect to each other, and that the mispricing will correct itself such that the gain on one side of the trade will more than compensate for the loss on the other side.

4. A *block trade* (K) is a simultaneous transaction involving a large number of securities, typically at least 10 000. Normally only institutional investors undertake trades of this magnitude.

5. A *spread trade* (S) is the purchase of one futures contract and the simultaneous sale of another in order to take advantage of relative price changes, for instance buying one futures contract and selling another futures contract with a different delivery month.

6. A *trade* (T) is a transaction in the future.

7. A *volatility trade* (V) is the simultaneous buying or selling of an option and its related future in an options contract. Volatility trades enable agents to take advantage of the (implied) volatility of the underlying contract, rather than the direction of price movement.

In this study, we will also restrict our attention to the three most prevalent transaction types, as indicated in Table 2.1, namely asks (45.9%), bids (43.5%) and trades (10.5%). Despite the fact that FTSE 100 futures are normally traded between 8:00 and 17:30 (LIFFE, 2003*c*), we discard all transactions (ca. 6%) with timestamps outside London Stock Exchange opening hours (8:00–16:30). Entries with timestamps before 8:00 were spurious, and almost certainly not related to actual transactions; in addition, preliminary analysis suggested that entries were significantly different when they had a time stamp after 16:30 .

	Type of transaction as percentage of total						
Month	**A**	**B**	**J**	**K**	**S**	**T**	**V**
January	41.819	43.732	0.004	0.002	0.016	14.281	0.146
February	43.412	43.683	0.000	0.000	0.025	12.778	0.102
March	44.077	43.798	0.004	0.005	0.184	11.840	0.093
April	44.710	45.459	0.004	0.001	0.011	9.735	0.080
May	46.145	43.913	0.004	0.001	0.014	9.850	0.072
June	44.077	43.091	0.006	0.011	0.136	12.608	0.072
July	44.563	41.113	0.003	0.003	0.008	14.238	0.072
August	45.644	43.987	0.003	0.002	0.007	10.306	0.051
September	45.896	43.342	0.002	0.005	0.083	10.633	0.039
October	48.774	41.701	0.001	0.001	0.006	9.473	0.044
November	48.592	43.973	0.002	0.002	0.018	7.377	0.035
December	46.016	45.700	0.002	0.006	0.087	8.156	0.033
Total	45.933	43.460	0.003	0.003	0.044	10.497	0.060

Table 2.1: Spread of entries over different transaction types

2.2 Seasonal patterns

In 2002, FTSE 100 futures were traded on 250 days. A total of 27 463 782 transactions—on average, 111 189 per day—were observed during this period. Table 2.2 gives the distribution of transactions over the twelve months under consideration.

Month	**Trading days**	**Total transactions**	**Transactions per day**
January	22	1 193 031	54 228.68
February[a]	18	1 181 516	65 639.78
March[b]	20	1 332 410	66 620.50
April	20	1 551 334	77 566.70
May	20	1 538 428	76 921.40
June	19	1 772 620	93 295.79
July	23	2 764 430	120 192.61
August	21	2 857 475	136 070.24
September	21	3 269 109	155 671.86
October	23	3 903 941	169 736.57
November	21	3 249 108	154 719.43
December	20	2 850 380	142 519.00
Total	248	27 463 782	111 189.40

[a]Tick data for 27 and 28 February were unavailable (Intelligent Financial Services Limited, 2003); as a consequence these days were excluded.

[b]Although tick data for 1 March was indicated as unavailable by Intelligent Financial Services Limited (2003), the entries for this day were in fact present; they were utilised accordingly.

Table 2.2: Number of transactions and number of trading days

Studying the number of transactions on each day that the market was open, we observe very significant changes in the activity level during the year. Despite the fact that LIFFE reported the highest ever number of total transactions for the month of January in its history in the first month of 2002 (LIFFE Press Office, 2002), we find the lowest levels at the beginning of the year. In contrast, market activity reaches its highest levels in September, October and November. This pattern corresponds to the findings of Rydberg and Shephard (2000, pp. 223–224), and is almost certainly seasonal.

To the naked eye, Figure 2.1 suggests that the daily number of transactions should increase over time. Fitting a Poisson generalized linear model with the natural logarithmic link function, we find that (the natural logarithm of) the number of transactions per day increases significantly over time.

However, the estimates for this gradient are very small (0.0036, 0.0033 and 0.0020, for asks, bids and trades, respectively; time is measured in days). Consequently, we have to conclude that the long-term trends in the number of transactions are negligible, at least on the timescale that we are interested in.

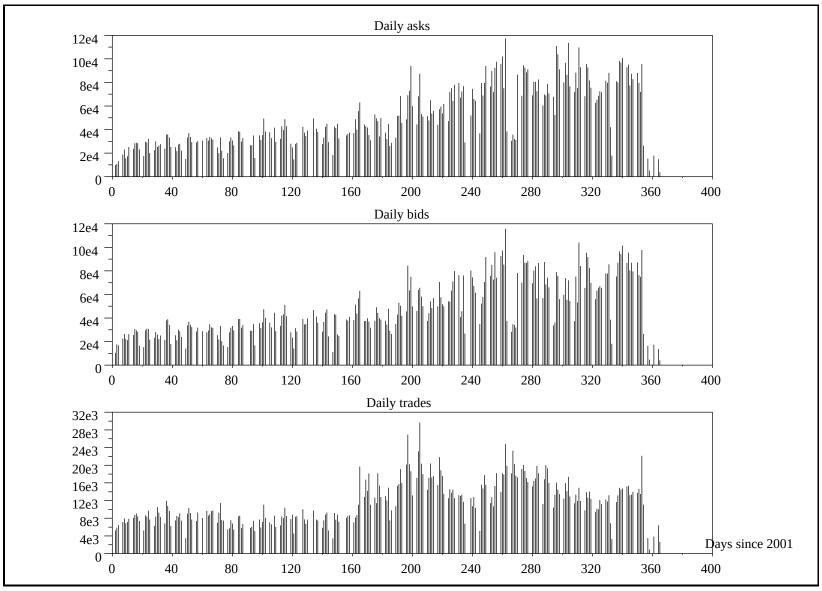

Figure 2.1: Daily transactions in FTSE 100 futures

Together, Table 2.1 and Figure 2.1 suggest the existence of strong interdependence between the daily number of bids, asks and trades. This is indeed the case: the correlation between the numbers of asks and bids is 94.46%, while the correlation between these two variables and the number of trades is 68.62% and 67.86%, respectively.

Table 2.3 illustrates another interesting feature of the data, namely its intra-weekly patterns. On average, the market is at its slowest on Mondays; according to ap Gwylim, Buckle and Thomas (1999, p. 168), this corresponds with significant lower volatility at the beginning of the week. Activity reaches a maximum in the middle of the week, after which it slows down again. As was the case with the long-term trend, we find that the daily number of transactions differ significantly between days of the week.

	Number	**Average number of transactions**		
Day	**recorded**	**Asks**	**Bids**	**Trades**
Monday	47	45 791.340	42 320.277	10 225.511
Tuesday	52	50 547.923	49 128.558	11 884.981
Wednesday	48	54 242.646	51 348.167	12 478.646
Thursday	50	55 539.220	52 918.280	12 628.520
Friday	51	48 112.804	44 732.961	10 858.882
Total	248	50 867.121	48 128.004	11 624.282

Table 2.3: Number of transactions in FTSE 100 index futures by day of week

Close inspection of Figure 2.1 reveals the presence of a cyclical time-of-the-month effect: it seems that activity is higher in die middle than at the beginning or end of a calendar month. Indeed, Figure 2.2, the correlogram of daily transactions (deseasonalised according to Table 2.3) exhibits noticeably higher correlation at lags 20 (one month) and 40 (two months). It is also interesting to note that this effect is most pronounced for bids, while it is almost invisible for trades.

We now consider the data at an intraday scale. Since time stamps are precise to one second, we group trades with the same time stamp and code together. The result is a sequence $\{N_k\}_{k\in\mathbb{N}}$ of binned counts; the number N_n is the number of transactions in the n^{th} second.

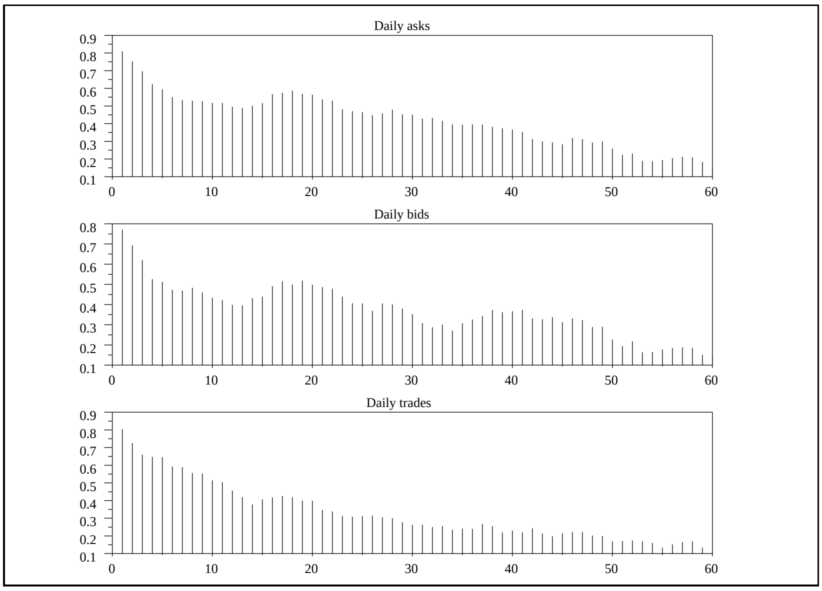

Figure 2.2: Autocorrelation of deseasonalised daily transaction counts of FTSE 100 index futures

We construct an estimate of the average number of transactions per second at every time of the day, for each day of the week, by first calculating the average number of transactions in each second, exponentially smoothing it and then fitting a natural cubic spline. Figures 2.3–2.5 are highly intriguing.

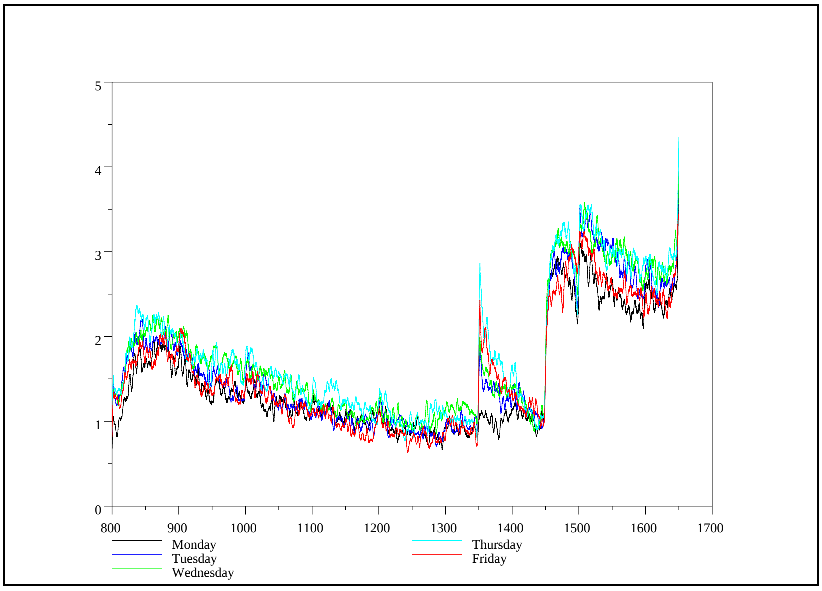

Figure 2.3: Estimated daily seasonal pattern of asks for FTSE 100 index futures

We find that business ambles along fairly evenly during the morning hours, slows down towards lunchtime and then picks up again; trading reaches its highest levels in the afternoon. This is in stark contrast to results on American markets pointing towards an inverse J-shaped pattern (cf. Goodhart and O'Hara (1997, p. 86), Rydberg and Shephard (2000, Figure 4)).

Figures 2.3–2.5 have a number of sudden jumps in common. The opening and closing of trade on the London Stock Exchange is most surely responsible for the rounded peaks at 8:30, and the dramatic increase in activity just before 16:30. Likewise, the near-vertical climb in the activity rate at 14:30 can be ascribed to the opening of American markets (Reuters, 2001), as can the quick gear-change around

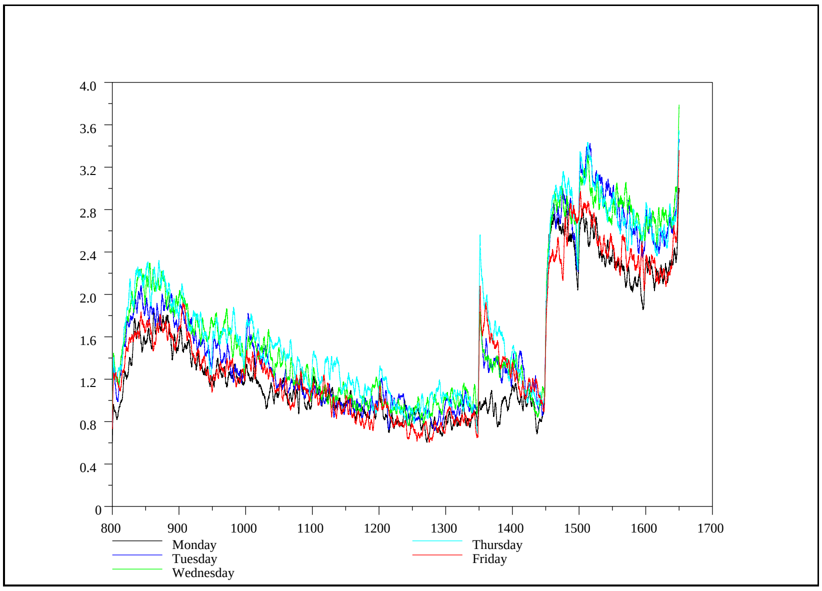

Figure 2.4: Estimated daily seasonal pattern of bids on FTSE 100 index futures

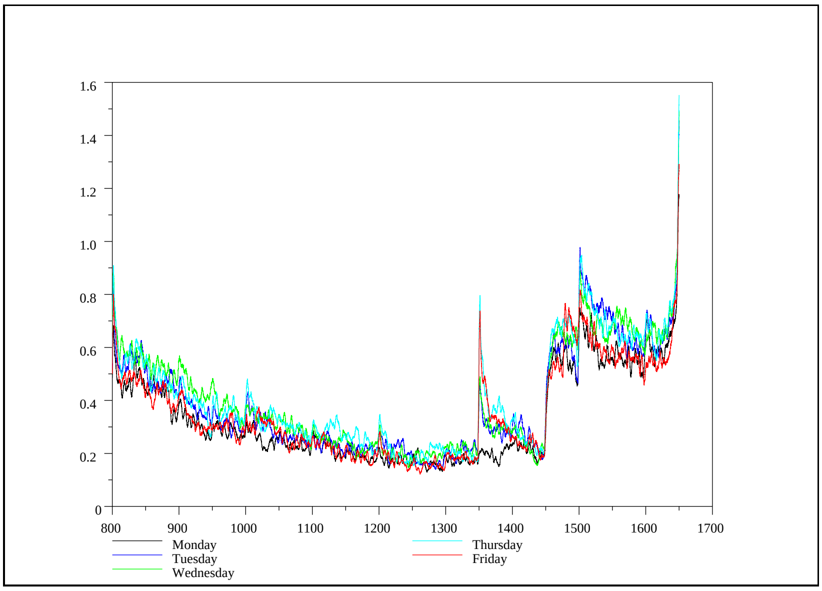

Figure 2.5: Estimated daily seasonal pattern of trades in FTSE 100 index futures

15:00. Except for the level, the pattern of activity between 15:00 and 16:00 compares favourably with the period of the same length, starting at 8:30.

The very pronounced peak—followed by a renormalisation of activity to the same level as over lunch—at 13:30, on every day of the week except on a Monday, coincides with the timing of U.S. macroeconomic announcements (Kim, 2003, Figure 1). The effects of these announcements on FTSE 100 index futures contracts extends, in both presence and form, to volume, price and volatility; it is well-documented by, among others, ap Gwylim et al. (1999, pp. 159–173) and Clare and Courtenay (2000, p. 269). In particular, comparison of Figures 2.3–2.5 to recent work done by Andersen, Bollerslev, Diebold and Vega (2002, Figure 2) on intraday volatility patterns is quite illuminating.

2.3 Serial dependence in transaction counts

We return, momentarily, to the daily transaction counts studied in Section 2.2. The slow decay in the autocorrelation function, illustrated in Figure 2.2, points to the presence of very long memory in the level of market activity. In fact, all the autocorrelations shown (up to 60 lags, or three months) are statistically different from zero. However, due to the fact that we have not taken any long-term seasonal effects into account (ideally, one would need several years worth of data to estimate it), this autocorrelation pattern may simply be a consequence of the presence of a short-term economic cycle. Most likely, as Rydberg and Shephard (2000, p. 226) also concludes, this autocorrelation pattern is due to a combination of long memory and economical trend.

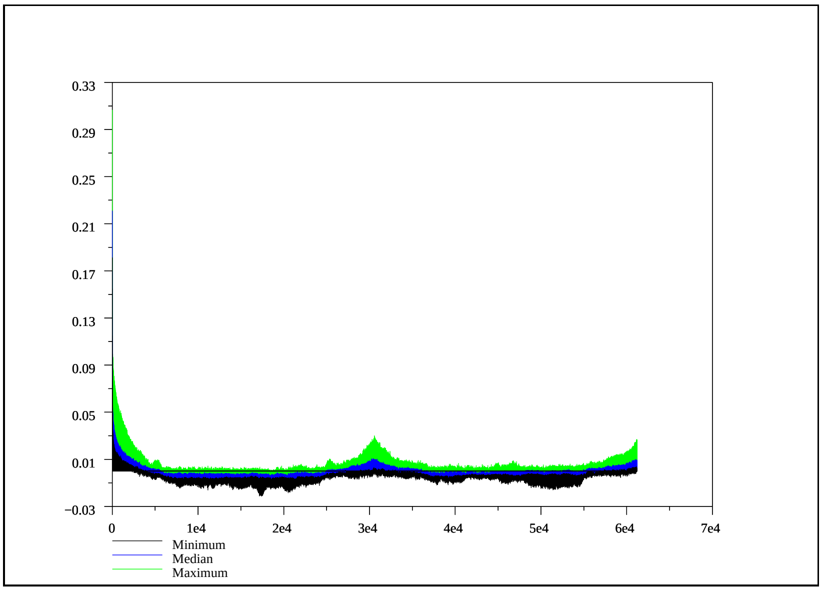

Figure 2.6: Minimum, median and maximum of autocorrelation of monthly second-by-second counts of asks for FTSE 100 index futures

We now turn to the dynamics of the time series $\{N_k\}_{k\in\mathbb{N}}$ of second-by-second bin counts. For each transaction type, we first expand the entries for each calendar month into a sequence consisting of 30 600 counts per working day; the result is a time series of length between 550 800 (February) and 703 800 (July and October). Having subtracted our estimates of the daily seasonal pattern (cf. Figures 2.3–2.5), we then calculate the autocorrelation function of each series up to 61 200 lags, representing two days.

The striking similarity between the resulting autocorrelation functions justifies our simplifying approach. Rather than reproducing the 36 correlograms (12 for each of asks, bids and trades), we display the minimum, median and maximum autocorrelation for each lag over the 12 months in Figures 2.6–2.8.

As was the case with the correlogram of daily transaction counts, we find that autocorrelations die out extremely slowly. In the short run, autocorrelations are positive; they are mostly irrelevant after

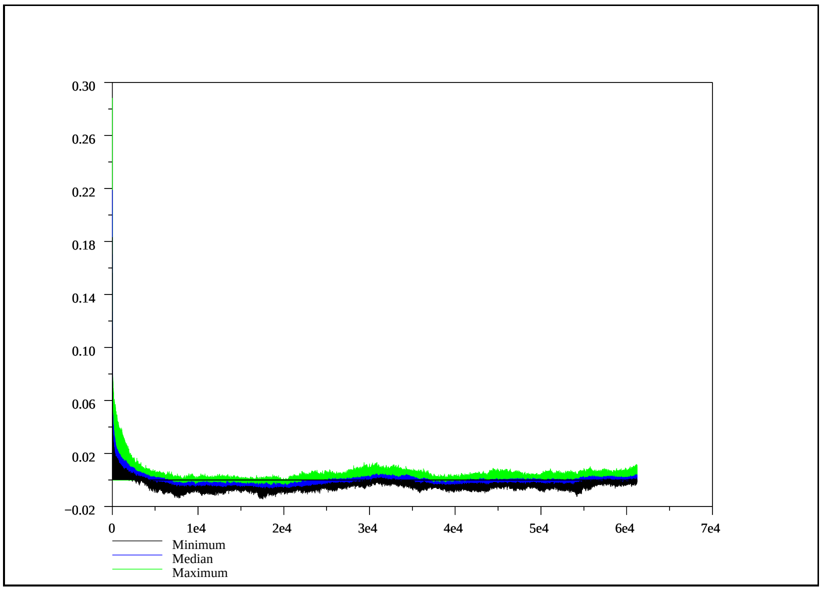

Figure 2.7: Minimum, median and maximum of autocorrelation of monthly second-by-second counts of bids on FTSE 100 index futures

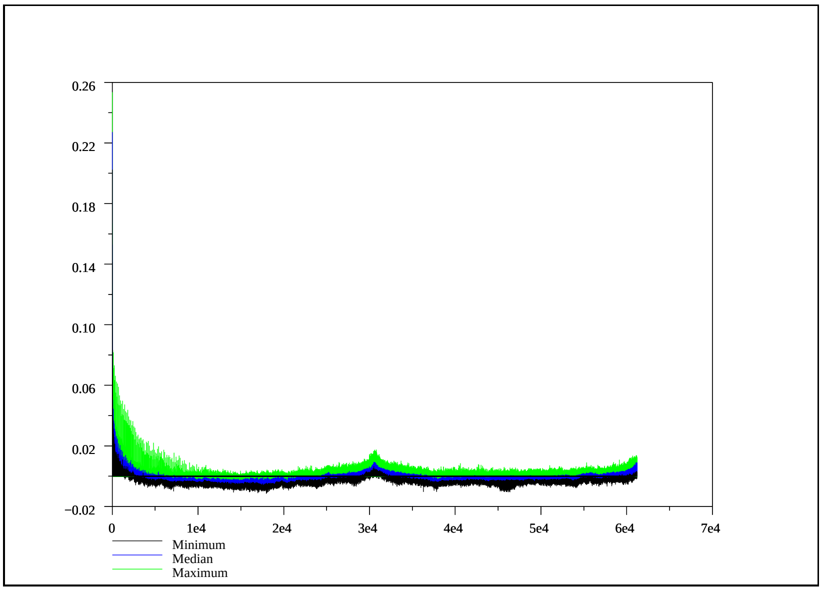

Figure 2.8: Minimum, median and maximum of autocorrelation of monthly second-by-second counts of trades in FTSE 100 index futures

5 000 lags (ca. 80 minutes). The only exceptions are local maxima around lags 30 600 (one day) and 61 200 (two days); we concur with Rydberg and Shephard (2000, p. 226) that this phenomenon should be ascribed to the neglected seasonal pattern.

2.4 Clustering of transactions

Having studied the gross features of the temporal dependence structure of transactions, we now turn to the short-term behaviour of transaction counts. Figure 2.9, which was constructed in the same way as Figures 2.6–2.8, features high autocorrelation at short lags; this decreases very rapidly, pointing to longer runs of high transaction counts than would be the case if the times of transactions were uniformly distributed over the course of the day. The only exception to this is the singular spike in the autocorrelation between trade counts at lag 60; the only plausible explanation for this is that, while asks and bids are submitted at any time, trades tend to be executed at round minutes.

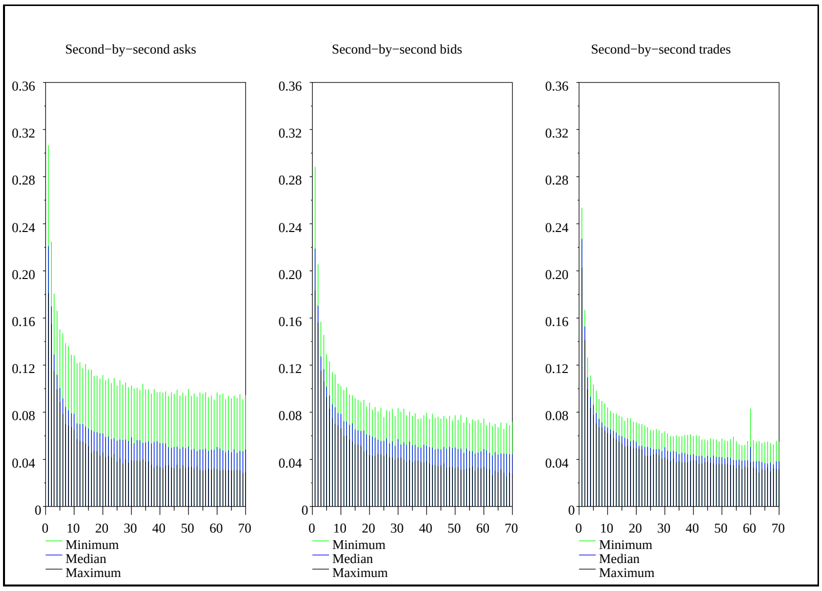

Figure 2.9: Short-term minimum, median and maximum autocorrelation of transactions in FTSE 100 index futures

This clustering, which is a common feature of transaction data, can be explained by the economics of market microstructure. According to Engle and Russell (1998, p. 1129), such theories often partition traders into two types, namely informed traders, who are assumed to possess information not publicly available, and liquidity traders, or those not having access to the same superior information set as informed traders. It is commonly assumed that liquidity traders arrive independently, according to a Poisson distribution. By contrast, informed traders will enter the market only after observing a private, potentially noisy signal.

Engle and Russell (1998, p. 1148) further explains that, in a rational expectations setting, monitoring the flow of orders allows specialists to detect such private information, and enter the market as informed traders. While prices change according to the spread of this superior information, informed traders will seek to trade as long as their information has value. As a consequence, we observe clustering of trading following an information event because of the changing numbers of informed traders.

This phenomenon is nowhere as evident as in the bursts of trades following U.S. macroeconomic announcements and the opening of U.S. markets, which we considered in Section 2.2. Whenever new relevant information becomes available, activity on the market increases; this persists for some time before returning to the previous background level.

2.5 Dependence of trades on bids and asks

Intuitively, we expect to observe strong interdependence between bids, asks and trades. In the normal course of business, asks and bids are expected to evolve in a similar way, with trades following closely. This relationship, however, is far from clear-cut; bids and asks are often submitted for various reasons other than the intention to trade. In comparison to placing bids and offers, agents should be expected to be much more cautious when trading: FTSE 100 index futures, for instance, trade at multiples of GBP 10, so a typical transaction would involve tens of thousands of pounds.

Having no means to identify trades with their constituent bids and asks, we are restricted to studying the covariation between trades and, respectively, bids and asks; an covariation at lag l of a sequence $(y_k)_{k=1}^m$ on a sequence $(y_k)_{k=1}^m$ is defined as

$$\rho_l := \frac{1}{m} \sum_{k=1}^{m-l} (x_k - \bar{x})(y_{k+l} - \bar{y}) .$$

Calculating these quantities in the same way as was done in the previous two sections (by dividing and conquering), we obtain Figure 2.10. In a sense, the smoothness of the curves are slightly disappointing. Evidently, more detailed analysis is needed to obtain sensible results about interdependence between different types of transactions on the same assets.

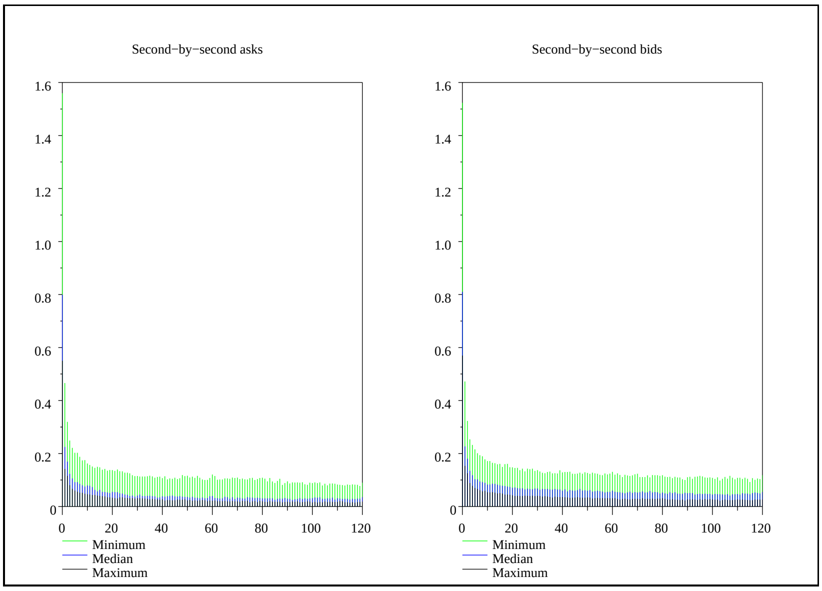

Figure 2.10: Covariation between trade counts and ask and bid counts at different lags

2.6 Trading methods

In their study of trades in IBM stocks on the New York Stock Exchange in 1995, Rydberg and Shephard (2000, pp. 224–225) found the correlogram to be completely dominated by a negative autocorrelation at lag 1. The autocorrelation function then increases from lag 2, to reach its maximum at 6; at higher lags their empirical autocorrelation function agrees with ours. The difference at microstructure level is due to the difference in trading methods: in 1995, the NYSE—and LIFFE—still operated on an open outcry basis. The negative first order autocorrelation—and subsequent increase—found by Rydberg and Shephard is most certainly due to the inability of the market maker to record trades quickly enough at active times of the market.

In contrast, the correlograms in Figure 2.9 decrease smoothly, with no evidence of reporting delay. Perhaps this single fact—that the time it takes to provide such straightforward information as time and price, has decreased from *several seconds* to *several milliseconds* in the space of only eight years—epitomises the dilemma we face when attempting to understand high frequency finance. As a consequence of our effort in developing technology to obtain faster, more accurate results, the very object of our study is becoming more and more sophisticated. As financial expertise becomes more prevalent, the single brilliant idea that may make an individual a millionaire becomes more and more elusive. Complete understanding of financial markets may forever lie just beyond the horison.

Chapter 3

A Markov-Cox model for transaction times

In this chapter, we develop a modelling and estimation framework for intraday binned transaction counts. In Section 3.1, we consider simple case of trades in a single asset, to ease the exposition of our modelling assumptions. Section 3.2 considers the general situation, where a number of assets are traded on the same market. We develop an iterative maximum likelihood estimation procedure in Section 3.3; in Section 3.4, we illustrate this method by applying it to binned counts of transactions in FTSE 100 index futures.

3.1 Trades in a single asset

Let us first consider the simplest case, namely that of transactions—trades, say—in a single asset. The problem of modelling times and frequencies of trades is at heart an issue of describing a countable random point process $\{t_k\}_{k\in\mathbb{N}}$ with very distinctive features. Generally speaking, the modelling of such a process an be approached in two different ways, namely by either modelling the occurrences of trades themselves, or modelling the durations between trades.

A prime example of the latter approach to modelling transaction processes is the Autoregressive Conditional Duration (ACD) model, developed by Engle and Russell (1998). This model, which has proved quite popular in the econometrics community, suggests that the expected duration between transactions follow a process of the type

$$\psi_i = \omega + \alpha x_{i-1} + \beta\psi_{i-1} \text{ for some } \alpha \geq 0 \text{ and } \beta \geq 0,$$

where $x_i := t_i - t_{i-1}$ denotes the previous observed duration. Note that the durations are the inverse of the intensity of the Poisson process. In some sense, the process is discrete, and the intensity with which jumps in prices occur can only change at the discrete time points. This model is based on the assumption thatthe standardized durations $\frac{x_i}{\psi_i}$ are independent and identically distributed with density function

$$g\left(\left.\frac{x_i}{\psi_i}\right| x_1, x_2, \ldots, x_{i-1}; \theta\right) \equiv g\left(\frac{x_i}{\psi_i}; \theta\right)$$

and expected value

$$\mathbb{E}\left(\left. x_i \right| x_1, x_2, \ldots, x_{i-1}; \theta\right) = \psi_i.$$

In our work, we take the first approach; we will model the times of trades as a point process. We begin by association the times of trades with a trade counting process. Many alternative definitions of counting processes exist (cf. Snyder (1975, p. 7), Engle and Russell (1998, p. 1129)); we follow Grandell (1997, pp. 51–52) in calling a process a *counting process* if it has state space $\mathbb{N}_0 \cup \{+\infty\}$ and non-decreasing right-continuous (hence càdlàg) paths.

The natural counting process $\{N_t\}_{t\geq 0}$ associated with a point process $T := \{t_k\}_{k\in\mathbb{N}}$ is defined by setting N_t to be the number of point in the set $T \cap [0,t]$. In the sequel, we will conveniently abuse notation by identifying a point process with its associated counting process, and vice versa.

In modelling trades, we intend to follow Rogers and Zane (1998, p. 2), Lin et al. (1999, pp. 45–46) and Rydberg and Shephard (2000, p. 219) and others in assuming that trades form a Poisson process. It is therefore imperative that we convince ourselves of the validity of such an assumption before proceeding to the development of our model. The following theorem (Snyder, 1975, Theorem 2.2.2) is central to the motivation of our modelling assumptions; we present it without proof.

Theorem 3.1 *Let $\{N_t\}_{t\geq 0}$ be the counting process associated with a countable point process $\{t_k\}_{k\in\mathbb{N}}$. Suppose that the following hold true:*

1. *$N_0 = 0$ almost surely.*
2. *The point process is* conditionally orderly*, i.e. at any time $u > 0$, there exists for any $\varepsilon > 0$ a $\delta > 0$ such that, for all $\Delta t \in (0, \delta)$,*
$$\mathbb{P}\left(N_{u+\Delta t} - N_u > 1 \,\middle|\, \{N_k\}_{k\in[0,u)}\right) \leq \varepsilon \mathbb{P}\left(N_{u+\Delta t} - N_u = 1 \,\middle|\, \{N_k\}_{k\in[0,u)}\right).$$
3. *For all $t > 0$,*
$$\lambda(t) \equiv \lambda\left(t \,|N_t, t_1, t_2, \ldots, t_{N_t}\right) := \lim_{\Delta t\to 0} \frac{\mathbb{P}\left(N_{t+\Delta t} - N_t = 1 \,|N_t, t_1, t_2, \ldots, t_{N_t}\right)}{\Delta t} \tag{3.1}$$
exists and is a finite, integrable function of t alone. In addition, if
$$\Lambda_t := \int_0^t \lambda(u)\, du \text{ for all } t \geq 0, \tag{3.2}$$
then $\Lambda_v - \Lambda_u < \infty$ for all finite u and $v > u$.

Then $\{N_t\}_{t\geq 0}$ is a Poisson counting process *with intensity λ (or mean measure induced by Λ), i.e. for any $u < v$ its increment $N_v - N_u$ is independent of $\{N_t\}_{t\in[0,u]}$, and has Poisson distribution with parameter $\Lambda_v - \Lambda_u = \int_u^v \lambda(t)dt$.*

Proof. Snyder (1975, pp. 51–54) ∎

If the three conditions above are satisfied, Theorem 3.1 guarantees that trades follow a Poisson process. The first two are easily dealt with; in short, requirement 2 means that, at any time, conditional on events before that time, the probability of two or more events occurring is infinitesimal relative to the probability of one event occurring in any sufficiently short interval of time. A moment's reflection convinces us that such an assumption is indeed reasonable. Requirement 3, however, poses somewhat of a challenge.

Requirement 3 implies not only that the limit (3.1) exists, but also that $\{t_k\}_{k\in\mathbb{N}}$ evolves without after-effects, at least infitesimally; a point process $T := \{t_k\}_{k\in\mathbb{N}}$ on $[0,\infty)$ is said to evolve *without after-effects* if, for any $u > 0$, the realization $T \cap [u,\infty)$ of events after u does not depend in any way on the sequence $T \cap [0,u]$ of events that have transpired in $[0,u]$ (Snyder, 1975, p. 42). This property expresses the independence of the past and future of the point process, at least in the short term.

Simply assuming that requirement 3 is true for the trades process directly contradicts our conclusions in Chapter 2 about the presence of clustering and long memory in the series. In addition, even if we could assume that trades evolve without after-effects, we face yet another challenge in using the Poisson modelling framework: the intensity. To illustrate this, we assume momentarily that requirement 3 above did indeed hold true. According to Theorem 3.1, trades follow a Poisson process; the only unclarified issue remaining is the form of the intensity (3.1).

We need only refer to Figure 2.5 to convince ourselves that λ cannot possibly be constant over time. Given that the mean intensity of trades (which is the maximum likelihood estimator of λ) varies so widely during the course of a day, it is inconceivable that trades follow a homogeneous Poisson process.

Allowing λ to be a deterministic funcion of time certainly provides us with a more flexible modelling framework. However, there are a number of reasons why an inhomogeneous Poisson process might fails do justice to the trade process:

1. The observed trade intensity varies between days. Figure 3.1 depicts second-by-second counts of trades in FTSE 100 futures on five Tuesdays in October 2002. Although all five days display the hallmark U-shape, it is clear that the intensity of trading differs between the five days; compare, for instance, the difference in activity in the period 13:00–14:00 between days.

2. Trades in different assets are often correlated, as are the changes in their intensities. If, in extending this model to multiple assets, we assume that the times of trades of each asset form a Poisson process with deterministic intensity, we are bound to encounter difficulties in explaining interdependence between different assets. The reason for this is that the intensities of the different assets are deterministic, and hence independent of each other.

3. A deterministic intensity function does not provide for clustering of transactions; this is even more the case with clustering due to unforeseen, unscheduled events.

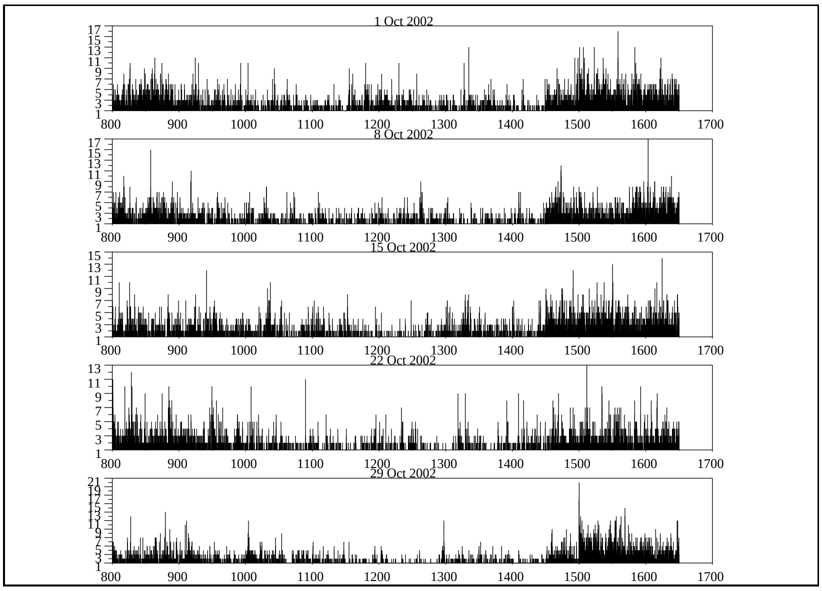

Figure 3.1: Second-by-second frequencies of trades in FTSE 100 (delivering December 2002) on five consecutive Tuesdays in October 2002

Thus, even if we could accept a modelling framework that does not allow for dependence between trade counts in disjoint intervals, determining the intensity of the counting process of trades is far from clear. Fortunately, we can remedy both these problems by allowing the intensity itself to be stochastic; conditional on the intensity, trades can be modelled as an inhomogeneous Poisson process. In this way, after-effects and dependencies are modelled through the intensity process, while, at the same time, we enjoy the benefit of working with a well-documented, well-behaved point process.

Such a model is easily expressed by means of a random measure. A *random measure* Λ on $[0, \infty)$ is a stochastic process such that $\Lambda_0 = 0$ almost surely, with non-decreasing right-continuous (hence càdlàg) paths satisfying $\Lambda_t < \infty$ for all $t < \infty$ (Grandell, 1997, p. 84). Note that a counting process, as defined above, is simply an integer-valued random measure.

In the same spirit as (3.2), the random measure Λ is linked to the intensity process λ via the integral equation

$$\Lambda_t = \int_0^t \lambda(u)\, du \text{ for all } t \geq 0. \tag{3.3}$$

Clearly, $\lambda \geq 0$ almost surely; in our work we will take λ to be strictly nonnegative. We also assume that the paths of λ are càdlàg; it follows from (3.3) that the paths of Λ are continuous.

Formally, we model the times of trades as a Cox (or doubly stochastic Poisson) process. A point process $T := \{t_k\}_{k\in\mathbb{N}}$ is called a *Cox process* with mean measure induced by Λ—or intensity process λ—if, conditional on Λ, T is a Poisson process with mean measure Λ (Kingman, 1993, p. 65). We defer the development of a model for the intensity process to the next section; our model allows the intensity to be governed by both intradaily seasonality and a latent independent Markov chain.

The remainder of this section is devoted to a related model for the process $\{N_k\}_{k\in\mathbb{N}}$ of binned trades, i.e. counts of trades in a sequence of time intervals of length Δt, not unlike the sequences of transaction counts considered in Chapter 2. In their BIN model, Rydberg and Shephard (1999, p. 9) assume that the Poisson parameter λ_n of the n^{th} binned count N_n follows a moving average process of the form

$$\lambda_n = \alpha + \sum_{k=1}^{p} \gamma_k N_{n-k} \text{ for some } \alpha > 0 \text{ and } \gamma_1, \gamma_2, \ldots, \gamma_p > 0.$$

This approach is intuitively appealing; for instance, at times of heightened activity, one would naturally want a model to predict that the number of transactions in the next bin would be high.

Rydberg and Shephard (2000, pp. 228–229) provides a simple yet striking example of how having a nondeterministic intensity function assists us in modelling the dependence structure of trades. Indeed, let Λ be the Cox mean measure. Conditional on $\Lambda_v - \Lambda_u$, we know that the count $N_v - N_u$ is a Poisson random variable with mean $\Lambda_v - \Lambda_u$; if

$$\lambda_n := \Lambda_{(n+1)\Delta t} - \Lambda_{n\Delta t} \text{ for } n \in \mathbb{N},$$

then, conditional on λ_n, the binned count N_n has a Poisson distribution with parameter λ_n. Moreover, since, conditionally on Λ, the Cox process has independent increments, it follows that the binned counts are conditionally independent.

We now calculate the unconditional correlation of binned counts. For simplicity, assume that λ is covariance stationary. Let μ, σ^2 and r denote, respectively, the mean, variance and autocorrelation function of the process λ. We also define

$$\rho(t) := \int_0^t r(u)\, du \text{ and } R(t) := \int_0^t \rho(u)\, du.$$

Assuming that λ_t is square integrable, we have

$$\operatorname{var}(\lambda_t) = \sigma^2 \int_0^t \int_0^t r(u-v)\, du\, dv = 2\sigma^2 R(u).$$

Consequently,

$$\operatorname{cov}(\lambda_n, \lambda_{n+s}) = \sigma^2 \Diamond R_s,$$

where

$$\Diamond R_s := R((s+1)\Delta t) - 2R(s\Delta t) + R((s-1)\Delta t).$$

The moments of N_n follow immediately; we have

$$\mathbb{E}(N_n) = \mathbb{E}(\lambda_n) = \mu \Delta t$$

and

$$\operatorname{var}(N_n) = \operatorname{var}(\lambda_n) + \mathbb{E}(\lambda_n) = 2\sigma^2 R(\Delta t) + \mu\Delta t.$$

Moreover,

$$\begin{aligned}\operatorname{cov}(N_n, N_{n+s}) &= \mathbb{E}(N_n N_{n+s}) - (\mathbb{E}(N_n))^2 \\ &= \operatorname{cov}(\lambda_n, \lambda_{n+s}) \\ &= \sigma^2 \Diamond R_s;\end{aligned}$$

it follows that

$$\operatorname{cor}(N_n, N_{n+s}) = \frac{\sigma^2 \Diamond R_s}{2\sigma^2 R(\Delta t) + \mu \Delta t}.$$

3.2 Transactions in multiple assets

Following Rogers and Zane (1998), we construct the Cox mean measure by using a Markov chain representing the level of activity of the market. On the one hand, this specific choice of underlying process is motivated by our observation that, whenever the intraday intensity changes, it remains at the new level for some time. For instance, studying Figure 3.1, it seems reasonable to suppose that activity is at one level in the morning, another in the middle of the day and yet another at the end of the day.

Another motivation for using Markov chains is that they are inherently simple, and their parameters are easy to interpret. For instance, for a Markov chain with state space $I := \{0,1\}^N$, an element $(x_1, x_2, \ldots, x_N) \in I$ can be thought of as a listing of the states of N shares, where we interpret $x_i = 1$ to mean that trading in share i is very active, and $x_i = 0$ to mean that share i is quiet.

We proceed in the following way. Let $X := \{X_t\}_{t \geq 0}$ be a continuous-time Markov chain with finite state space $\{1, 2, \ldots, N\}$. For simplicity of exposition, we assume X to be stationary and ergodic. Let Q be the matrix of transition rates, and let

$$P(\Delta t) := e^{-\Delta t Q} \text{ for } \Delta t > 0 \tag{3.4}$$

be the semigroup of transition matrices.

We have M different transactions processes in a market, labelled $1, 2, \ldots, M$; this may be thought of as trades in different assets, or even different types of transactions (asks and bids, for instance) in the same asset. Suppose that the intensity of transactions in each asset $l = 1, 2, \ldots, M$ can be expressed in the form

$$\lambda^l_{k,t} := f^l_k \, s^l_t \text{ for } k = 1, 2, \ldots, N, \tag{3.5}$$

where f^l_k is the mean intensity of transactions in asset l when X is in state k, and s^l_t is the seasonal index of asset l at time t. Combining (3.5) with (3.3), we have

$$\Lambda^l_t \equiv \Lambda^l_t(X) = \int_0^t f^l_{X_u} s^l_u \, du \text{ for all } t \geq 0. \tag{3.6}$$

Following Rogers and Zane (1998, p. 4), we suppose that these transactions are generated, in a way that shall soon be made clear, by standard Poisson counting processes $\tilde{N}^1, \tilde{N}^2, \ldots, \tilde{N}^M$, independent of each other and of X; a standard Poisson process is a simply a homogeneous Poisson process with unit intensity. That the assumption of independence among the Poisson processes is reasonable follows, on the one hand, from the martingale characterisation of Poisson processes (Meyer, 1976, pp. 288–289), and on the other, from the fact that independent Poisson processes are almost surely disjoint (Kingman, 1993, Disjointness Lemma). The point is that, as long as our different transaction processes do not intersect (i.e. trades in different assets taking place at the *exact same moment*), it is safe to assume that $\tilde{N}^1, \tilde{N}^2, \ldots, \tilde{N}^M$ are independent.

For each asset $l = 1, 2, \ldots, M$, we now define the counting process N^l by

$$N^l_t := \tilde{N}^l \left(\Lambda^l_t(X) \right) = \tilde{N}^l \left(\int_0^t f^l_{X_u} s^l_u \, du \right).$$

This is a Cox process with mean measure induced by Λ^l (Kingman (1993, Mapping Theorem), Cox and Miller (1965, pp. 153–154)). By this construction, we now have M Cox processes, whose interdependence is explained completely by the Markov chain X; conditional on X, the Poisson processes $N^1, N^2, \ldots, N^M$ are mutually independent, each with deterministic mean measure $\Lambda^l(X)$.

3.3 Estimation methodology

Having calculated the seasonal indices directly from the data, our model has $N(M+N-1)$ parameters, consisting of the $N(N-1)$ off-diagonal entries of Q, as well as N Poisson intensities for each of the M assets. Given that, at each time point, we have M observables, it is clear that the number of parameters exceeds the number of data points by far.

Conventional statistical wisdom would have a number of objections to a modelling framework such as this: in the first instance, due to the myth that, if the number of parameters exceeds the number of data points, the model must necessarile be degenerate. Rogers and Yousaf (2002, p. 380) provides convincing evidence to the contrary; similarly, the fit we obtain is by no means perfect. In considering this model, our main aim is to fit its strong structural properties; the number of parameters is a by-product of this objective, rather than an end in itself.

Secondly, some of the parameters will be indeterminate; this is to be expected, as our estimation procedure maximizes a real-valued function of many variables, and there is no reason to believe that this maximum should be unique. Experience with generalised method of moments estimation based on Laplace transforms of the history of transactions, suggested by Rogers and Zane (1998), has taught us that care is needed in ensuring that our estimation method is as proficient as possible in distinguishing between rival models of similar (but different) merit. In addition, it is well known that Markov chains do not have unique Q-matrices; interchanging rows (and corresponding columns) results in the same Markov chain, albeit with altered state labels.

Finally, we observe that the estimates of many of the parameters will be subject to large error; this is simply the consequence of some parameters having little or no influence on the model values for the observables. For instance, at a time when we believe strongly that the Markov chain X is in state i, say, we should accept that the Poisson intensities associated with state $j \neq i$ cannot be estimated with high precision. Rogers and Yousaf (2002, pp. 380–381) also points out that this type of situation—where the parameter space is very large in relation to the sample space—is quite common in the finance industry, and that resulting problems with parameter stability have proven very hard to deal with in practice. We take note of these objections and warnings; however, our methods will be justified by the quality of the fit that they achieve, and the stability of the estimates they produce.

Our model is parametrised by a vector θ; as usual, we regard θ as a stacked vector consisting of the $N(N-1)$ off-diagonal entries of Q, together with the MN Poisson intensities

$$f_1^1, f_2^1, \ldots, f_N^1, \ldots, f_1^M, f_2^M, \ldots, f_N^M.$$

We will be very lenient in the range Ω of values of θ; our only absolute requirement is that Ω is contained in $[0,\infty)^{N(M+N-1)}$.

As necessitated by the format of our data, we can only observe the Cox processes $N^1, N^2, \ldots, N^M$ periodically; we fix this period $\Delta t > 0$, and let $t_k := k\Delta t$ for $k \in \mathbb{N}_0$. We also let n_k^l be the observed number of transactions in the l^{th} asset in the interval $(t_{k-1}, t_k]$. For the sake of brevity, we write

$$\mathbf{n}_k^l \equiv \left(n_1^l, n_2^l, \ldots, n_k^l\right), \mathbf{n}_k \equiv \left(\mathbf{n}_k^1, \mathbf{n}_k^2, \ldots, \mathbf{n}_k^M\right) \text{ and } n_k \equiv \left(n_k^1, n_k^2, \ldots, n_k^M\right)$$

for the sequences of observed trades,

$$s_{t_k}^l \equiv s_k^l, \lambda_{l,k}^m \equiv f_l^m s_k^m \text{ and } \lambda_{l,k} \equiv \left(\lambda_{l,k}^1, \lambda_{l,k}^2, \ldots, \lambda_{l,k}^M\right)$$

for the Poisson intensities (cf. (3.5)), as well as

$$X_k \equiv X_{t_k} \text{ and } \mathbf{X}_k \equiv (X_0, X_1, \ldots, X_k)$$

for the Markov process.

We adopt a Bayesian standpoint. Suppose that $\mu_0 : \{1, 2, \ldots, N\} \to [0,1]$ is the initial law of X, and the initial distribution of θ is given by the density $\phi_0 : \mathbb{R}^{N(M+N-1)} \to [0,\infty)$. The likelihood of $(\mathbf{X}_k, \mathbf{n}_k, \theta)$ is

$$H_k\left(\mathbf{X}_k, \mathbf{n}_k, \theta\right) \propto \phi_0(\theta)\mu_0\left(X_0\right) \prod_{j=1}^{k} p_{X_{j-1},X_j}\left(\Delta t \,|\theta\right) \pi\left(n_j \,\middle|\lambda_{X_j,j}\Delta t\right), \tag{3.7}$$

where $\{p_{i,j}(\Delta t\,|\theta)|\, i = 1, 2, \ldots, N, j = 1, 2, \ldots, N\}$ are the elements of the matrix $P(\Delta t\,|\theta)$ defined in (3.4), and $\pi : \mathbb{N}_0^M \to [0, 1]$ is the multivariate Poisson probability mass function, i.e.

$$\pi\left(n_k\,|\lambda_{x,k}\Delta t\right) = \prod_{l=1}^{M} \frac{\left(\lambda_{x,k}^l \Delta t\right)^{n_k^l}}{\left(n_k^l\right)!} e^{-\lambda_{x,k}^l \Delta t} \text{ for } x = 1, 2, \ldots, N. \tag{3.8}$$

For $m = 1, 2, \ldots, M$, the quantity $\lambda_{l,k}^m \Delta t$ is our best approximation of the Cox mean measure $\Lambda_{t_k}(l) - \Lambda_{t_{k-1}}(l)$ (cf. (3.6)); note how our independence assumption simplifies the structure of the Poisson probability in (3.8).

We are primarily interested in the posterior distribution of (X_k, θ) given $\mathbf{n}_k$; we introduce the notation

$$L_k\left(x, \theta\,|\mathbf{n}_k\right) := \sum_{\mathbf{X}_k \in \{0,1,\ldots,N\}^k \times \{x\}} H_k\left(\mathbf{X}_k, \mathbf{n}_k, \theta\right). \tag{3.9}$$

Combining (3.7) and (3.9), the recursive relationship

$$L_{k+1}\left(x, \theta\,|\mathbf{n}_{k+1}\right) = \sum_{l=1}^{N} L_k\left(l, \theta\,|\mathbf{n}_k\right) p_{l,x}\left(\Delta t\,|\theta\right) \pi\left(n_{k+1}\,|\lambda_{l,k+1}\Delta t\right) \tag{3.10}$$

follows immediately. However, for the Markov chain model in mind, even the simplified expression (3.10) is far too complicated to allow exact calculation, let alone detailed analysis. As a matter of necessity, we make some simplifying assumptions.

We follow Rogers and Yousaf (2002, pp. 384–385) in using a conditional-independence calibration. In our model, it is easy to imagine the situation where we have information with regards to the process over a very long time span; we postulate that the likelihood can be factorised as

$$L_k\left(x, \theta\,|\mathbf{n}_k\right) = \mu_k\left(x\,|\mathbf{n}_k\right) \phi_k\left(\theta\,|\mathbf{n}_k\right). \tag{3.11}$$

On the one hand, having seen a great amount of data, we expect to have a fairly accurate idea of the parameter θ; the values of θ will largely be determined by the long-run average behaviour of the market. On the other hand, as a result of the ergodicity of the Markov chain, the posterior distribution of X_k will be more influenced by recent events. As a consequence, an assumption of approximate conditional independence of the posterior distributions of X_k and θ is within reason.

We assume moreover, for $k \in \mathbb{N}$, that ϕ_k takes the form

$$\phi_k\left(\theta\,|\mathbf{n}_k\right) \propto \exp\left\{-\frac{1}{2}\left(\theta - \hat{\theta}_k\right) \cdot S_k\left(\theta - \hat{\theta}_k\right)\right\} \tag{3.12}$$

for a positive symmetric matrix S_k. If we suspect that we have very nearly identified the true value of θ, then such a quadratic approximation to the likelihood is quite natural. On the one hand, it stems from a normal appproximation, on the other, from the Taylor expansion of the log-likelihood (with the gradient being negligible near its maximum).

Combining (3.10) with (3.11) and (3.12), we obtain, for $k \in \mathbb{N}$,

$$\begin{aligned} &L_{k+1}\left(x, \theta\,|\mathbf{n}_{k+1}\right) \\ &\quad \propto \sum_{l=1}^{N} \mu_k\left(l\,|\mathbf{n}_k\right) p_{l,x}\left(\Delta t\,|\theta\right) f\left(n_{k+1}\,|\lambda_{l,k+1}\Delta t\right) \exp\left\{-\frac{1}{2}\left(\theta - \hat{\theta}_k\right) \cdot S_k\left(\theta - \hat{\theta}_k\right)\right\}. \end{aligned} \tag{3.13}$$

Armed with initial values $\hat{\theta}_0$, S_0 and μ_0, we update the maximum likelihood estimate iteratively. At step $k+1$, assume $\hat{\theta}_k$, D_k, S_k and μ_k known. Optimising

$$\sum_{l=1}^{N} L_{k+1}\left(l, \theta\,|\mathbf{n}_{k+1}\right) \tag{3.14}$$

over $\theta \in \Omega$, we find the new maximum likelihood estimate $\hat{\theta}_{k+1}$ of θ. Substituting $\hat{\theta}_{k+1}$ into (3.14) and differentiating, we find the the Hessian S_{k+1}. Finally, μ_{k+1} is calculated from

$$\mu_{k+1}(x, |\mathbf{n}_{k+1}) \propto \sum_{l=1}^{N} \mu_k(l\,|\mathbf{n}_k) \int_{\Omega} p_{l,x}(\Delta t\,|\theta) f(n_{k+1}\,|\lambda_{l,k+1}\Delta t) \exp\left\{-\frac{1}{2}\left(\theta - \hat{\theta}_k\right) \cdot S_k \left(\theta - \hat{\theta}_k\right)\right\} d\theta; \quad (3.15)$$

we scale the entries of μ_{k+1} to have unit sum.

Having designed an estimation method, the next logical step is to investigate methods of comparing the fits of different models, for instance to determine the optimal number of states. As a point of departure, one may consider the one-step prediction

$$\hat{n}_{k+1}^{l} := \sum_{m=1}^{N} \mu_k(m\,|\mathbf{n}_k)\,\lambda_{m,k}^{l}. \quad (3.16)$$

Comparing $\hat{n}_{k+1}^{l}$ with the observed value n_{k+1}^{l} provides us with a measure of the fit of the model. However, we cannot attach very much weight to such an error estimate, since n_{k+1}^{l} is a Poisson variable: its variance is the same as its mean.

Direct comparison of maximum likelihood values across models proves to be the best measure for comparison of fit, provided that we take into account the fact that this cannot be done directly, due to the missing constant of proportionality in (3.13). Experience has shown that, when ignoring this constant, likelihood values decrease dramatically when the number of states is increased; this creates the illusion that increasing the number of states worsens the fit, whereas this is emphatically not the case. This misconception can be remedied by calculating the missing constant of proportionality for each model, and dividing the maximum likelihood by it.

Comparing maximum likelihood values at each timestep may be misleading, because a single maximum likelihood value is certainly not representative of the performance of the model over the whole period of estimation. In fitting this model, we will, instead, compare rolling likelihoods, where a rolling likelihood is simply the product of one-period likelihoods over the period of estimation.

3.4 Fitting the model to FTSE 100 index futures

The lazy—or impatient—statistician has a natural inclination to aggregate temporal data to some fixed time interval. However, Engle and Russell (1998, pp. 1127–1128) warns that choosing too short an interval may introduce heteroskedasticity into the data; on the other hand, choosing too long an interval risks losing the microstructure features of the data, mitigating the advantages of moving to high frequency data in the first place.

Having experimented with various bin sizes, our eventual choice of the bin size Δt as 60 seconds was not so much influenced by statistical and economical as by practical considerations. Using a very computer-intensive estimation procedure on a huge data set (cf. Table 2.2), we found, as could be expected, that small values of Δt resulted in the estimation process taking agonisingly slow, often without any noticeable benefit in accuracy. On the other hand, choosing Δt too big often brought perils of a different kind, namely having to numerically deal with infinity; more often than not, this was caused by the difficulty of calculating the Poisson probability mass function (3.8) accurately for large intensities (we were reluctant to use the normal approximation).

For the implementation of the estimation procedure, we use Scilab 2.6 (INRIA Meta2 Project and ENPC Cergrene, 2001), an open source scientific software package developed by the French National Institute for Research in Computer Science and Control (INRIA) and the National School of Civil Engineering (ENPC) (Scilab Consortium, 2003). The heart of the estimation procedure is an optimisation routine; for this we chose the Feasible Sequential Quadratic Programming (FSQP) algorithm by Lawrence, Zhou and Tits (1997); we accessed this from Scilab using the interface by Delebecque (2000).

We concentrate on testing the model on the data, rather than obtaining exact values for the estimates. To this end, we split the data into subsets of manageable size, and experiment with different flavours of the model on these subsets. Given such a subset, we first aggregate observations into bins of Δt seconds, and then proceed in the following way.

We set all the off-diagonal entries of the transition rate matrix Q to the same convenient constant, namely one hour. We are free to do this, as estimation has proven fairly insensitive to the initial estimates for the transition rates of the Markov chain. The initial law μ_0 of the Markov chain is then calculated as its equilibrium distribution; this turns out to be the uniform distribution on $\{1, 2, \ldots, N\}$. Likewise, the initial value S_0 of the Hessian does not influence the estimation procedure noticeably beyond the first few iterations; for convenience, we set it equal to the identity matrix. Through (3.12), this naturally extends to taking the initial distribution of θ as normal with identity covariance, centered around the initial estimate of θ.

In contrast, care is needed in choosing reasonable initial estimates for the Poisson intensities for each of the M transaction sequences. For transaction sequence l, we first find the intraday and intraweek indices $\hat{d}^l_1, \hat{d}^l_2, \ldots, \hat{d}^l_{510}$ and $\hat{w}^l_1, \hat{w}^l_2, \ldots, \hat{w}^l_5$; together, they constitute the seasonal index. Deseasonalising and rearranging each sequence, we find the N equispaced percentiles $p_{\frac{100}{N+1}}, p_{\frac{200}{N+1}}, \ldots, p_{\frac{100N}{N+1}}$; after dividing by Δt, they become the initial frequency estimates $\hat{f}^l_1, \hat{f}^l_2, \ldots, \hat{f}^l_N$.

Having found initial values for the estimates, we follow the steps outlined in Section 3.3. However, we need to make some minor adjustments to ease the fitting exercise.

The normal approximation (3.12) is highly accurate whenever our parameter estimate is close to the true value. However, far out in its tail, its gradient often fails to provide information about the whereabouts of the true mean. This problem is aggravated by our use of central finite difference approximations for derivatives. In situations where this may occur (for instance, in the initial stages of the iteration), we address this problem by first performing the optimisation with, instead of (3.12), the rather more heavy-tailed Cauchy density

$$\phi_k\left(\theta\,|\mathbf{n}_k\right) \propto \frac{1}{1 + \left(\theta - \hat{\theta}_k\right) \cdot S_k \left(\theta - \hat{\theta}_k\right)}.$$

We then substitute the optimising value of θ for $\hat{\theta}_k$, and proceed as usual.

In updating the law of X, we also approximate (3.15) by assuming that the posterior distribution of θ is the point mass at $\hat{\theta}_{k+1}$, i.e. substituting

$$\mu_{k+1}\left(x, |\mathbf{n}_{k+1}\right) \propto \sum_{l=1}^{N} \mu_k\left(l\,|\mathbf{n}_k\right) p_{l,x}\left(\Delta t\,|\theta\right) f\left(n_{k+1}\,|\lambda_{l,k+1}\Delta t\right)$$

for (3.15). In this way, we avoid repeatedly integrating over a large number of dimensions.

Finding the most efficient way of keeping the Hessian updated remains a moot point; apart from the optimisation itself, this is the singularly most expensive calculation in the whole estimation process. Several methods for approximating the Hessian exist, most of which involve diagonal matrices. Rogers and Zane (1998, p. 11) suggests using the diagonal matrix of squares of the parameter estimates; this method is not suitable to our case, since the difference in magnitude of our parameters (typically, the Poisson intensities will be much larger than the Markov transition rates) may occasionally lead to a badly scaled matrix. Likewise, the fact that our parameter estimates are typically highly correlated prohibits us from following Rogers and Yousaf (2002, p. 383) in only calculating the diagonal entries.

We obtain a slight speed improvement by requiring, at first, that the Hessian be updated at every iteration; as estimation continues, we update the Hessian less and less frequently. The motivation for this is that, once the basic structure of the Hessian has been established, we expect that its entries will stabilise to growing approximately linearly with time. Our situation is analogous to the estimation of the mean of a distribution using a sequence of noisy observations of the mean; having seen $n-1$ observations, the n^{th} observation carries a weight of $\frac{1}{n}$ in the estimation. In calculating the Hessian in a conditional-independence calibration setting, the most recent observations tend to carry relatively little weight in comparison to the average over earlier times (Rogers and Yousaf, 2002, p. 384).

Finally, to enable comparison between different models, we integrate the total likelihood (3.14) by Monte Carlo methods, adopting a simple example by Brooks and King (2003, p. 60).

We present two examples. In the first instance, we fit our model to trades in FTSE 100 futures (with maturity data March 2002) during February 2002. As an example of what happens during estimation, we first take a graphical tour of the estimation of one model—a Markov chain with four states—on 23 February, the 16$^{\text{th}}$ day of our estimation period. The observed 60-second transaction counts for this day are given in Figure 3.2, as are the 510 estimates of the long-term intensities $\hat{f}_1, \hat{f}_2, \hat{f}_3, \hat{f}_4$. In comparison, Figure 3.3 contains the seasonally adjusted trade intensities; notice how close these estimates follow the development of trades over the course of the day.

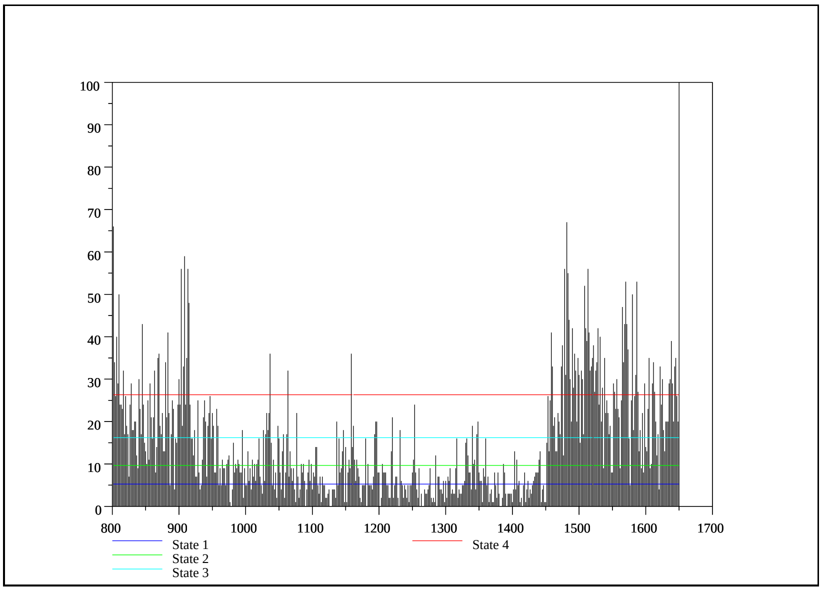

Figure 3.2: Transaction counts and estimated intensities for trades in FTSE 100 futures (maturing March 2002) on 23 February 2002

Figures 3.4 and 3.5 depict, respectively, the estimated law μ of the Markov chain and the estimated transition rates. The colouring of these graphs are consistent with Figure 3.3; in Figure 3.5, the graph of each transition rate is coloured according to the destination state. Studying these two graphs closely, we notice an inconsistency which has proven difficult to explain. Even though the estimated transition rates are quite low, Figure 3.4 leads us to believe that the Markov chain changes states very often. Many factors may cause this effect; most probably, our estimation procedure either finds it difficult to distinguish the present state of the Markov chain, or a great number of the transition rates are indeterminate. Future research will attempt to provide an explanation for this problem, if not a solution.

Having fitted several Markov models, we now turn to the evaluation thereof. Table 3.1 summarises rolling likelihood values for four different times in the month of February, for a number of different Markov models. As we would expect, the rolling likelihood increases with the number of states used. Choosing the best model entails weighing the increase in the rolling likelihood against the number of parameters to be estimated. Clearly, a three-state model fits the data much better than a model with two states; balancing cost against our measure of fit, we opt for a four-state Markov chain.

Although rolling likelihoods do not provide absolute measures of how well the models describe our data, Figure 3.6 is reassuring; the four-state model's one-step predictions are, on average, reasonably accurate. Note the heavy tails: this is almost certainly due to outlier bursts of trades.

In our second example, we attempt to model bids, asks and trades during May 2002 pertaining to futures maturing in June 2002. We treat these three sequences of transactions as if they were trades on different assets. The results of a number of different models are given in Table 3.2. Even at a first glance, the patterns in Table 3.2 differ fundamentally from the patterns we saw in Table 3.1. The estimation of chains with too few states have proven numerically unstable (if not downright impossible); estimation of three assets with only two possible arrangements of intensities, for instance, cannot possibly return positive probabilities when one or more of the sequences of transactions deviates from the expected pattern. As a consequence, zero Poisson probabilities are returned, a feasible estimate of the law of the Markov chain becomes impossible and the process derails.

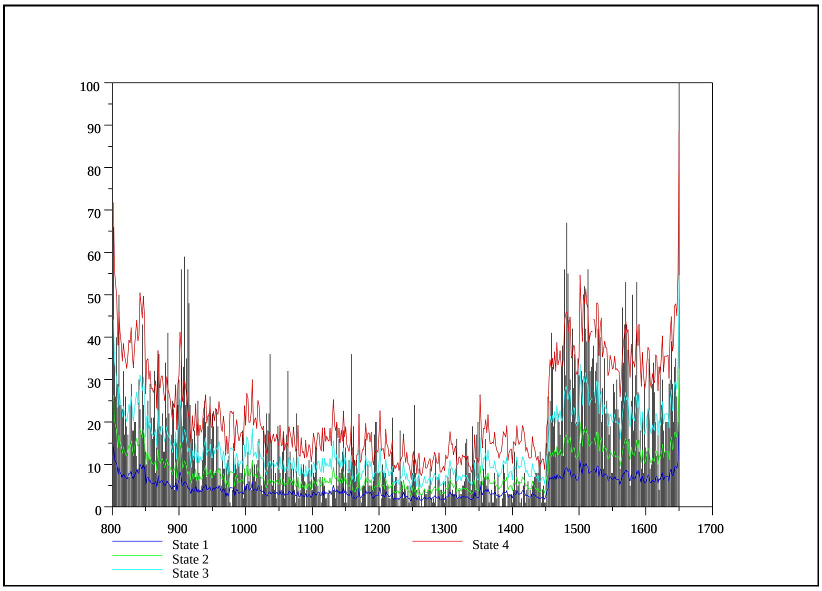

Figure 3.3: Transaction counts and estimated intensities for trades in FTSE 100 futures (maturing March 2002) on 23 February 2002

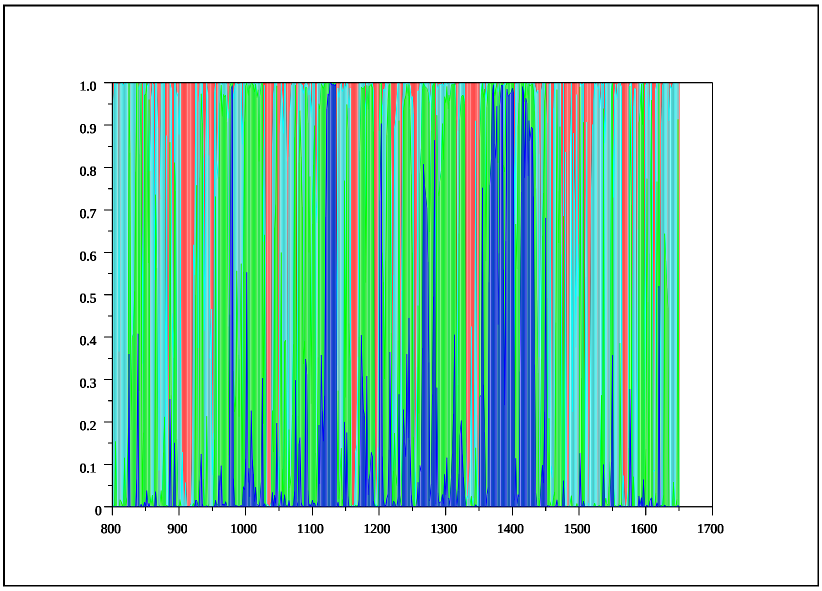

Figure 3.4: Estimated law of four-state Markov chain fitted to trade counts in FTSE 100 futures (maturing March 2002) on 23 February 2002

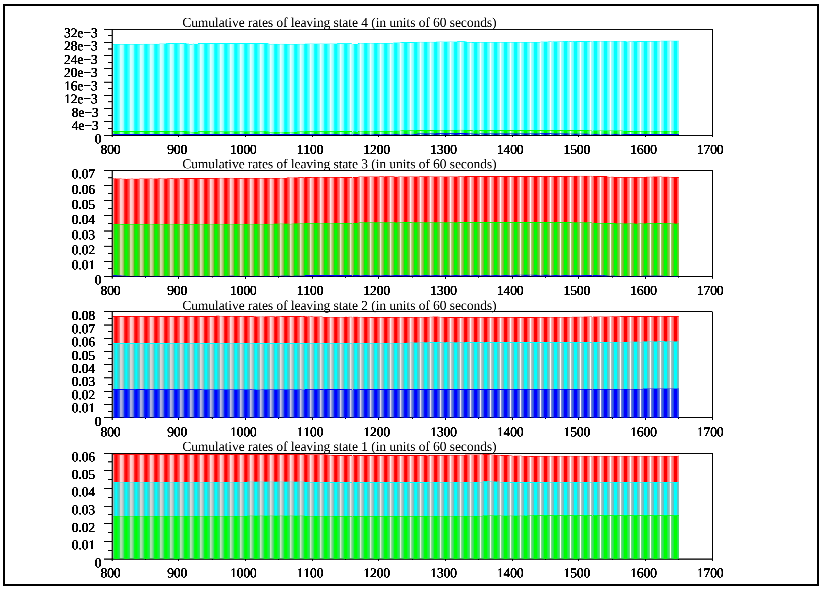

Figure 3.5: Estimated transition rates of four-state Markov chain fitted to trade counts in FTSE 100 futures (maturing March 2002) on 23 February 2002

Number of states	**Number of parameters**	**Rolling likelihood**			
		1 Feb, 10:51	**4 Feb, 11:51**	**6 Feb, 15:51**	**26 Feb, 16:30**
2	4	-157.904	-776.928	-3 646.925	-13 373.978
3	9	-120.462	-565.280	-2 631.432	-9 631.049
4	16	-94.502	-452.923	-2 021.057	-7 001.032
5	25	-83.659	-390.319	-1 733.944	-5 837.117
6	36	-82.280	-400.784	-1 609.711	-5 574.074
7	49	-82.758	-372.249	-1 473.718	-5 029.464
8	64	-77.552	-347.037	-1 377.439	-
9	81	-73.843	-323.613	-	-
10	100	-70.746	-	-	-[a]

[a]Missing values are due to the estimation process being stopped: no additional benefit could be derived from these time-consuming calculations.

Table 3.1: Comparison between Markov-Cox models of trades in February 2002

Number of states	**Number of parameters**	**Rolling likelihood**			
		7 May, 15:51	**8 May, 11:40**	**28 May, 15:17**	**31 May, 16:30**
2	8	-7 906.231	-9 708.793	-105 725.565	-[a]
3	15	-5 905.276	-6 576.533	-63 142.731	-[b]
4	24	-5 160.894	-5 743.364	-[c]	-
5	35	-4 584.647	-5 129.131	-19 447.203	-24 811.603
6	48	-4 091.453	-4 575.658	-17 564.977	-22 583.834
7	63	-3 693.891	-4 133.671	-15 634.816	-18 653.703
8	80	-3 366.553	-[d]	-	-

[a]Estimation of a two-state chain crashes repeatedly, due to inability to calculate Poisson probabilities.
[b]Estimation of a three-state chain crashes repeatedly, due to inability to calculate Poisson probabilities.
[c]Estimation process was stopped in view of the results of two- and three-state estimation.
[d]Estimation process was stopped, for reasons of economy.

Table 3.2: Comparison between Cox-Markov models of trades in May 2002

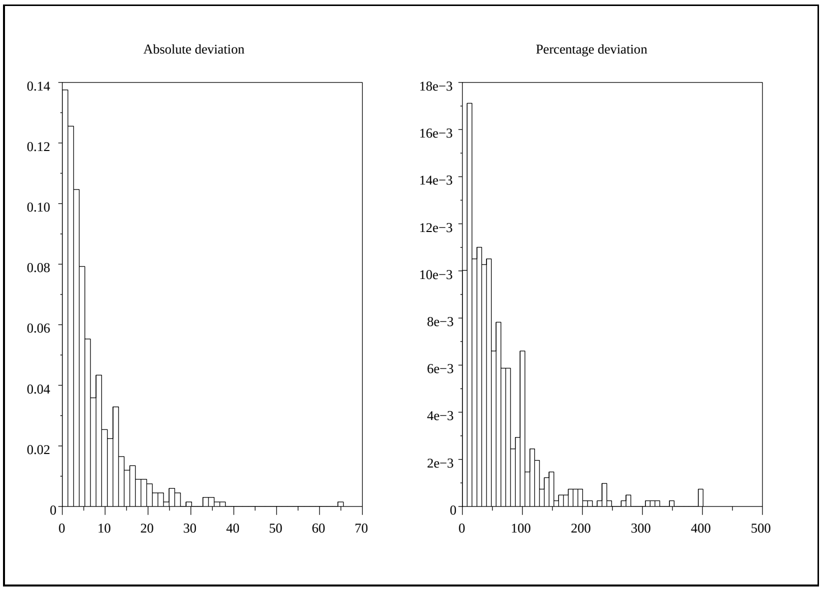

Figure 3.6: Histograms of absolute and percentage deviation, fitting a four-state model to trades in February

Since the rolling likelihoods for the models that were able to complete the month successfully are actually very close to each other, choosing the best model reduces to the availability and cost of computing power. Figure 3.7 is a histogram of the mean percentage deviation from one-step predicted values; once again, we observe a number of huge outliers.

We conclude this chapter with a number of snapshots from the estimation of the six-state Markov chain.

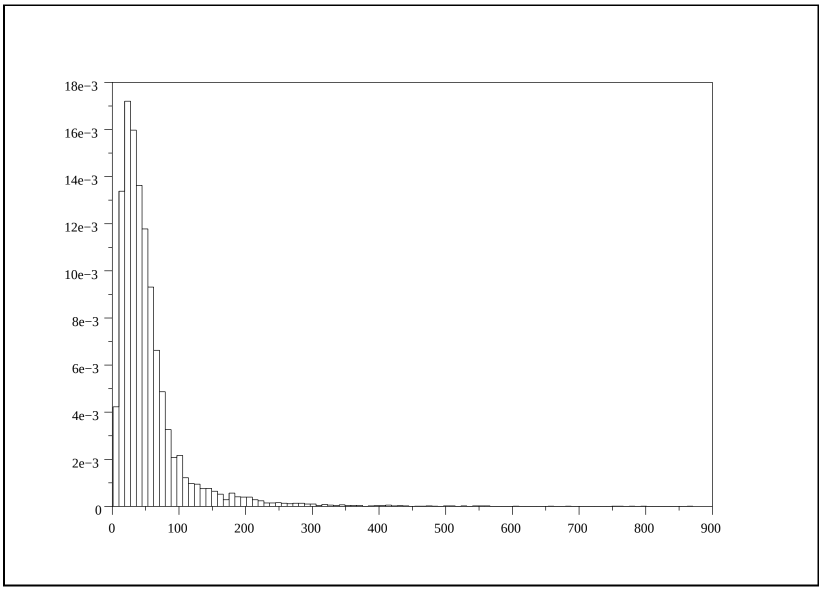

Figure 3.7: Histogram of average percentage deviation of a six-state model, fitted to bids, asks and trades in May 2002

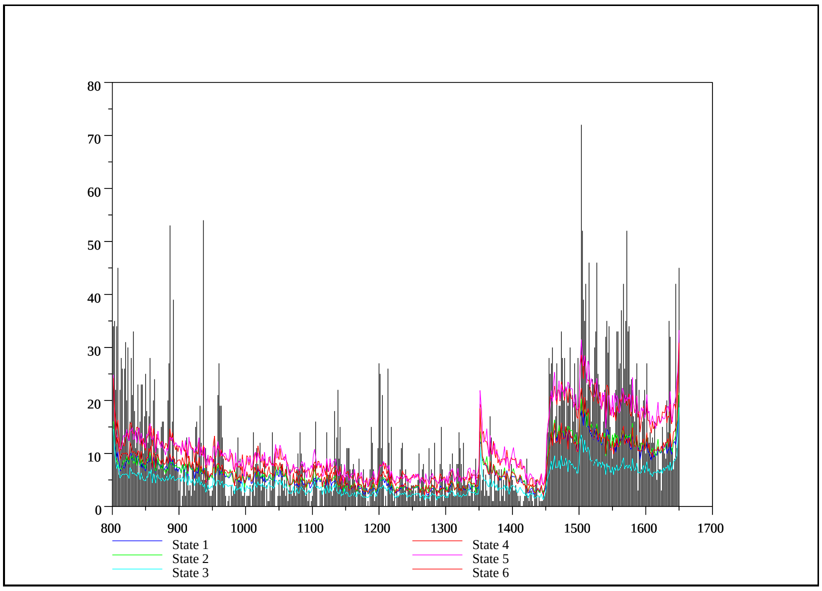

Figure 3.8: Estimated seasonally adjusted ask intensity for a six-state Markov chain, 1 May 2002

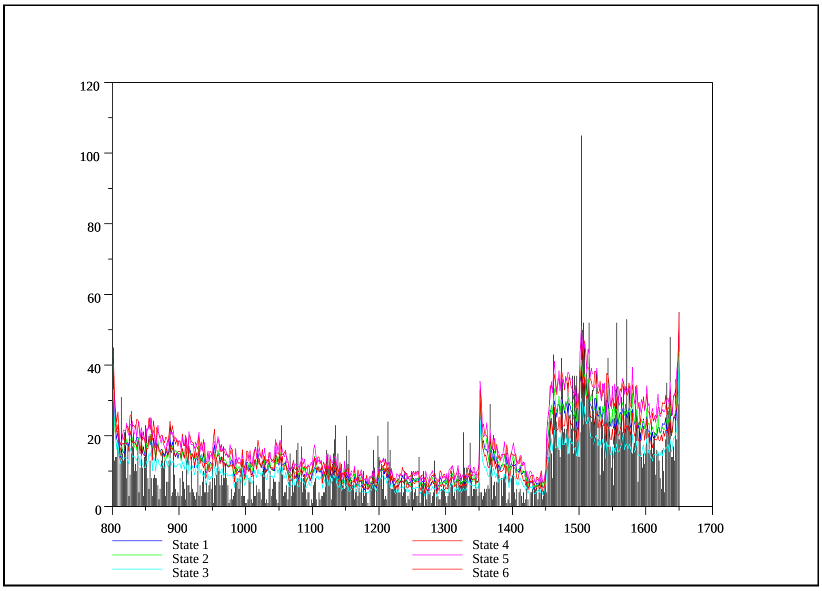

Figure 3.9: Estimated seasonally adjusted bid intensity for a six-state Markov chain, 1 May 2002

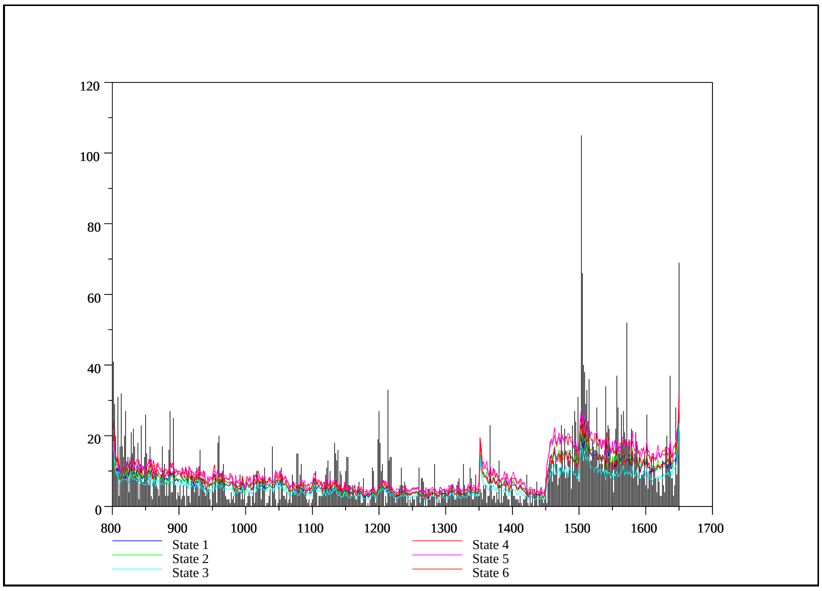

Figure 3.10: Estimated seasonally adjusted trade intensity for a six-state Markov chain, 1 May 2002

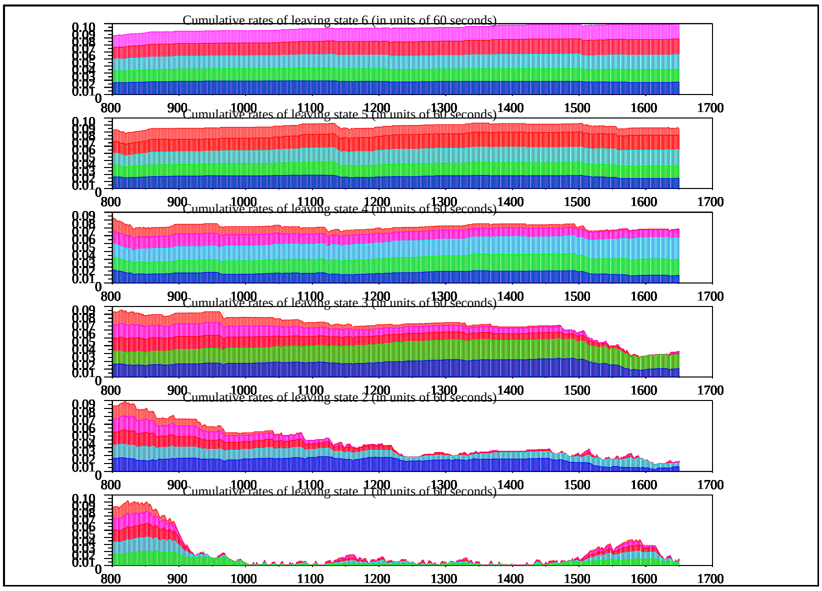

Figure 3.11: Estimated transition rates of a six-state Markov chain, 1 May 2002

Chapter 4

Conclusion

Our study was concerned with modelling the distinguishing features of high frequency transactions data. To this end, we investigated the salient features of transaction counts in FTSE 100 index futures. Several interesting patterns have emerged, most notably the changes in cluster patterns brought about by the technological development of market trading structures.

Several exciting research opportunities remain. For instance, given that activity levels on American markets reach their peak in the morning and the indication is that London derivatives markets reach their peak in the afternoon, the period 15:00–19:00 GMT must be the the busiest part of the international trading day. It would be interesting to study the influence these markets on opposite sides of the globe have on each other.

We have also suggested a new framework for the modelling of times of transactions on financial markets, using mutually independent Cox processes for the modelling of individual trade processes. In this framework, temporal dependence as well as interdependence between assets are completely explained by a latent independent Markov chain. This Markov chain represents the state of the market; it also drives the individual Cox intensities.

Having fitted models to several subsets of our data, they seem to perform satisfactorily. However, evaluation of the performance of this class of models is hampered by the absence of suitable representative measures of fit. Although we are able to compare the merits of models with different numbers of states, we are as yet unable to quantify the extent to which our model explains movement in transaction counts.

In addition, although other studies have found a clear connection between the number and times of trades and price movements, we are not clear on the best way of extending our modelling framework to allow for price movements. Further research into application of the theory of marked Poisson processes may yet yield fruitful results.

As is evident from Chapter 2, we are in the extraordinary position that, instead of having the usual problem with too little data, we are at risk of being overwhelmed by the sheer volume of information. Fortunately, continuous development in computing power will alleviate this problem, if not completely solve it. As it is, a computationally intensive estimation procedure such as ours is simply not economically viable, at least not on the high-frequency timescale that we expect these results will be required. Some benefit may be derived from porting our implementation from an interpreted to a compiled platform; we are also looking forward to experimenting with a number of new optimisation and utility algorithms.

Bibliography

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Vega, C. (2002), *Micro effects of macro announcements: real-time price discovery in foreign exchange*, Working paper 8959, National Bureau of Economic Research, Cambridge (MA).
*http://papers.nber.org/papers/w8959.pdf

ap Gwylim, O., Buckle, M. and Thomas, S. (1999), 'The intra-day behaviour of key market variables for LIFFE derivatives', in P. Lequeux, ed., *Financial markets tick by tick: Insight in financial markets microstructure*, Wiley, Chichester, chapter 6, pp. 151–189.

Baillie, R. T. and Dacorogna, M. M., eds (1997), *Journal of Empirical Finance: High Frequency Data in Finance*, Vol. 4, Elsevier, pp. 69–293.

Brooks, S. and King, R. (2003), *Monte Carlo inference*, sixth edn, Statistical Laboratory, University of Cambridge.

Clare, A. and Courtenay, R. (2000), 'Financial market reactions to interest rate announcements and macroeconomic data releases', *Bank of England Quarterly Bulletin* pp. 266–273.
*http://www.bankofengland.co.uk/qb/qb000302.pdf

Cox, D. R. and Miller, H. D. (1965), *The theory of stochastic processes*, Chapman and Hall, London.

Dacorogna, M. M., Gençay, R., Müller, U. A., Olsen, R. B. and Pictet, O. V. (2001), *An introduction to high-frequency finance*, Academic Press, London and San Diego.

Delebecque, F. (2000), *Scilab-FSQP interface*, Contributed toolbox, Scilab Group. Accessed online 28 June 2003.
*ftp://ftp.inria.fr/inria/projects/meta2/scilab/contrib/fsqp/fsqp-1.1.tgz

Engle, R. F. (2000), 'The economics of ultra-high-frequency data', *Econometrica* **68**(1), 1–22.

Engle, R. F. and Russell, J. R. (1998), 'Autoregressive conditional duration: A new model for irregularly spaced transaction data', *Econometrica* **66**(5), 1127–1162.

FTSE (2003), *FTSE UK Index Series*. Accessed online 25 June.
*http://www.ftse.com/indices_marketdata/uk_series/index_home.jsp

Goodhart, C. A. E. and O'Hara, M. (1997), 'High frequency data in financial markets: Issues and applications', *Journal of Empirical Finance* **4**, 73–114.

Grandell, J. (1997), *Mixed Poisson processes*, Vol. 77 of *Monographs on statistics and applied probability*, Chapman, London.

Hunter, R. (1998), 'Bad news for open outcry', *Derivatives Strategy* . Accessed online 17 June 2003.
*http://www.derivativesstrategy.com/magazine/archive/1998/0598fea3.asp

INRIA Meta2 Project and ENPC Cergrene (2001), *Scilab*. Accessed online 28 June 2003.
*http://www-rocq.inria.fr/scilab/

Intelligent Financial Services Limited (2003), *Tick Data*. Accessed online 18 June.
*http://www.liffe-style.com/tickinfo.aspx

Jones, G., Kaul, G. and Lipson, M. (1994), 'Transactions, volume and volatility', *Review of Financial Studies* pp. 631–651.

Kim, S.-J. (2003), 'The spillover effects of U.S. and Japanese public information news in advanced Asia-Pacific stock markets', *Pacific Basin Finance Journal* . Forthcoming.
*http://banking.web.unsw.edu.au/staff/profiles/skim/USJAPNews-Kim.pdf

Kingman, J. F. C. (1993), *Poisson processes*, number 3 in 'Oxford studies in probability', Oxford University Press, Oxford.

Lawrence, C. T., Zhou, J. L. and Tits, A. L. (1997), *User's Guide for CFSQP Version 2.5: A C Code for Solving (Large Scale) Constrained Nonlinear (Minimax) Optimization Problems, Generating Iterates Satisfying All Inequality Constraints*, Technical Report TR-94-16r1, Institute for Systems Research, University of Maryland.
*http://www.aemdesign.com/aemdesign/download-cfsqp/cfsqp-manual.pdf

Lequeux, P., ed. (1999), *Financial markets tick by tick: Insight in financial markets microstructure*, Wiley Trading Advantage Series, Wiley, Chichester.

LIFFE (2003*a*), *FTSE 100 Index Future*. Accessed online 25 June.
*http://www.liffe.com/liffeinvestor/introduction/how/products/ftse100futs/what.htm

LIFFE (2003*b*), *Glossary*. Accessed online 17 June.
*http://www.liffe.com/help/glossary/index.html

LIFFE (2003*c*), *Understanding Contract Specifications*. Accessed online 25 June.
*http://www.liffe.com/liffeinvestor/introduction/how/products/ftse100futs/specs.htm

LIFFE Press Office (2002), *Record January Volume for LIFFE*, Press release, 4 February, LIFFE. Accessed online 16 June 2003.
*http://www.liffe.com/press/releases/020204.ht

Lin, S.-J., Knight, J. and Satchell, S. E. (1999), 'Modelling intra-day equity prices and volatility using information arrivals – A comparative study of different choices of informational proxies', in P. Lequeux, ed., *Financial markets tick by tick: Insight in financial markets microstructure*, Wiley, Chichester, chapter 2, pp. 27–64.

MacGregor, P. (1999), 'The sources, preparation and use of high frequency data in the derivatives markets', in P. Lequeux, ed., *Financial markets tick by tick: Insight in financial markets microstructure*, Wiley, Chichester, chapter 11, pp. 305–323.

Meyer, P. A. (1976), 'Un cours sur les intégrales stochastiques', in P. A. Meyer, ed., *Séminaire de Probabilités X: Université de Strasbourg*, number 511 in 'Lecture Notes in Mathematics', Springer-Verlag, Berlin, pp. 245–400.

Reuters (2001), *Chicago Stock Exchange may eye early-morning trade*, News report, 18 July. Accessed online 1 July 2003.
*http://www.thestandard.com/wire/0,2231,7960,00.html

Rogers, L. C. G. and Yousaf, F. A. (2002), 'Markov chains and the potential approach to modelling interest rates and exchange rates', in H. Geman, D. Madan, S. R. Pliska and T. Vorst, eds, *Mathematical Finance - Bachelier Congress 2000: Selected Papers from the First World Congress of the Bachelier Finance Society, Paris, June 29–July 1, 2000*, Bachelier Finance Society, Springer, Berlin and Heidelberg, pp. 375–406.

Rogers, L. C. G. and Zane, O. (1998), *Designing and estimating models of high-frequency data*, Unpublished paper, Department of Mathematics, University of Bath.
*http://www.statslab.cam.ac.uk/~chris/papers/tick.pdf

Rydberg, T. H. and Shephard, N. (1999), *Modelling trade-by-trade price movements of multiple assets using multivariate compound Poisson processes*, Working paper 1999-W23, Nuffield College, Oxford. Draft.
*http://www.nuff.ox.ac.uk/users/rydberg/multi.pdf

Rydberg, T. H. and Shephard, N. (2000), 'A modelling framework for the prices and times of trades made on the New York Stock Exchange', in W. J. Fitzgerald, R. L. Smith, A. T. Walder and P. C. Young, eds, *Nonlinear and nonstationary signal processing*, Cambridge University Press, Cambridge.

Scilab Consortium (2003), *Launch of the Scilab Consortium dedicated to scientific computing*, Press release, 22 May. Accessed online 29 June.
*http://www.inria.fr/presse/pre98.en.html

Snyder, D. L. (1975), *Random Point Processes*, Wiley-Interscience, New York.

Turan, S. (2003). Correspondence by electronic mail, 10 March.

WebFinance (2003), *The biggest, best investing glossary on the web.* Accessed online 17 June.
*http://www.investorwords.com