LIKELIHOOD ANALYSIS OF A DISCRETE BID/ASK PRICE MODEL FOR A COMMON STOCK QUOTED ON THE NYSE

Aurora Manrique
Departamento de Economía e Historia Económica,
Universidad de Salamanca, 37008 Salamanca, Spain
amg@gugu.usal.es
and
Neil Shephard
Nuffield College, Oxford University, Oxford OX1 1NF, UK
neil.shephard@nuf.ox.ac.uk

December 2, 1997

Abstract

This paper develops methods to perform likelihood inference for Hasbrouck's (1996) discrete bid/ask price model for a stock in a security market. Our approach is to sample from the distribution of the continuous latent variables, (the bid cost, the ask cost and the efficient price), given the most recent updated values of the latent points at the previous and next period, the discrete observed variables, (the bid and ask quotes), and some parameter estimates, within a Metropolis-Hastings algorithm. The method is applied to estimate the bid and ask quotes for Alcoa for all trading days in 1994, a data set previously analyzed in Hasbrouck (1996).

Some key words: Bid/ask quotes; Discreteness; Gaussian limited dependent process; Latent variables; Markov chain Monte Carlo; Metropolis-Hastings algorithm.
JEL classification: C13; C22; G1

1 Introduction

1.1 Security markets

This paper deals with the estimation of a dynamic model of discrete bid and ask quotes for a stock in a security market proposed by Hasbrouck (1996).

The basic trading process in a security market, (e.g. the New York Stock Exchange) can be described as follows. A market maker or quote setter in the security, posts a bid price (a price at which he/she is willing to buy), an ask price (a price at which he/she is willing to sell), and a quote size (the maximum number of shares he/she is willing to buy or sell at these prices). The Exchange regularly disseminates (communicates) the highest bid and the lowest ask price per stock, as well as the number of shares that can be bought or sold at these prices.

Specialists and floor brokers form the trading crowd. They are going to transact buying and selling securities. A floor broker represents an order (an intent to buy or sell a security) of one of his/her customers such as a pension fund or private individual. Floor brokers can represent orders in all securities quoted in the market. All trading in a given stock, is centralized at the assigned trading post and panel location for that security. The specialist representing the individual stock stays at the trading post. He/she determines the quotation size to be displayed.

He/she can also buy for his/her own account but only after exposing the orders to the trading crowd. See Hasbrouck and Sofianos (1992) for details on the role and regulation of specialists in the NYSE. Most orders reach the post where the security is traded electronically (the SuperDot system in the NYSE). However, large or difficult orders are walked to the post by floor brokers.

When a floor broker receives an order, he/she can either leave it with the specialist or bid for/offer the stock on behalf of his/her customer. The trader (specialist or floor broker), exposes the order to the crowd as a quote. Trade occurs (the order is executed), when a specialist or floor broker accepts one of the quotes, either hitting the bid or lifting the offer. Orders may be executed against each other, against the specialist's own inventory, or against an order represented by another floor broker. Before buying or selling for his/her own account, the specialist has to expose orders to the trading crowd. If an order is executed, transactions must be reported by the specialist or floor broker representing the seller. Trade takes place at or within the posted quotes but never at worse prices than those displayed. Orders may be executed at a better price than the quoted price.

Example. *The NYSE quotes stock XYZ at* \$20 *bid for* 30,000 *shares,* 20,000 *offered at* \$20 1/4*. Floor broker A comes with a market order (it will be executed at the most advantageous price obtainable after being made public to the crowd) to buy* 5,000 *shares. In an attempt to do better than the offer price, he/she bids* \$20 1/8 *for* 5,000*. Floor broker B hits the bid and the two brokers complete the transaction. Floor broker A got price improvement: instead of buying at* \$20 1/4 *(the posted offer), he/she bought at* \$20 1/8*. (Hasbrouck, Sofianos, and Sosebee (1993, p. 14)).*

Orders may also be stopped by the specialist. By stopping the order, the market maker guarantees execution at the prevailing quote while attempting to get a price improvement. Orders representing both the buying and selling side of a transaction may also be crossed at a more advantageous price, subject to certain rules.

Posted quotes may be revised in the absence of trade. The quote change also needs to be reported. All the information regarding quotes and trades is kept in the audit trial for surveillance operations. See Hasbrouck, Sofianos, and Sosebee (1993) for a detailed description of NYSE systems, trading rules and variations of the basic trading process.

By posting the bid and ask quotes, the quote setter is risking that other traders may act on this information to his/her disadvantage. For example, if the quotation size is large, either on the bid or on the ask side, some dealers may try to profit from the free trading option provided by the large size displayed. They will then quote on the same side attempting to get their orders filled ahead of the large size displayed. Most financial markets force these dealers to improve the price if they wish to acquire order precedence. In practice, the benefits from posting quotes and being perceived as an active market participant is likely to outweigh the direct costs of quote formation. For example, if a bank in the foreign exchange market continuously refuses to quote bid and ask prices or quotes uncompetitive spreads, other traders may reduce their contact and exclude the bank from the network that makes up the market (Lyons (1991)).

The difference between the bid and ask quotes is the spread. Incoming order flow depends on the quoted spreads. Thus, a market maker may diminish the frequency of incoming order arrival simply by widening the bid/ask spread. Large bid/ask spreads make trading expensive, especially for small traders. Spreads are expected to decrease with trading activity because the fixed costs of market making are spread over more traders, and to increase with price volatility because dealers are risk averse.

In most markets quoted spreads must be some multiple of the tick size (the minimum price variation), usually 1/8. This constraint appears to be binding for low-price stocks. For example on the NYSE a stock priced at \$5 or more per share trades in ticks of 1/8 dollar. Discreteness then appears as an institutional feature of market data and reflects costs of negotiation. In the absence of discreteness restrictions, traders would haggle over very small amounts of money before hitting the bid or lifting the offer.

For high frequency data (e.g. intraday changes in the bid and ask prices), discreteness could considerably distort the econometric analysis and we have to model it explicitly.

1.2 Econometric issues of bid/ask price models

Some econometric studies of discreteness focus on transaction prices, (see for example, Hausman, Lo, and MacKinlay (1992), who present an ordered Probit model of price transaction changes, where the underlying continuous price variable is mapped into a set of discrete prices using breakpoints). Others analyze the discrete bid/ask spread. Papers along these lines include Harris (1994) or Bollerslev and Melvin (1994). Harris (1994) explores how price levels (and indirectly the minimum price increment), are related in cross sections to relative spreads (spread expressed as a fraction of price), quotation sizes and trading volumes. He uses intraday stock quotation spread frequencies to estimate a cross sectional discrete spread model, in which the discrete quote is obtained by rounding to the nearest tick, random draws from a continuous distribution. Bollerslev and Melvin (1994) present empirical evidence that the size of the bid/ask spread in the foreign exchange market is positively related to the exchange rate uncertainty. Since the observed spread is assumed to only take on a fixed number of discrete values, they use an ordered Probit model to estimate the relationship between exchange rate volatility and spread. Exchange rate volatility is measured as the conditional variance of the ask price estimated by an MA(1)-GARCH$(1, 1)$ model.

In a recent paper, Hasbrouck (1996) estimates a model of discrete bid and ask prices. He considers discrete bid and ask quotes as arising from three continuous random variables: the efficient price of the security, a cost of quote exposure on the bid side and a similar cost of quote exposure on the ask side. The bid quote is the efficient price less the bid cost, rounded down to the next tick. The ask quote is the efficient price plus the ask cost, rounded up to the next tick. Exposure costs reflect fixed transaction costs and asymmetric information costs associated with market making. A limitation of using these simple rounding functions is that the analysis cannot be extended to clustering (tendency of transaction prices and quotes to combine in integers, halves and quarters, in decreasing frequency). Other models in the literature, (see for example, Glosten (1994), Chordia and Subrahmanyan (1995), Bernhardt and Hughson (1995) or Cordella and Foucault (1996)) incorporate discreteness effects by allowing a restricted set of traders and permissible interactions. These models typically focus on information costs without an explicit rounding construction to get the discrete bid and ask observations.

1.3 Our approach

In this paper we follow Hasbrouck (1996) in mapping continuous latent variables (efficient price, ask cost and bid cost) into discrete observed variables (bid and ask prices) by a ceiling and floor function. The bid and ask are modelled separately, although they are assumed to follow the same dynamics. This model is a special case of a Gaussian limited dependent process. We say y_t, $t = 1, ..., n$, is a Gaussian limited dependent process if $y_t = h(s_t)$, where at any time y_t is observed, $s_t = (s_{t1}, ...s_{tp})'$ is an unobserved p-dimensional vector, $h(\cdot)$ is not a one-to-one function and s follows a model which can be put in Gaussian state space (e.g., West and Harrison (1997) and Harvey (1989)). Note that s_t can be, and will in practice often be, non-stationary. If y_t is multivariate, then it could include some variables which are identical to elements in the s_t vector. This is a particular form of a non-Gaussian state space form where the 'measurement equation' is a deterministic and non-invertible function of a Gaussian state space form. Examples of this setup also include Tobit structures, Probit structures and disequilibrium processes. An extensive analysis of Gaussian limited dependent processes is given in Manrique

(1997). Some examples of this setup (Tobit processes and disequilibrium models) have been discussed in Manrique and Shephard (1997).

As in Hasbrouck (1996), we assume that the three latent variables evolve independently. A Metropolis-Hastings algorithm can be designed to sample from the density of the latent variables at any time, given the values of the latent variables at any other time, the observations and a fixed parameter value. We postulate a very stylized model which can be criticized on several grounds, as it does not take into account some intricate typical features of market data. In this setup, we consider a simpler Metropolis-Hastings algorithm which draws from the density of the latent variable at any time, given the values of the latent variables for the previous and next period, this period's observation and some parameter estimates. The model is then estimated for Alcoa, a representative stock in the New York Stock Exchange, using the ask and bid prices for this stock for all trading days in 1994. This data has previously been analyzed in Hasbrouck (1996).

The outline of the paper is as follows. Section 2 presents the underlying economic model to generate the bid and ask quotes using the same construction as in Hasbrouck (1996). The Metropolis-Hastings algorithm is given in section 2.1. Section 3 is devoted to parameter inference. Section 3.1 outlines the approximated maximum likelihood approach in Hasbrouck (1996), who uses Kitagawa's (1987) non-linear filtering method. The Metropolis-Hastings algorithm for Bayesian inference is in section 3.2. Results of a Monte Carlo study are discussed in section 4, first with a single data set and then repeating the experiment 100 times. In section 5 the sampler is applied to estimate the bid and ask quotes for Alcoa for all trading days in 1994. Section 6 examines the limitations of the model and suggests possible alternative formulations to incorporate some typical features exhibited by the market data. A slightly more general model is discussed in detail. Section 7 gives some conclusions. All the computations were generated using the high level programing language *Ox* of Doornik (1996).

2 Underlying economic model

2.1 Basic model

A simple bid/ask spread model arises from the following trading process. Assume that m, the implicit efficient price of the security, (the expectation of the security's terminal value conditional on all public information, including the transaction price history), is known to all participants.

A market maker or quote setter posts a price at which he/she is willing to buy (the bid price) and a price at which he/she is willing to sell (the ask or offer price). The difference between the bid and ask quotes is the spread. Transactions occur when active traders arrive and accept one of the quotes (hitting the bid or lifting the offer). Buyers and sellers are assumed to arrive independently and with equal probability. The incoming order flow depends on the quotes.

The agent establishing the bid quote is assumed to be subject to a non-negative cost of quote exposure, $\beta \geq 0$, for small trades, which is assumed to reflect fixed transaction costs, costs associated with being the market maker and asymmetric information costs. With no discreteness restrictions, this agent would quote a bid price of $m - \beta$. Similarly, the agent establishing the ask quote is assumed to be subject to a non-negative cost of quote exposure $\alpha \geq 0$ also for small trades. With no discreteness restrictions, this agent would quote an ask or offer price of $m + \alpha$. The quote exposure costs are defined implicitly by the conditions that ensure the quote setters or market makers' zero expected profits and no ex post regret. Clearly, the bid quote must be lower than or equal to the ask price.

However, in most markets, assets are only traded at a fixed number of prices as quoted spreads must be some multiple of the tick size, usually 1/8. This institutional feature has to be explicitly taken into account for modelling high frequency data of the same scale or smaller than the tick size. Constrained by discreteness, the market maker will post a bid price of $b = Floor\,(m - \beta)$, where $Floor\,(\cdot)$ rounds its argument down to the next tick since $m - \beta$ is the highest price at which he/she is willing to buy. Similarly, as he/she is not willing to sell at any price lower than $m + \alpha$, he/she will quote an offer price of $a = Ceiling\,(m + \alpha)$, where $Ceiling\,(\cdot)$ rounds its argument up to the next tick.

This construction is taken from Hasbrouck (1996). In this setup, random variation in m is enough to induce randomness in the spread even when the bid and ask exposure costs are equal and constant. For example, if $\alpha = \beta = 1/4$, the spread is one tick, if the decimal part of m is between 1/4 and 3/4, and is two ticks otherwise (Hasbrouck (1996, p. 7)).

The market maker avoids the possibility of loss on the incoming trade by asymmetrically rounding up on the ask and down on the bid. If the efficient price of the security is rounded to the nearest tick, in order to avoid the expectation of losing money, a market maker posting bid and offer quotes must round his/her bid price down and his/her offer price up. With symmetric rounding (all prices rounded up, all rounded down, or all rounded to the nearest integer), one or both sides of the quotes might be associated with an expected loss. For example, if the efficient price is 5, and the cost is 1.1, nearest integer rounding yields a bid of 4 and an offer of 6, both of which yield expected losses (Hasbrouck (1996, p. 4)). Furthermore, symmetric rounding may result in identical bid and ask prices if α and β are small, while asymmetrically rounding up on the ask and down on the bid implies that with probability 1 the bid and ask prices will be different. Even with no costs associated with market making, i.e., $\alpha = \beta = 0$, the bid and ask prices could only be identical when the efficient price m is exactly an integer, which happens with probability 0.

This construction is motivated by most simple models of dealer behaviour. Following Glosten and Milgrom (1985), it is typically assumed that there are informed and uninformed traders in the population. For instance, in the foreign exchange market there are liquidity traders and informed traders. The first ones do not speculate but only buy or sell currencies according to what they need for their normal business activity. The second ones profit both from information received by dealing with liquidity traders and from information asymmetries regarding the determinants of the spot exchange rate. m is the expectation of the final value of the security, conditional on all public information. The quote exposure costs are defined by the conditions that $m - \beta$ and $m + \alpha$ ensure the quote setter's zero expected profits and no ex post regret. Due to the asymmetric rounding, a Glosten-Milgrom dealer will achieve a profit (both ex ante and ex post) at each trade. These profits will typically not lead to competitive price cutting. The discreteness restriction ensures that such action, if feasible, will result in a loss. Furthermore, despite these profits, new entrants are not expected to join the market as local time priority is usually enforced so that the probability of execution diminishes with the length of the queue.

In summary, the three underlying continuous variables in the model are the efficient price, m_t, the ask cost, α, and the bid cost,β. The observed variables are the discrete bid and ask quotes. These variables are related by

$$\begin{aligned} b &= Floor\,(m - \beta)\,, \\ a &= Ceiling\,(m + \alpha)\,. \end{aligned} \tag{1}$$

2.2 Modelling framework

A simple dynamic bid/ask model is considered here. At any time, bid and ask prices are observed while the underlying efficient price, the bid cost and the ask cost are unobserved. We can write

$$\begin{aligned} s_t &= (s_{1t}, s_{2t}, s_{3t})'\,, \\ y_t &= [Floor\,\{\exp(s_{1t}) - \exp(s_{2t})\}\,, Ceil\,\{\exp(s_{1t}) + \exp(s_{3t})\}]'\,, \end{aligned} \tag{2}$$

where s_{1t} is interpreted as the logarithm of the efficient price $\log(m_t)$, s_{2t} as the logarithm of the bid cost $\log(\beta_t)$, s_{3t} as the logarithm of the ask cost $\log(\alpha_t)$, y_{1t} as the bid price b_t and y_{2t} as the ask price a_t. We assume s_t follows a model which can be put in Gaussian state space form, (e.g., West and Harrison (1997) and Harvey (1989)). The model for s_t is therefore formulated in terms of the unobserved states. At any time, s_t is linearly related to the unobserved state α_t, but corrupted through the addition of Gaussian noise. Let $s = (s_1, ..., s_n')$ and $y = (y_1, ..., y_n)'$. As before, the Floor function rounds its argument down to the next tick while the Ceil function rounds it up to the next tick.

3 Parameter estimation

3.1 Previous work

Hasbrouck (1996) describes a non-linear state space estimation method for estimating a dynamic bid/ask price model.

The estimation approach follows Hamilton (1984), Hamilton (1994) and Harvey (1989). It is based on a recursive likelihood calculation. This likelihood is numerically approximated as in Kitagawa (1987). Throughout, y^t will denote $(y_1', ..., y_t')'$, and following Hasbrouck (1996), s_{1t} will be interpreted as the efficient price m_t, s_{2t} as the bid cost β_t, and s_{3t} as the ask cost α_t, so that $y_t = (b_t, a_t)' = \{Floor(s_{1t} - s_{2t}), Ceil(s_{1t} + s_{3t})\}'$, and $s_t = (s_{1t}, s_{2t}, s_{3t})'$ as the unobserved state.

Assume that $f(s_t|y^t;\theta)$, the density function of the current state, conditional on current and past information and a fixed parameter value, is known. Then $f(s_{t+1}|y^{t+1};\theta)$ is recursively calculated as follows. Given θ, the density of the state in the next period, conditional on current and past information is

$$f\left(s_{t+1}|y^t;\theta\right) = \int f(s_{t+1}|s_t;\theta) f\left(s_t|y^t;\theta\right) ds_t, \tag{3}$$

where $f(s_{t+1}|s_t;\theta)$ is the state transition density. The conditional probability of observing y_{t+1} given y^t, $\Pr(y_{t+1}|y^t;\theta)$, is defined as the integral of the density of next period's state, given y^t over Q_{t+1}, where Q_{t+1} is the feasible region of the possible values for s_{t+1} consistent with observation $y_{t+1} = \{b_{t+1}, a_{t+1}\}'$. That is,

$$\Pr\left(y_{t+1}|y^t;\theta\right) = \int_{s_{t+1}\in Q_{t+1}} f\left(s_{t+1}|y^t;\theta\right) ds_{t+1}. \tag{4}$$

The joint density of next period's state and next period's observed quotes given y^t is

$$f\left(s_{t+1}, y_{t+1}|y^t;\theta\right) = \begin{cases} f(s_{t+1}|y^t;\theta) & \text{if } s_{t+1} \in Q_{t+1}, \\ 0 & \text{otherwise}, \end{cases} \tag{5}$$

and hence,

$$f\left(s_{t+1}|y^{t+1};\theta\right) = \begin{cases} f(s_{t+1}, y_{t+1}|y^t;\theta)/\Pr(y_{t+1}|y^t;\theta) & \text{if } s_{t+1} \in Q_{t+1}, \\ 0 & \text{otherwise}, \end{cases} \tag{6}$$

which completes the update.

Maximum likelihood estimates are obtained by maximizing the log likelihood function,

$$\log f(y_1, ..., y_n;\theta) = \sum_{t=1}^{n-1} \log \Pr\left(y_{t+1}|y^t;\theta\right) + \log p(y_1;\theta), \tag{7}$$

where $p(y_1;\theta)$ is the density function of the first observation.

Analytical computation of the integrals in this calculation will typically be infeasible in non-Gaussian cases. Kitagawa (1987) suggests approximating $f\left(s_t|y^t;\theta\right)$ by a numerical grid, say C_t^i, $i=1,...$ in the s_t space. Then, the state transition densities become the discrete transition probabilities $\Pr\left(C_{t+1}^i|C_t^j;\theta\right)$ for $i=1,...$ and $j=1,...$ and the conditional probability of observing y_{t+1} given y^t is

$$\Pr\left(y_{t+1}|y^t;\theta\right)=\sum_j \Pr\left(C_{t+1}^j|y^t;\theta\right)\frac{Vol\left(C_{t+1}^j\cap Q_{t+1}\right)}{Vol\left(C_{t+1}^j\right)}, \tag{8}$$

where Vol denotes volume. Note that the intersection $C_{t+1}^j\cap Q_{t+1}$ is empty for most of the cells in the s_{t+1} space. There are several approaches to approximate $\Pr\left(C_{t+1}^i|C_t^j;\theta\right)$, see Hasbrouck (1996) for details.

In this framework, it seems difficult to add endogenous variables such as quote sizes or trades. This will require expanding the set of state variables, which "may run into the curse of dimensionality because of the requirement that the integration of the conditional probabilities be computed numerically over all variables", Hasbrouck (1996, p. 22).

3.2 Bayesian inference

An alternative approach to efficient inference, which is computationally attractive, is to use a Bayesian analysis carried out by Markov chain Monte Carlo (MCMC) methods, (see e.g., Chib and Greenberg (1994), or Gilks, Richardson, and Spiegelhalter (1996)). In this context, this involves deriving methods for simulating from $\theta, s|y$. It turns out that we can simulate from the distributions of $s_t|s_{\backslash t}, y;\theta$, where $s_{\backslash t}=\left(s_1,...,s_{t-1},s_{t+1},...,s_n\right)'$, and $\theta|s$, (note that $s|\theta$ follows a Gaussian state space form, and so the likelihood function can be evaluated using the Kalman filter (e.g., Kalman (1960)). It then follows that we can set up a simple Markov chain Monte Carlo sampler to estimate the model by drawing from $\left(s',\theta'\right)'|y$. Assuming a prior distribution for the unknown parameters θ, the MCMC sampler will proceed as follows:

1. Initialize θ.
2. Sample $s\sim s|y;\theta$, i.e., $s_t\sim s_t|s_{\backslash t}, y_t;\theta$, $t=1,..,n$ inside a Metropolis-Hastings algorithm.
3. Sample $\theta\sim\theta|s$.
4. Repeat from step (2).

This type of algorithm will, (under some rather weak regularity conditions; e.g. Tierney (1994)), converge to a draw from $\left(s',\theta'\right)'|y$ using Markov chain Monte Carlo results. Averaging subsets of these simulations will lead to Bayesian estimators of the parameters. The resulting estimators, based on the mean, median or mode of the posterior density of $\theta|y$, are typically efficient, (if viewed as estimators from a sampling viewpoint), as shown, for example, by Barndorff-Nielsen and Cox (1994, Ch. 4) for a wide class of prior distributions.

In this section, we detail how to perform step (2) of the sampler. In section 4.2.2, we will outline how to sample from the posterior distribution, $\theta\sim\theta|s$, in a simple example.

The scan sampler of de Jong (1997) can be used to deal with the time series dimension of this problem. The scan sampler is a powerful tool which enables us to sample from the distribution of the latent variable s_t, at any time, given all the latent points $s_{\backslash t}$ in only $O\left(np^3\right)$ operations. It is important to note that at every iteration s_t is updated, given some parameter estimates, conditional on the most recent values of all the other points. As we can easily sample from the distribution of $s_t|s_{\backslash t};\theta$ using the scan sampler, the only unresolved difficulty left is sampling from

$$\Pr\left(s_t|s_{\backslash t}, y;\theta\right) \propto \Pr\left(s_t|s_{\backslash t};\theta\right) \Pr\left(y_t|s_t;\theta\right). \tag{9}$$

Let $f\left(s_t\right)$ denote the density of $s_t|s_{\backslash t}, y_t;\theta$ and Q_t the set of values s_t, consistent with observation $y_t = \{b_t, a_t\}'$, i.e.,

$$Q_t = \{s_t\,|b_t \leq \exp\left(s_{1t}\right) - \exp\left(s_{2t}\right) < b_t + \delta,\ a_t - \delta < \exp\left(s_{1t}\right) + \exp\left(s_{3t}\right) \leq a_t\}, \tag{10}$$

where δ is the minimum tick size, (in this case, 1/8). Note that $f\left(s_t\right) = 0$ if $s_t \notin Q_t$.

If we write $s_t|s_{\backslash t};\theta \sim N\left(\gamma_t, \Sigma_t\right)$, sampling from f would involve sampling s_t from a normal distribution truncated to $R_t = [\log\left(b_t\right), \log\left(a_t\right)] \times \left(\log\left(a_t - b_t\right), \infty\right) \times \left(\log\left(a_t - b_t\right), \infty\right)$, and accepting s_t only if $s_t \in Q_t$. However, we can propose from the truncated distribution without checking whether s_t is in Q_t. This can be done inside a Metropolis-Hastings algorithm. At step i, a proposal $\overset{i}{s_t}$ is drawn from some density $q\left(s_t\right)$ and accepted with probability

$$\min\left\{1, \frac{f\left(s_t^i\right) q\left(s_t^{i-1}\right)}{f\left(s_t^{i-1}\right) q\left(s_t^i\right)}\right\}, \tag{11}$$

where $s_t^{i-1} \in Q_t$. If the proposal is not accepted we write $s_t^i = s_t^{i-1}$ and make another proposal. We take a proposal density q, which is proportional to f in Q_t. Hence, if $s_t \in Q_t$ then the probability of accepting it is 1 (the constants of proportionality in (11) disappear) while if $s_t \notin Q_t$ the probability of acceptance is 0 as $f\left(s_t\right) = 0$. In other words, we sample from the true density f and accept the proposal if it satisfies

$$b_t = \mathit{Floor}\left\{\exp\left(s_{1t}\right) - \exp\left(s_{2t}\right)\right\}, \qquad a_t = \mathit{Ceil}\left\{\exp\left(s_{1t}\right) + \exp\left(s_{3t}\right)\right\}, \tag{12}$$

(note that (12) is simply (10) rewritten). Clearly, the proposals will be accepted with high probability if the region Q_t is almost as large as the region R_t from which we sample.

The Metropolis-Hastings algorithm we set up (e.g., Metropolis, Rosenbluth, Rosenbluth, Teller and Teller (1953), Hastings (1970) or Tierney (1994)), is as follows:

1. Initialize s at say s^0; set $k = 0$.
2. Set $t = n$ and then run for $t = n, ..., 1$.
   (a) Set $i = 1$, generate a candidate value s_t^i from q.
   (b) If (11) is 1 write $s_t^{k+1} = s_t^i$ and go to step (3). If (11) is 0 set $i = i + 1$.
   (c) If $i = 50$ write $s_t^{k+1} = s_t^i$ if (11) is 1 or $s_t^{k+1} = s_t^k$ if (11) is 0, and go to step (3).
3. Let $t = t - 1$. If $t > 1$, then repeat from step (2a). If $t = 1$, then write $k = k + 1$ and repeat from step (2).

In short, for each t we perform up to 50 Metropolis rejection steps using a proposal s_t from q and accepting this value when it is in Q_t. If, after 50 rejections the proposal is still outside Q_t, we take $s_t^{k+1} = s_t^k$.

The simulations are not continuous with respect to the model parameters as the acceptance probability in the Metropolis-Hastings algorithm is less than 1. This rules out the use of simulated EM algorithm (e.g., Qian and Titterington (1991), Chan and Ledolter (1995), Ruud (1991) or Tanner (1996, Ch. 4)).

It is important to note that, unlike the estimation method used in Hasbrouck (1996), this procedure allows us to add new state variables to the process straightforwardly. The only difficulty will be to design an appropriate Metropolis-Hastings algorithm to sample from $s_t|s_{\backslash t}, y_t; \theta$ in step (2), and maybe sampling additional parameters in step (3).

Note also that the sampling from $\theta|s$ to estimate another Gaussian limited dependent processes may be generically difficult, even though we can evaluate the likelihood $f(s;\theta)$ and can sample from the full conditionals one at a time, rather than from the full density. Then, it will be possible to add another line in the simulation procedure and to sample the unobserved states α given s and θ, (e.g., using the signal simulation smoother of de Jong and Shephard (1995)). The advantage with the resulting sampler is that it is usually easy to sample from $\theta|\alpha, s$ than from $\theta|s$. Despite the fact that we are adding an unnecessary conditioning variable, the sampler seems to work reasonably well in most cases, (see Manrique (1997) for details).

4 A Monte Carlo study

4.1 Monte Carlo setup

For sake of simplicity, we assume that the logarithm of the efficient price is a random walk and the logarithm of the bid and ask exposure costs are stationary AR(1) processes, i.e.,

$$\begin{array}{rcl} \log(m_t) & = & \log(m_{t-1}) + \sigma_{\varepsilon t}\varepsilon_t, \\ \log(\beta_t) - \mu_t & = & \phi\left\{\log(\beta_{t-1}) - \mu_{t-1}\right\} + \sigma_\nu \nu_{1t}, \\ \log(\alpha_t) - \mu_t & = & \phi\left\{\log(\alpha_{t-1}) - \mu_{t-1}\right\} + \sigma_\nu \nu_{2t}. \end{array} \tag{13}$$

where $|\phi| < 1$, ε_t, ν_{1t} and ν_{2t} are independent, serially uncorrelated with $N(0,1)$ distributions, and $\log(\beta_0) = \log(\alpha_0) = 0$. We take $\log(m_0) = 4$.

The assumption of independent bid and ask exposure costs is appropriate to a market setting in which the bid and ask quotes are set by limit orders, (orders which must be executed at a specified price -the limit price- or at a better price), of different traders. If the quotes reflect the interest of a single quote setter, it would be better to allow for positive correlation between the two costs.

The unconditional mean and variance of the bid and ask costs are then

$$\begin{array}{rcl} \overline{\mu} = E(\beta_t) = E(\alpha_t) & = & \exp\left[\mu + (1/2)\left\{\sigma_\nu^2/(1-\phi^2)\right\}\right], \\ Var(\beta_t) = Var(\alpha_t) & = & \overline{\mu}^2 \cdot \exp(\sigma_\nu^2) - 1. \end{array} \tag{14}$$

In this setup, $\theta = \left(\sigma_\varepsilon^2, \mu, \sigma_\nu^2, \phi\right)'$. The parameter values are chosen so that the artificial data looks rather similar to the real data which will be considered in the next section. This data is also analyzed in Hasbrouck (1996). There are 6,780 observations, 6,528 intraday observations and 251 overnight observations from 252 trading days. An overnight observation is the first observation of the day, and intraday observations are the remaining observations during the day. The first observation of the first day is included within the intraday observations. Bid and ask quotes are reported at the end of fifteen minute intervals. The first observation of a day is typically at 9:30 and the last one is at 16:00. This gives 27 daily points with the exception of 7 days. The bid quotes take values between \$64 3/8 and \$90 1/8, while the ask quotes are between \$64 3/4 and \$90 1/4.

We take $\sigma_\varepsilon^2 = 0.00001$, $(\sigma_\varepsilon = 0.00316)$, $\mu = -3.715$, $\sigma_\nu^2 = 1.05$, $(\sigma_\nu = 1.025)$, and $\phi = 0.4$. A very small value of σ_ε^2 is required to avoid huge variability in the efficient price. The value of the autoregression coefficient is taken as being quite close to Hasbrouck's estimate (0.37). These values give an unconditional mean and variance of 0.0455 and 0.00218 respectively for the bid and ask costs. The simulated bid quotes take values between \$50 1/4 and \$74 1/8, while the

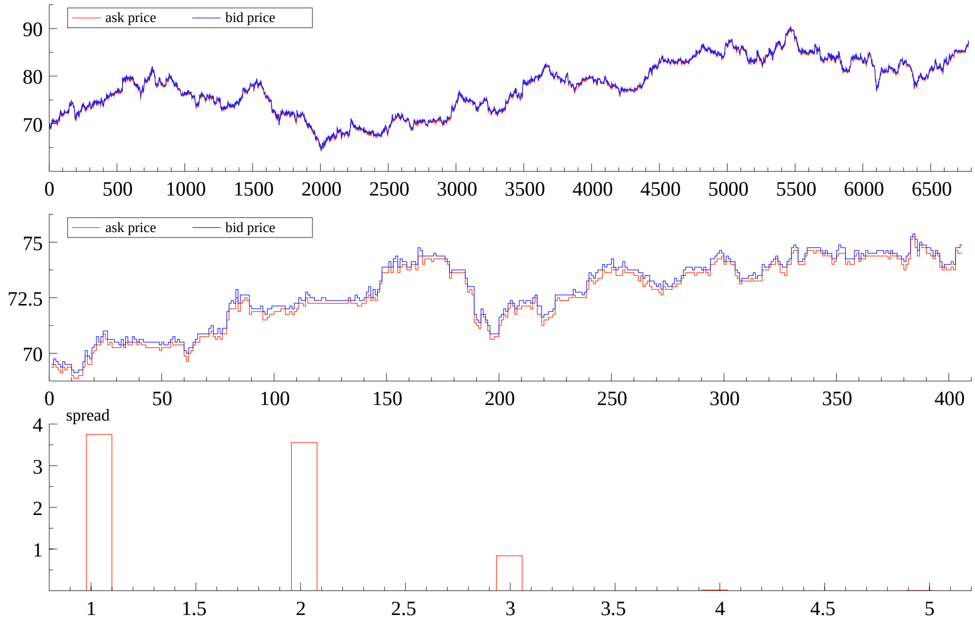

Figure 1: *Alcoa data for all trading days in 1994. There are* 6,780 *observations, typically at the end of fifteen minutes intervals. Top graph: bid and ask quotes. Middle graph: bid and ask quotes for the first* 15 *trading days. Bottom graph: histograms of the bid/ask spread. The middle graph is simply an enlargement of the first part of the top graph, containing the first* 405 *osbervations, to illustrate how the bid and ask quotes evolve over time.*

simulated ask quotes are between $50 3/8 and $74 1/4. The fact that these values are not in the same intervals as those for the market data does not matter for inference purposes.

Summary statistics for the bid/ask spread, (in ticks), both for the Alcoa series and for the simulated series are reported in Table 1. Figure 1 plots the bid and ask quotes for the whole year (top graph), the bid and ask quotes for the first 15 days, i.e., 405 observations (middle graph) and the spreads (bottom graph) for the Alcoa data. Figure 2 plots the efficient price, bid and ask quotes for the whole period, bid and ask quotes for the first 189 observations and the spreads for the simulated data. There is clearly variation in the spread for both the Alcoa data and the simulated data. The standard deviation of the spread for the simulated data is larger than that for the market data. A weakness of our model is that the frequency distribution of spreads of more than 4 ticks for the simulated data, is much larger than that for the market data. Moreover, it occasionally produces very big spreads compared to the observed spreads for the market data.

The first differences of the bid and ask quotes for Alcoa data are plotted in Figure 3. The maximum of both series is 19, the minimum is equal to -15 for the bid quote and to -16 for the ask quote. There is no change on the bid quote in 37.9% of the observations, the change (in absolute value) is 1 tick in 36.2% of the cases and is larger than 1 tick for the remaining 25.6% of the observations. The corresponding numbers for the first differences of the ask quotes are 38%, 36.2% and 25.8%. In short, both series are practically identical. It is therefore legitimate to start by assuming that the bid and ask exposure costs processes follow the same dynamics. A more detailed analysis of these series is given in section 5.

Our model does not replicate these series. Instead, the maximum of the first differences of

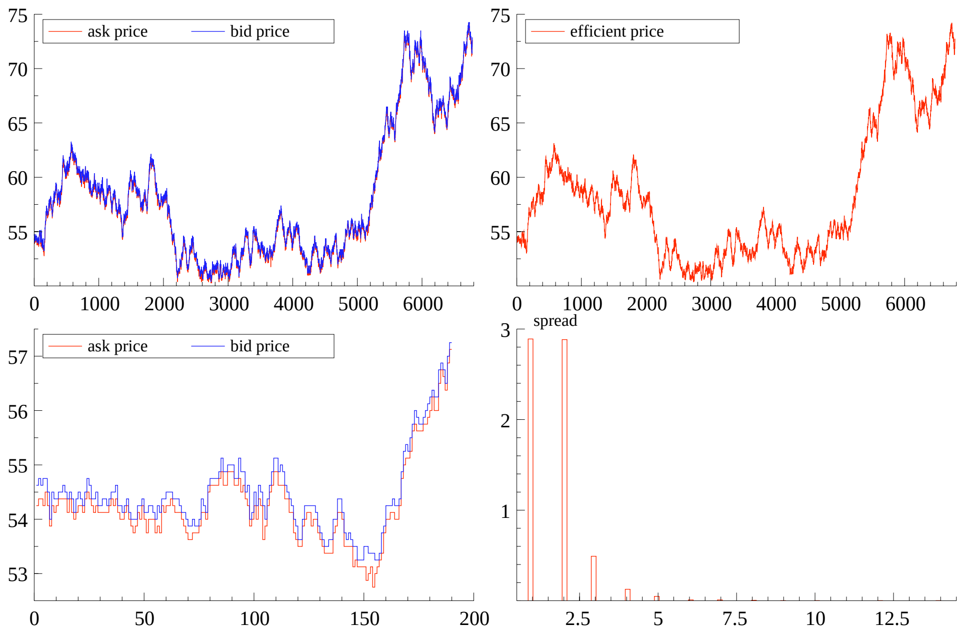

Figure 2: $6,780$ *simulated data points using* $\sigma_\varepsilon = 0.00316$, $\mu = -3.715$, $\phi = 0.4$, *and* $\sigma_\nu = 1.025$. *Top graphs: bid and ask quotes; efficient price series. Bottom graphs: bid and ask quotes for the first* 189 *observations; histograms of the bid/ask spread. The bottom left graph is simply an enlargement of the first part of the top left graph to illustrate how the bid and ask quotes evolve over time.*

the bid and ask quotes is 9, while the minimum is -7 for the bid series and -11 for the ask series. There is no change on the bid quote in 25.1% of the cases; the change (in absolute value) is 1 tick, (larger than 1 tick), in 40.4%, (34.5%), of the observations. The percentages for the ask quotes series are almost the same.

4.2 Bayesian inference

This section deals with Bayesian inference of the model defined by (13). We use the MCMC sampler outlined in section 3.2. Section 4.2.1 deals with sampling from $s_t|s_{\backslash t}, y_t; \theta$ in step (2) of the sampler. Section 4.2.2 details how to sample from the posterior distribution $\theta|s$, in step (3). Section 4.2.3 looks at a Monte Carlo experiment using a single data set. Section 4.2.4 presents the results of 100 Monte Carlo replications.

To start the scan sampler, we need to set the initial values of the logarithms of the efficient price, bid and ask costs. To do so, we take the efficient price equal to the middle point between the bid and ask quotes, the bid cost equal to the efficient price less the bid price, and the ask cost equal to the ask price less the efficient price. That is, $m_t = (1/2)(b_t + a_t)$, $\beta_t = m_t - b_t$, $\alpha_t = a_t - m_t$, $t = 1, ..., n$. Note that the bid and ask costs are then just the spread divided by two, $\beta_t = \alpha_t = (a_t - b_t)/2$. A different set of initial values with $m_t \sim U[b_t, a_t]$ was also considered. In any case, the burn in period does not need to be longer than 250 iterations.

The initial values for the parameters in the bid and ask costs equations are drawn from uniform distributions on $[-4.215, -3.215]$ for μ, on $[0.525, 1.025]$ for σ_ν, and on $[-0.2, 1]$ for ϕ. The sampler is practically independent of these values. However, it appears to be very sensitive to meaningless choices for the variance of the efficient price. As an initial value for this variance,

	Alcoa data	Simulated data
n	6,780	6,780
Min	1	1
Max	5	14
Mean	1.65	1.71
St dev	.674	.871
Distribution		
1 tick	45.9%	44.7%
2 ticks	43.5%	44.6%
3 ticks	10.3%	7.6%
4 or more ticks	0.3%	3.1%

Table 1: Descriptive statistics for bid/ask spreads, (in 1/8 ticks), for Alcoa for all trading days in 1994 and for the simulated data using $\sigma_\varepsilon = 0.00316$, $\mu = -3.715$, $\phi = 0.4$, and $\sigma_\nu = 1.025$. The frequency distributions of the spread correspond to the histograms in Figure 1 and Figure 2.

we take the variance of the first differences of the series given by $(1/2)(b_t + a_t)$.

4.2.1 Sampling from $s_t|s_{\backslash t}, y_t; \theta$

The assumption of state variables evolving independently facilitates the sampling from the density of $s_t|s_{\backslash t}, y_t; \theta$, where θ is fixed. Writing $s_{1t}|s_{\backslash 1t} \sim N(\gamma_{1t}, \Sigma_{1t})$, $s_{2t}|s_{\backslash 2t} \sim N(\gamma_{2t}, \Sigma_{2t})$, and $s_{3t}|s_{\backslash 3t} \sim N(\gamma_{3t}, \Sigma_{3t})$, sampling from this density will involve sampling from three truncatedunivariate distributions

$$\begin{array}{lcl} s_{1t}|s_{\backslash 1t}, y_t & \sim & TN_{s_{1t}<[\log(b_t),\log(a_t)]}(\gamma_{1t}, \Sigma_{1t}), \\ s_{2t}|s_{\backslash 2t}, y_t & \sim & TN_{s_{2t}>\log(a_t-b_t)}(\gamma_{2t}, \Sigma_{2t}), \\ s_{3t}|s_{\backslash 3t}, y_t & \sim & TN_{s_{3t}>\log(a_t-b_t)}(\gamma_{3t}, \Sigma_{3t}), \end{array} \tag{15}$$

subject to the constraints in (12) using the Metropolis-Hastings algorithm described above.

However, due to the conditional independence structure of the model considered, we can draw s_t, $t = 1, ..., n$, given the values of the latent variables for the previous and next period, this period's observation and some parameter estimates, within a Metropolis-Hastings algorithm. Our target density will therefore be $s_t|s_{t-1}, s_{t+1}, y_t; \theta$. Although the general scan sampler of de Jong (1997) is typically better, there is not much difference with respect to convergence in this particular setup.

In our problem, we have

$$\begin{array}{lcll} s_{11}|s_{12}; \theta & \sim & N\left(s_{12}, \sigma_\varepsilon^2\right), & \\ s_{1t}|s_{1t-1}, s_{1t+1}; \theta & \sim & N\left\{\frac{1}{2}\left(s_{1t-1} + s_{1t+1}\right), \sigma_\varepsilon^2/2\right\}, & t = 2, ..., n-1, \\ s_{1n}|s_{1n-1}; \theta & \sim & N\left(s_{1n-1}, \sigma_\varepsilon^2\right), & \end{array} \tag{16}$$

and, for $i = 2, 3$, since $s_{i0} = 0$,

$$\begin{array}{lcll} s_{it}|s_{it-1}, s_{it+1}; \theta & \sim & N\left(s_{it}^*, v_s^2\right), & t = 1, ..., n-1, \\ s_{in}|s_{in-1}; \theta & \sim & N\left\{\mu + \phi\left(s_{in-1} - \mu\right), \sigma_\nu^2\right\}, & \end{array} \tag{17}$$

where

$$s_{it}^* = \mu + \frac{\phi\left\{(s_{it-1} - \mu) + (s_{it+1} - \mu)\right\}}{(1 + \phi^2)} \qquad \text{and } v_s^2 = \frac{\sigma_\nu^2}{(1 + \phi^2)}, \tag{18}$$

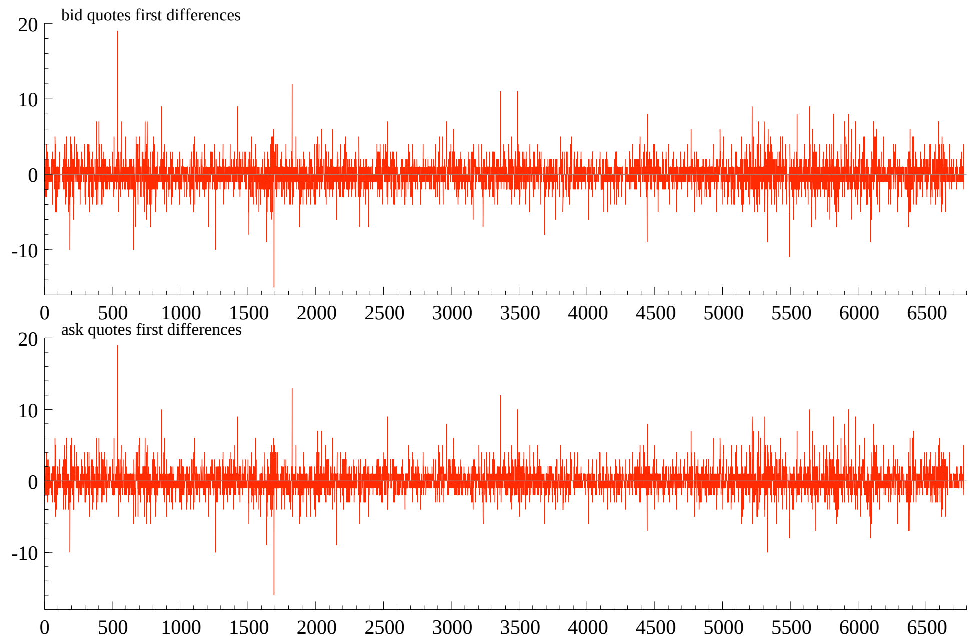

Figure 3: *First differences, (in* $1/8$ *ticks), of the bid and ask quotes for Alcoa for all trading days in 1994. There are* $6,780$ *observations.*

(see e.g. Kim, Shephard, and Chib (1997)). Note that a different end condition for $s_{i1}|s_{i2};\theta$, $i=2,3$, should be added if we do not take $s_{i0}=0$. It follows that $s_{it}|s_{it-1},s_{it+1},y_t;\theta$, $i=1,2,3$ can be drawn from univariate truncated distributions $TN_{s_{1t}<[\log(b_t),\log(a_t)]}$, $TN_{s_{2t}>\log(a_t-b_t)}$, $TN_{s_{3t}>\log(a_t-b_t)}$, using a proposal q, proportional to the target density and accepting with probability (11). We draw from truncated normal distributions without rejection using the probability integral transform theorem of Feller (1971), (e.g., Devroye (1986, p. 39)). The simulations are not continuous with respect to the model parameters as the acceptance probability in the Metropolis-Hastings algorithm is less than 1.

4.2.2 Sampling from $\theta|s$

In this framework $\theta=(\sigma_\varepsilon^2,\mu,\sigma_\nu^2,\phi)'$. Given the structure of this model, the parameter θ can be partitioned into $\theta=(\theta_1,\theta_2')'$ where θ_1 denotes the unknowns in the efficient price equation and θ_2 denotes the unknowns in the cost price equations, that is, $\theta_1=\sigma_\varepsilon^2$ and $\theta_2=(\mu,\sigma_\nu^2,\phi)'$. Thus, step (3) of the Markov chain Monte Carlo sampler for Bayesian estimation can be divided into two parts.

- Sample θ_1 from $\theta_1|s_{1.}$ where $s_{1.}=(s_{11},...,s_{1n})'$ is the $n\times 1$ vector of the logarithms of the efficient prices.
- Sample θ_2 from $\theta_2|s_{2.},s_{3.}$ where $s_{2.}=(s_{21},...,s_{2n})'$ is the $n\times 1$ vector of the logarithms of the bid prices and $s_{3.}=(s_{31},...,s_{3n})'$ is the $n\times 1$ vector of the logarithms of the ask prices.

When we update the parameters we use the following conditional structure:

1. $\sigma_\varepsilon^2|s_{1.}$,

2. $\mu | s_{2.}, s_{3.}, \sigma_\nu^2, \phi$,

3. $\sigma_\nu^2 | s_{2.}, s_{3.}, \mu, \phi$,

4. $\phi | s_{2.}, s_{3.}, \mu, \sigma_\nu^2$.

All the parameters except for ϕ have straightforward conjugate distributions which we use.

Sampling σ_ε^2
We use a non-informative prior distribution for σ_ε^2, $\pi\left(\sigma_\varepsilon^2\right) \propto \sigma_\varepsilon^{-2}$. Then σ_ε^2 is sampled from

$$\sigma_\varepsilon^2 | s_{1.} \sim \chi_{n-1}^{-2} \sum_{t=2}^{n} \left(s_{1t} - s_{1t-1}\right)^2 . \tag{19}$$

Sampling μ, σ_ν^2 and ϕ
We use a non-informative prior distribution for μ, $\pi(\mu) \propto c$. This yields the posterior conditional density

$$\mu | s_{2.}, s_{3.}, \sigma_\nu^2, \phi \sim N\left[\widehat{\mu}, Var\left(\widehat{\mu}\right)\right], \tag{20}$$

where

$$\widehat{\mu} = \frac{1}{2(n-1)(1-\phi)} \left\{ \sum_{t=2}^{n} \left(s_{2t} - \phi s_{2t-1}\right) + \sum_{t=2}^{n} \left(s_{3t} - \phi s_{3t-1}\right) \right\}, \tag{21}$$

and

$$Var\left(\widehat{\mu}\right) = \frac{\sigma_\nu^2}{2(n-1)(1-\phi)^2}. \tag{22}$$

We use $\chi_q^{-2} L_1$ for σ_ν^2, with $q = 5$ and $L_1 = q$. Then σ_2^2 is sampled from

$$\sigma_\nu^2 | s_{2.}, s_{3.}, \mu, \phi \sim \chi_{n-1+q}^{-2} \left[\sum_{t=2}^{n} \left\{ \left(s_{2t} - \mu\right) - \phi\left(s_{2t-1} - \mu\right) \right\}^2 + \sum_{t=2}^{n} \left\{ \left(s_{3t} - \mu\right) - \phi\left(s_{3t-1} - \mu\right) \right\}^2 + L_1 \right]. \tag{23}$$

Following Kim, Shephard, and Chib (1997) we use $2Beta(\zeta_1, \zeta_2) - 1$ as a prior family for ϕ to enforce the stationarity condition. Although a flat prior $\pi(\phi)$ may be attractive as it leads to an analytically tractable conditional density, the stationarity condition would not be achieved. This implies $E(\phi) = \{2\zeta_1/(\zeta_1 + \zeta_2)\} - 1$ and $Var(\phi) = 4\zeta_1\zeta_2 \Big/ \left[(\zeta_1 + \zeta_2)^2 (\zeta_1 + \zeta_2 + 1) \right]$. Alternative priors (restricted for the stationary region) for autoregressive models are discussed in Marriott and Smith (1992). We take $\zeta_1 = 10$, $\zeta_2 = 2$ so that ϕ has a prior mean of 0.66 and a standard deviation of 0.207.

To sample ϕ from the posterior distribution, we use an accept/reject algorithm as in Shephard and Pitt (1997) and Kim, Shephard, and Chib (1997), although a more general procedure such as Chib and Greenberg (1994), based on the Metropolis-Hastings algorithm, can also be used. By Bayes theorem, the posterior of $\phi | s_{2.}, s_{3.}, \mu, \sigma_\nu^2$ is proportional to the product of the likelihood function $f\left(s_{2.}, s_{3.} | \mu, \sigma_\nu^2, \phi\right)$, which is quadratic in ϕ, and the prior density, which is not quadratic in ϕ. The log-likelihood function can be written as

$$\begin{aligned} \log f\left(s_{2.}, s_{3.} | \mu, \sigma_\nu^2, \phi\right) &= \text{constant} - \tfrac{1}{2\sigma_\nu^2} \textstyle\sum_{t=2}^{n} \left\{ \left(s_{2t} - \mu\right) - \phi\left(s_{2t-1} - \mu\right) \right\}^2 \\ &\quad - \tfrac{1}{2\sigma_\nu^2} \textstyle\sum_{t=2}^{n} \left\{ \left(s_{3t} - \mu\right) - \phi\left(s_{3t-1} - \mu\right) \right\}^2 + \log f\left(s_{21}, s_{31} | s_{20}, s_{30}, \mu, \sigma_\nu^2, \phi\right). \end{aligned} \tag{24}$$

The log of the prior density is first order Taylor expanded around some point $\overline{\phi}$, (we take $\overline{\phi} = 0.55$ in our calculations). This is combined with (24) in order to form a Gaussian covering density, say $N\left\{\widehat{\phi}, Var\left(\widehat{\phi}\right)\right\}$. Then ϕ can be sampled using an accept/reject algorithm, i.e., drawing a proposal ϕ^* from this Gaussian distribution and accept it, (if in the stationary region), with a probability which depends on ζ_1, ζ_2, ϕ^* and $\overline{\phi}$. If ϕ^* is rejected, we propose a new value.

4.2.3 Illustration

We first look at a Monte Carlo experiment on a single data set. The efficient price, bid and ask prices, and spread series are the same as in Figure 2.

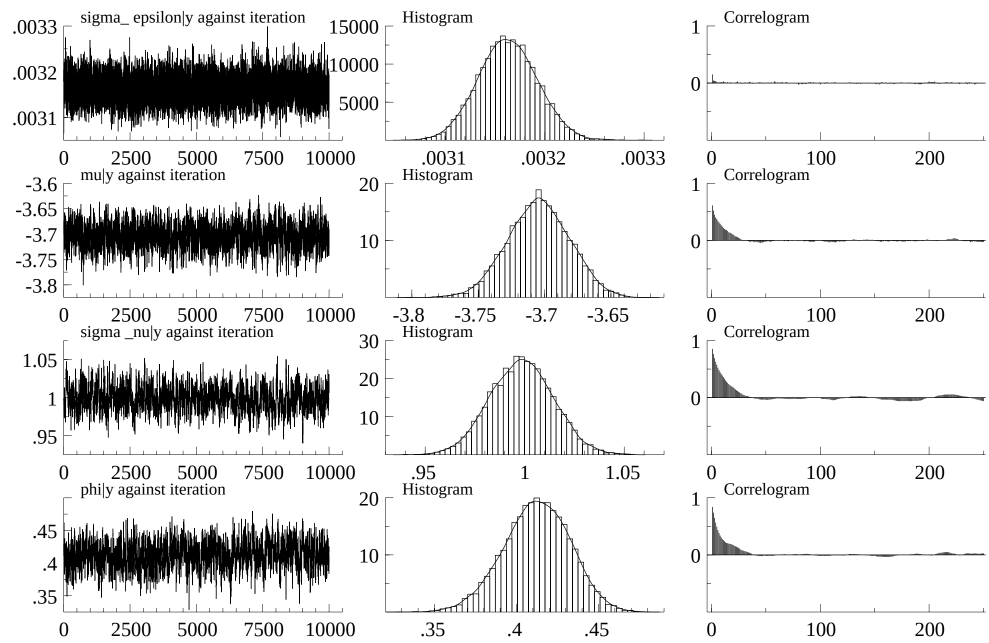

Figure 4: *Metropolis-Hastings sampler. Left graphs: the simulated values of θ against iteration number. Middle graphs: histograms of the resulting marginal distributions and estimated densities using a non-parametric density estimator. Right graphs: the corresponding correlograms for the iterations. In total* $10,000$ *iterations were drawn, discarding the first* 250. *The true parameter values are* $\sigma_\varepsilon = 0.00316$, $\mu = -3.715$, $\phi = 0.4$, *and* $\sigma_\nu = 1.025$.

Figure 4 and Table 2 present the results. We consider a burn in period of 250 iterations. The next $10,000$ iterations are recorded.

The summary statistics of Table 2 report inefficiency factors of the sampler. These are estimated as the variance of the sample mean from the MCMC sampling scheme relative to the variance of a hypothetical sampler which draws independent random variables from the posterior, (the posterior variance divided by the number of iterations). The variance of the sample mean from the MCMC sampler is estimated using a Parzen kernel (see Priestley (1981, Ch. 6)) to account for the serial correlation in the draws. Let $\widehat{\tau}_N$ denote this estimator. In our computations we follow Andrews (1991, pp. 849). The suggested $\widehat{\tau}_N$, using N samples of the chain, is then

$$\widehat{\tau}_N = 1 + \frac{2N}{N-1}\sum_{i=1}^{B_N} K\left(\frac{i}{B_N}\right)\widehat{\rho}(i), \tag{25}$$

where $\widehat{\rho}(i)$ is the estimate of the autocorrelation function of the chain, and the kernel is

$$\begin{aligned} K(x) &= 1 - 6x^2 + 6x^3, & x \in [0, \tfrac{1}{2}], \\ &= 2(1-x)^3, & x \in [\tfrac{1}{2}, 1], \\ &= 0, & \text{elsewhere}. \end{aligned} \tag{26}$$

The term B_N, the bandwidth, is selected empirically and will be stated in all calculations. In most cases, B_N will be selected based on an assessment of the correlogram.

	TRUTH	Mean	MCse	Ineff	Covariance and *Correlation*			
$\sigma_\epsilon\|y$	.003162	.003163	4.668e-007	2.5	8.722e-010	*.029*	*-.071*	*.053*
$\mu\|y$	-3.715	-3.705	.0007625	5.4	2.023e-008	.0005415	*-.517*	*-.155*
$\sigma_\nu\|y$	1.0247	.9986	.0006604	17.5	-3.323e-008	-.0001901	.0002498	*-.337*
$\phi\|y$	.4	.4123	.0008365	17.4	3.128e-008	-7.234e-005	-.0001067	.0004012

Table 2: Summaries of Figure 4. MCse denotes the Monte Carlo standard error of the simulation estimator of the mean of the posterior density. Throughout, these standard errors are computed using 250 lags and 10,000 iterations. Numbers in italics are correlations rather than covariances. Ineff denotes the estimated inefficiency factor.

For example, the inefficiency factor for θ_i, $i = 1, ..., 4$, where θ_i denotes the i-th component of $\theta = (\sigma_\varepsilon, \mu, \sigma_\nu, \phi)'$, is $\{MCse\,(\theta_i)\}^2 / \{Var\,(\theta_i|y)\,/N\}$, where $MCse\,(\theta_i)$ is computed using (25) and $Var\,(\theta_i|y)$ denotes the posterior variance of θ_i given y.

The inefficiency factor can be a useful diagnostic, although not the only one in measuring how well the chain mixes. If we require the Monte Carlo error in estimating the mean of the posterior to be no more than one percentage of the variation of the error due to the data, i.e., $\{MCse\,(\theta_i)\}^2 / Var\,(\theta_i|y) \leq 0.01$, then inefficiency factors of less than 18 for all the parameters suggest that these models can be estimated quite precisely with 1,800 iterations of the Markov chain Monte Carlo sampler. This number will give an idea of how many iterations should be used at the sampling experiment.

Results also suggest that, after the burn in period, at each iteration the simulated parameters are quite close to their true values. There is quite a strong negative correlation between μ and σ_ν, $(-.517)$ and between σ_ν and ϕ, $(-.337)$.

4.2.4 Sampling behaviour

The experiment is now repeated 100 times to study the sampling behaviour of the Bayesian estimators. We take a burn in period of 100 iterations and record the next 1,000 iterations. Results are reported in Figure 5 and Table 3. The histograms of the estimates of μ, σ_ν, and ϕ are broadly symmetric, while the histogram of σ_ε is slightly skewed to the left. All the parameter estimates are within small intervals of their true values. The means of the estimates are very close to the true parameters; the estimated standard errors are quite small, especially for σ_ε.

	σ_ϵ	μ	σ_ν	ϕ
mean	.0031882	-3.7175	1.0292	0.39326
st. deviation	(4.273e-005)	(.0242)	(.0180)	(.0211)
TRUTH	.003162	-3.715	1.0247	.4

Table 3: Summaries of Figure 5. Sampling behaviour of the Bayesian estimators using 1,000 iterations of the Metropolis-Hastings sampler. We perform 100 Monte Carlo replications. Main figures are the means of the estimates; figures in brackets are the estimated standard deviations of the replications.

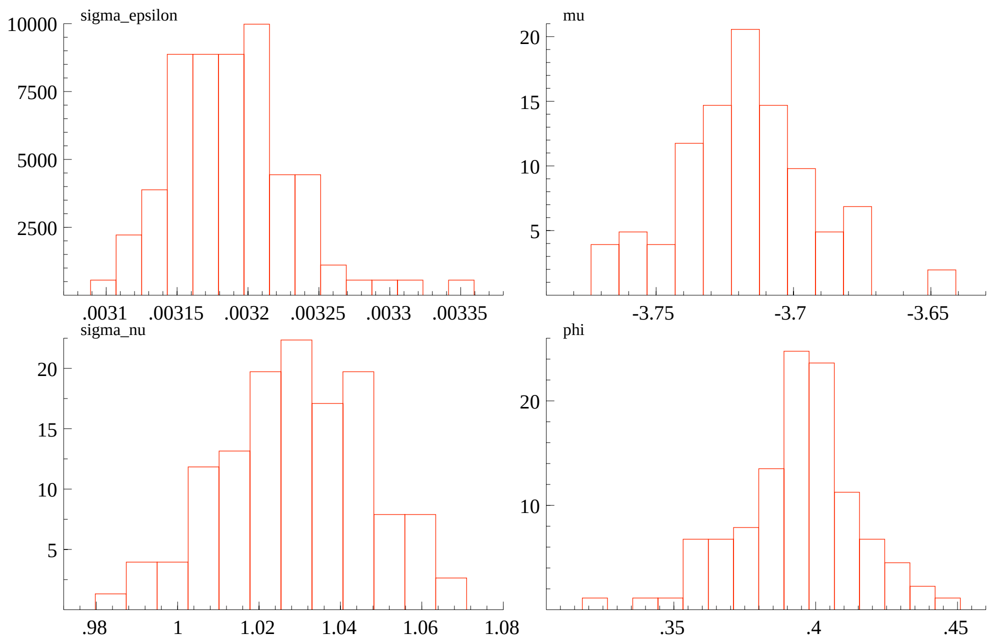

Figure 5: *Histograms of the Bayesian estimates using* 1000 *iterations of the Metropolis-Hastings sampler after a burn in period of* 100 *iterations. We perform* 100 *Monte Carlo replications. The true parameters are* .003162*,* -3.715*,* 1.0247*, and* 4.

5 Data estimation

We analyze the NYSE bid and ask quotes for Alcoa for all trading days in 1994. This data has previously been analyzed in Hasbrouck (1996).

The first differences of the bid and ask quotes for intraday and overnight observations are plotted in Figure 6. The changes in the quotes are given in ticks. The bid and ask series are again virtually identical. Summary statistics are given in Table 4. The frequency distributions of the changes in the quotes for intraday and overnight observations are clearly different. For overnight observations the size of the change, (in absolute value), is larger than 1 tick in more than 50% of the cases, while this happens only in 24.6% of the intraday observations. Moreover, the changes overnight are occasionally very large for both quotes. This is reflecting the information accumulated when the market is closed. Our model ignores this feature of the market data. It would be more realistic to allow for larger variation in the efficient price at the beginning of the day. We incorporate this feature at the end of next section.

The initial values for the efficient price, and the bid and ask quotes the same as in our previous Monte Carlo experiment, i.e., $m_t = (1/2)(b_t + a_t)$, $\beta_t = m_t - b_t$ and $\alpha_t = a_t - m_t$.

We take the variance of the first differences of the series given by $(1/2)(b_t + a_t)$ as an initial value for σ_ε^2, Hasbrouck's (1996) estimate for ϕ, twice the mean of $\{\log(\alpha_t) - \phi\}/(1-\phi)$, for μ, and twice the variance of the first differences of $\{\log(\alpha_t) - \mu\}$ for σ_ν^2, (note that the bid and ask prices are initialized at the same value and so $\alpha_t = \beta_t$,$t = 1, ..., n$). This gives initial values of $(\sigma_\varepsilon, \mu, \sigma_\nu, \phi)' = (.002636, -4.7077, .57744, .37)'$.

We use $\chi_p^{-2} L_0$ with $p = 3$ and $L_0 = 0.01 \times p$ as the prior distribution for σ_ε^2. The mean estimates for σ_ε appear to be slightly sensitive to the choice of prior, (e.g. a prior with $p = 2$ yields a mean estimate for $\sigma_\varepsilon$of .003117). However, there is almost no difference in the estimates for σ_ε^2, (.000011269or .000009716).

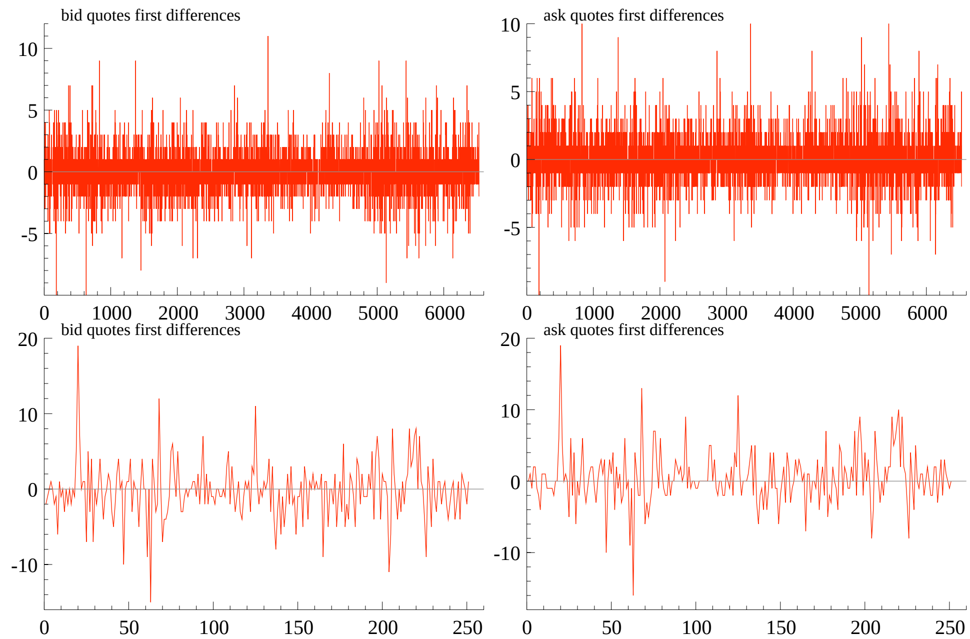

Figure 6: *First differences (in* $1/8$ *ticks) of the bid and ask quotes for Alcoa for all trading days in 1994. Top graphs: intraday observations. Bottom graphs: overnight observations. There are* $6,780$ *observations,* $6,528$ *intraday observations and* 251 *overnight observations from* 252 *trading days.*

Results are presented in Figure 7 and Table 5. We consider a burn in period of 750 iterations. The next $3,000$ iterations are recorded. Inefficiency factor of .6 for σ_ε indicates negative autocorrelation. Actually, the correlogram reflects very small negative autocorrelation. Inefficiency factors of less than 22 suggest that the model can be estimated reasonably precisely with $2,200$ iterations of the sampler. The Monte Carlo standard errors (except for σ_ε) are larger than those in Table 2. Note that they are now computed using 500 lags instead of 250. The posterior correlations between μ, σ_ν and ϕ are of the same magnitude as those in Table 2. As in Hasbrouck (1996) both the disturbance standard deviation σ_ν and the autoregression coefficient ϕ are strongly positive. His estimates for these parameters are $.37(.03)$ and $.86(.03)$, (standard errors in brackets). Despite the simplicity of the model specified here, we get similar estimates. The autoregression coefficient suggests that 38.1% of the excess log cost persists at the following period, (fifteen minutes later).

To conclude this section, we simulate $6,780$ data points setting the parameters equal to their estimates, with $\log(m_0) = 4.25$. This gives bid and ask quotes, (in dollars), within $[62\ 3/8, 90]$ and $[62\ 1/2, 89\ 3/4]$ respectively. Figure 8 plots the bid and ask quotes for the whole year (top graph), the bid and ask quotes for the first 15 days, i.e., 405 observations (middle graph) and the spreads (bottom graph) for the simulated data using the parameter estimates. We still occasionally observe very large spreads (e.g., in Figure 8, the spread is 7 ticks at the 270-th observation). The spread mean (1.67), and standard deviation (723), are closer to the values for the Alcoa data than before. The distribution frequency of the spread is as follows. There is a spread of 1, 2, 3 and 4 or more ticks in 42.3%, 47.1%, 6.1% and 1.7% of the observations, respectively. The percentage of 4 or more ticks spreads is still quite large (2.9%) compared to only .3% for the Alcoa series. The maximum change in the bid (ask) quote is now 9(8) and the

	Intraday		Overnight	
	bid changes	ask changes	bid changes	ask changes
n	6,528	6,528	251	251
Min	-10	-10	-15	-16
Max	11	10	19	19
Mean	0.03	0.002	-0.21	0.5
St dev	1.58	1.57	3.7	3.7
Distribution of the absolute values				
no change	38.6%	38.7%	18.3%	18.7%
1 tick change	36.8%	36.7%	27.5%	24.3%
more than 1 tick change	24.6%	24.6%	54.2%	57%

Table 4: Descriptive statistics for the first differences, (in 1/8 ticks), of the bid and ask quotes for Alcoa for all trading days in 1994. There are 6,780 observations, 6,528 intraday observations and 251 overnight observations (252 trading days).

	Mean	MCse	Ineff	Covariance and *Correlation*			
$\sigma_\epsilon \| y$	.003357	4.2147e-007	.6	9.0994e-010	*-.012*	*.004*	*.003*
$\mu \| y$	-3.6478	.001177	8.3	-8.3554e-009	.0004998	*-.593*	*-.130*
$\sigma_\nu \| y$	.90746	.0014245	21.6	2.1024e-009	-.0002226	.00028234	*-.371*
$\phi \| y$	.38107	.001862	16.6	2.4342e-009	-7.2913e-005	-0.000156	.00062523

Table 5: Summaries of Figure 7. Mean estimates of the bid and ask quotes for Alcoa for all trading days in 1994. MCse denotes the Monte Carlo standard error of the simulation estimator of the mean of the posterior density. Throughout, these standard errors are computed using 500 lags and 3,000 iterations. Numbers in italics are correlations rather than covariances. Ineff denotes the estimated inefficiency factor.

minimum change is $-9(-8)$. There is no change in the bid quote in 19.4% of the observations, the change (in absolute value) is 1 tick in 35.2% of the cases and is greater than 1 tick in the remaining 45.4%. Again this is quite different from the frequency distribution exhibited by the data. In the next section we analyze more realistic specifications which may account for these weaknesses in the fitted model.

We also look at the autocorrelation function of the first difference of the log ask and log bid quotes for both the Alcoa series and the simulated series using the parameter estimates. This is shown in Figure 9 and does not appear not to be very informative. There is a very small but persistent correlation in both series.

6 Extensions

6.1 Alternative specifications of the model

Our model is very limited as it does not take into account many intricate features of the data. We discuss here some possible extensions of the basic formulation in (13), review the specification considered in Hasbrouck (1996) and then study a slightly more general version of the model in (13).

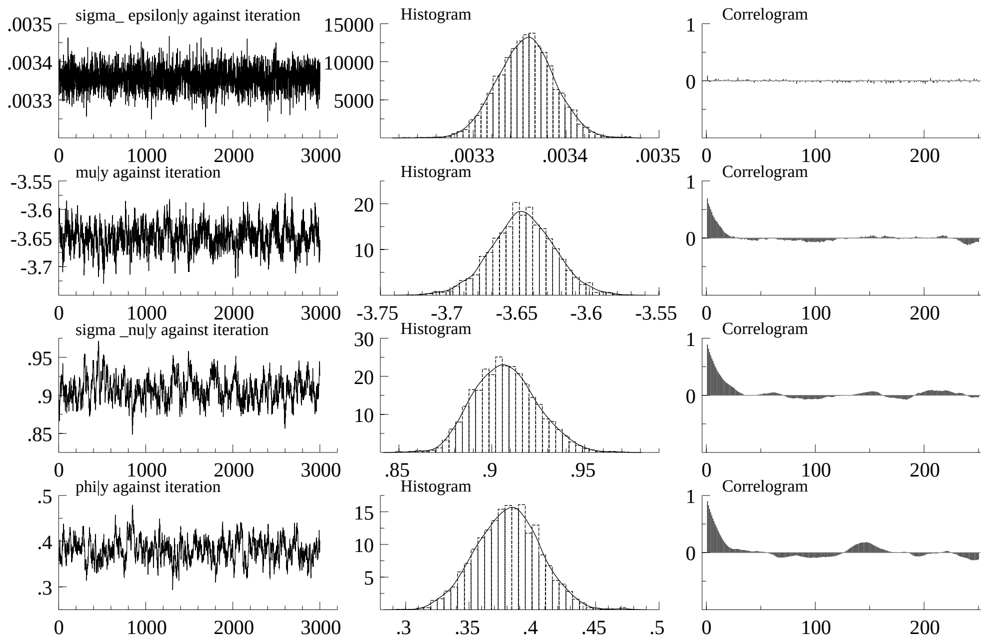

Figure 7: *Metropolis-Hastings sampler using the bid and ask quotes series for Alcoa in 1994. Left graphs: the simulated values of θ against iteration number. Middle graphs: histograms of the resulting marginal distributions and estimated densities using a non-parametric density estimator. Right graphs: the corresponding correlograms for the iterations. In total* $3,000$ *iterations were drawn, discarding the first* 750.

To allow for the intraday 'U' pattern frequently exhibited by market data, (where observations are typically higher at the start and end of each day), we could assume that μ in the bid and ask costs equations varies according with the time of day the observation takes place. Hasbrouck (1996) suggests using exponential decay functions of the type

$$\mu_t = k_1 + k_2^{open} \exp\left(-k_3^{open} \tau_t^{open}\right) + k_2^{close} \exp\left(-k_3^{close} \tau_t^{close}\right), \tag{27}$$

where τ_t^{open} is the time in hours that has passed since the market opened, and τ_t^{close} is the time remaining before the market closes, also in hours.

Asset returns tend to be leptokurtic. To allow for leptokurtosis we could relax the Gaussian assumption and suppose that the increments of the efficient price are independent and identically distributed (i.i.d.) draws from a thick tailed distribution.

We could also allow for a time varying variance in the efficient price equation, using a (parametric) autoregressive conditional heteroskedastic (ARCH) model or a stochastic volatility model. ARCH type models are surveyed in Bollerslev, Engle, and Nelson (1994, Ch. 49). Reviews of the literature on stochastic volatility models are given in Taylor (1994), Shephard (1996) and Ghysels, Harvey, and Renault (1996).

Stochastic volatility models specify a latent stochastic process for the variance. A simple stochastic volatility model for the log of the efficient price is written as

$$\begin{array}{lll} \Delta \log(m_t) & = & e^{h_t/2}\varepsilon_t, \\ h_{t+1} & = & \gamma + \varphi(h_t - \gamma) + \sigma_\eta \eta_t, \end{array} \tag{28}$$

where $h_t$is the log volatility at time t, $|\varphi| < 1$,$\varepsilon_t$and η_t are uncorrelated normal white noise

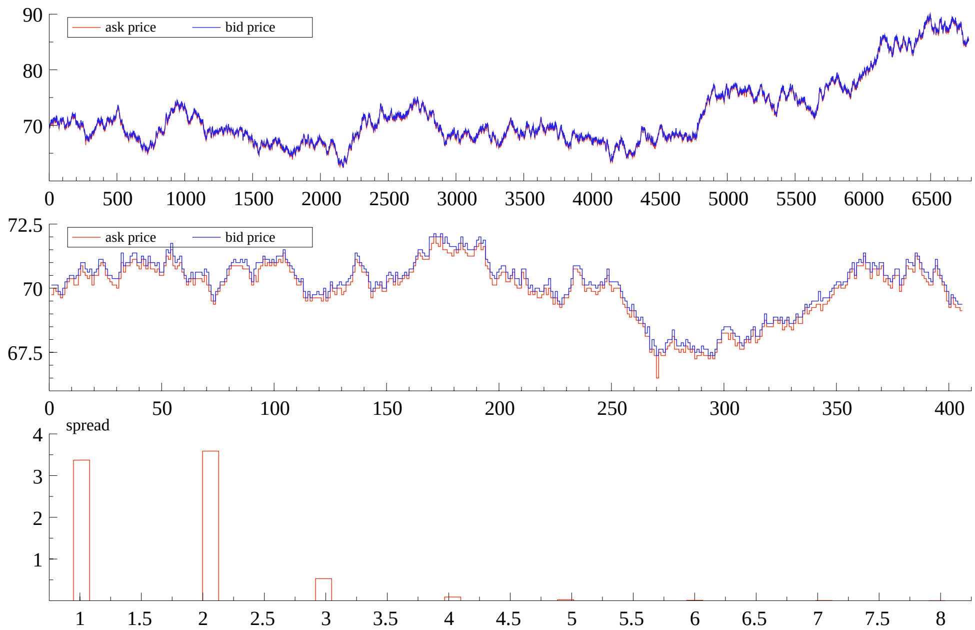

Figure 8: *Simulated data constructed with the parameter estimates which result from estimating the bid and ask quotes for Alcoa in 1994 inside a Metropolis-Hastings sampler. Top graph: bid and ask quotes. Middle graph: bid and ask quotes for the first* 405 *observations. Bottom graph: histograms of the bid/ask spread. The middle graph is simply an enlargement of the first part of the top graph to illustrate how the bid and ask quotes evolve over time.*

shocks, and $h_1$is drawn from the stationary distribution $h_1 \sim N\left\{\gamma, \sigma_\eta^2/(1-\varphi^2)\right\}$. We can think of $\exp(\gamma/2)$ as the modal instantaneous volatility, φ as the persistence in the volatility, and σ_η as the volatility of the log-volatility.

Simulation based methods, in particular, Markov chain Monte Carlo methods, can be designed to analyze these models. Gibbs samplers that draw the log volatilities one at a time are given in Jacquier, Polson, and Rossi (1994) and Shephard (1993). These samplers typically converge very slowly. In a recent paper, Kim, Shephard, and Chib (1997) propose a much more efficient adapted Gibbs sampler which draws all the log volatilities at once. They use an approximated offset mixture model and then correct for the (minor) approximation error by using an importance reweighting procedure.

Alternatively, an ARCH specification can be used. In the basic ARCH(p)model formulated by Engle (1982), the conditional variance at time t, $\sigma_{\varepsilon t}^2$, is a linear function of past $q$squared observations $\{\log(m_t)\}^2$. In the generalized ARCH, GARCH(p,q) model suggested by Bollerslev (1986), $\sigma_{\varepsilon t}^2$ is a linear function of $p$past variances and $q$squared innovations. In the exponential ARCH (GARCH) model of Nelson (1991), $\log(\sigma_{\varepsilon t}^2)$ depends on the size and sign of lagged standardized residuals. Other parametric formulations have been suggested in the literature.

A GARCH$(1,1)$model for the log of the efficient price is written as

$$\log(m_t)\,|\log(m_1),...,\log(m_{t-1}) \sim \left(0, \sigma_{\varepsilon t}^2\right) \tag{29}$$

for some distribution which needs to be specified with

$$\sigma_{\varepsilon t}^2 = \varphi_0 + \varphi_1\{\log(m_{t-1})\}^2 + \varphi_2\sigma_{\varepsilon t-1}^2, \tag{30}$$

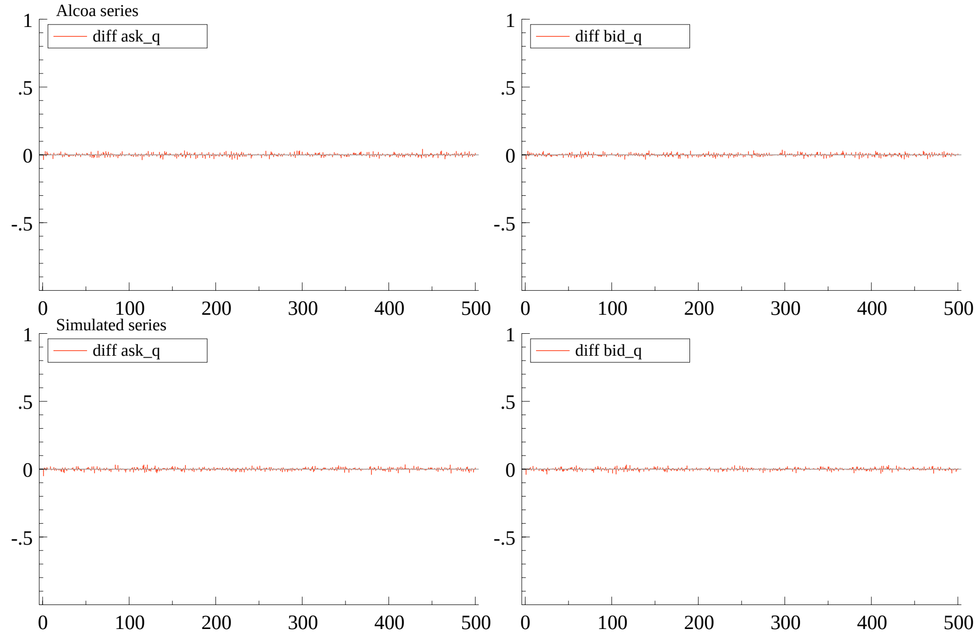

Figure 9: *Autocorrelation function of the first differences of the log ask and log bid quotes series. Top graph: data for Alcoa in 1994. Bottom graph: simulated data constructed with the parameter estimates which result from estimating the bid and ask quotes for Alcoa in 1994 inside a Metropolis-Hastings sampler.*

where $\varphi_0 \geq 0$,$\varphi_1 \geq 0$ and $\varphi_2 \geq 0$. The model is covariance stationary if $\varphi_1 + \varphi_2 < 1$. This model can be written as an ARMA$(1,1)$ model for $\{\log(m_t)\}^2$

$$\{\log(m_t)\}^2 = \varphi_0 + (\varphi f_1 + \varphi_2)\{\log(m_{t-1})\}^2 + \upsilon_t - \varphi_2 \upsilon_{t-1}, \tag{31}$$

with $\upsilon_t = \{\log(m_t)\}^2 - \sigma_{\varepsilon t}^2$. In many applications with high frequency market data, $\varphi_1 + \varphi_2$ tends to be close to 1.

The standardized innovations $\zeta_t = \varepsilon_t / \sigma_{\varepsilon t}$ are assumed to be i.i.d. If the normal distribution is used, then the unconditional distribution for $\log(m_t)$, although leptokurtic, does not capture all the leptokurtosis present in high speculative prices. Bollerslev (1987) suggests using a t_ι -Student distribution with $\iota > 2$ degrees of freedom. Nelson (1991) suggests using a generalized error distribution (GED) with tail-thickness parameter ι

$$f_{GED}(\zeta_t;\iota) = \frac{\iota \exp\left(-\frac{1}{2}\left|\frac{\zeta}{\lambda}\right|^{\iota}\right)}{\lambda 2^{(1+1/\iota)}\Gamma(1/\iota)}, \qquad \text{where } \lambda = \sqrt{\frac{2^{(-2/\iota)}\Gamma(1/\iota)}{\Gamma(3/\iota)}}. \tag{32}$$

This distribution is the standard normal for $\iota = 2$, has thicker tails than the normal for $\iota > 2$ and thinner tails than the normal for $\iota < 2$. Nelson (1991) finds that this distribution is insufficient to successfully fit data on US stock index returns due to the fact that the data has many more large standardized residuals than those postulated by a GED distribution. Instead, Bollerslev, Engle, and Nelson (1994, Ch. 49) suggest using a generalized t -distribution which nests both the Student's t and the GED distributions.

In principle, it should be straightforward to generalize the Markov chain Monte Carlo method for Bayesian estimation given in section 3.2 to the stochastic volatility case or to the GARCH$(1,1)$ case, (with normally distributed innovations for simplicity). The parameters of interest in the efficient price equation will be $\theta_1 = \left(\gamma, \varphi, \sigma_\eta^2\right)'$in the stochastic volatility formulation or $\theta_1 = (\varphi_0, \varphi_1, \varphi_2)'$in the GARCH$(1,1)$ formulation. Prior distributions as discussed in Kim, Shephard, and Chib (1997) could be used. In particular, they take a diffuse prior on γ,a $2Beta\,(\zeta_1, \zeta_2) - 1$for φ and $\chi_p^{-2}L$for σ_η^2,for some constants $p$and L.For the GARCH model, rewritten as an ARMA model, they take the same prior on $\varphi_1 + \varphi_2$as for φ,a $Beta$distribution for $\varphi_2 / (\varphi_1 + \varphi_2) |\, \varphi_1 + \varphi_2 = r_\varphi$and a diffuse inverse chi-squared distribution for $\varphi_0 / (\varphi_1 + \varphi_2) |\, \varphi_1, \varphi_2$.

Before analyzing, as a way of illustration, a version of (13) model which incorporates the effect on the efficient price of information accumulated when the market is closed, we briefly describe the model in Hasbrouck's (1996).

6.1.1 Hasbrouck's (1996) specification

The model considered in Hasbrouck (1996) is defined by three independent processes, m_t,β_t,α_t,

$$\begin{array}{lll} \log(m_t) & = & \log(m_{t-1}) + \sigma_{\varepsilon t}\varepsilon_t, \\ \log(\beta_t) - \mu_t & = & \phi\left\{\log(\beta_{t-1}) - \mu_{t-1}\right\} + \sigma_\nu \nu_{1t}, \\ \log(\alpha_t) - \mu_t & = & \phi\left\{\log(\alpha_{t-1}) - \mu_{t-1}\right\} + \sigma_\nu \nu_{2t}. \end{array} \tag{33}$$

Following Nelson (1991), he uses a GED distribution for the standardized innovations ε_t,and an EGARCH specification for the variance

$$\log\left(\sigma_{\varepsilon t}^2\right) = \eta_t + \varphi\left\{\log\left(\sigma_{\varepsilon t-1}^2\right) - \eta_{t-1}\right\} + \gamma\left(|\varepsilon_{t-1}| - E\,|\varepsilon_{t-1}|\right), \tag{34}$$

with $E\,|\varepsilon| = \lambda 2^{(1/\iota)}\Gamma(1/\iota)/\Gamma(3/\iota)$ .To avoid inclusion of $\sigma_{\varepsilon t}^2$as another state variable he replaces $|\varepsilon_{t-1}|$ with $E_{t-1}\,|\varepsilon_{t-1}| = E\left[|u_{t-1}|\,|s_0, s_1, ..., s_{t-1}\right]/\sigma_t$. The deterministic components are modelled as

$$\eta_t = \begin{cases} l_1 + l_2^{open}\exp\left(-l_3^{open}\tau_t^{open}\right) + l_2^{close}, & \text{if } t \text{ is an intraday interval} \\ \eta^{overnight}, & \text{if } t \text{ is an overnight interval} \end{cases} \tag{35}$$

for η_t,and

$$\mu_t = k_1 + k_2^{open}\exp\left(-k_3^{open}\tau_t^{open}\right) + k_2^{close}\exp\left(-k_3^{close}\tau_t^{close}\right), \tag{36}$$

for μ_t,where $\tau_t^{open}$and $\tau_t^{close}$are defined as before.

Hasbrouck (1996) uses the method described in section 3.1 to estimate the bid and ask quotes for Alcoa for all trading days in 1994. He defines the states as $s_t = \{m_t, \alpha_t, \beta_t\}'$instead of as the log variables, so that the observed bid and ask prices are simply

$$y_t = \left\{Floor\,(s_{1t} - s_{3t})\,, Ceiling\,(s_{1t} + s_{2t})\right\}'. \tag{37}$$

6.2 A slightly more general model

As a way of illustration, we consider a version of (13) model which incorporates the effect on the efficient price of information accumulated when the market is closed. In (13), we assumed that the increments to the efficient price all have the same variance, independently of the time of day at which the observation is recorded, that is, $\Delta\log(m_t) \sim N\left(0, \sigma_\varepsilon^2\right)$ for all $t = 2, ..., n$.Instead, we now assume that the variance of the difference between the prices at the first observation of each day and the price at the last observation of the previous day is larger than the variance of the difference in prices for intraday observations.

The resulting model can be specified as follows. Assume there are Rtrading days with $n_r$observations in the r-th day of trading, i.e., $\sum_{r=1}^{R} n_r = n$.The new model is written as

$$\begin{array}{ll} \Delta \log(m_t) \sim N\left(0, \sigma_{\varepsilon_1}^2\right), & \text{for } t = n_1+1, n_1+n_2+1, ..., n_1+...+n_{R-1}+1 \\ \Delta \log(m_t) \sim N\left(0, \sigma_{\varepsilon}^2\right), & \text{for any other value of } t \text{ greater than } 1 \end{array} \tag{38}$$

with $\sigma_{\varepsilon_1}^2 > \sigma_{\varepsilon}^2$.In our case, $R = 252$and in most days, (245) ,$n_r = 27$.For simplicity, let us denote $N_r = n_1+...+n_r$,$r = 1, ..., R$,with $N_1 = 0$.Note that the observations of day $r$are $N_{r-1}+1, ..., N_r$.

6.2.1 A Monte Carlo experiment

As an illustration, we carry out a Monte Carlo experiment on a single data set generated in (38), taking $\sigma_{\varepsilon_1}^2 = 2\sigma_{\varepsilon}^2$.We take $n_r$and $R$to be the same as in the data for Alcoa. In this setup, $\theta = (\sigma_{\varepsilon}^2, \sigma_{\varepsilon_1}^2, \mu, \sigma_{\nu}^2, \phi)'$and $\theta_1 = (\sigma_{\varepsilon}^2, \sigma_{\varepsilon_1}^2)'$.We take $\sigma_{\varepsilon}^2 = 0.00001$,$(\sigma_{\varepsilon} = 0.00316)$,$\sigma_{\varepsilon_1}^2 = 0.00002$,$(\sigma_{\varepsilon_1} = 0.00447)$,$\mu = -3.715$,$\sigma_{\nu}^2 = 1.05$,$(\sigma_{\nu} = 1.025)$,and $\phi = 0.4$.The bid and ask quotes, spreads and the efficient prices series, (with $\log(m_0) = 0$), are plotted in Figure 10. This model also occasionally yields very large spreads, (for example, the spread is 9ticks at the last observation of day 10). As before, the distribution frequency of 1 and 2 ticks spreads is approximately the distribution for the data. However, we get too many observations with more than 3 ticks spreads (3.1%) and too few with 3 ticks spreads (7.8%) . We are able to replicate the distribution of the changes in bid and ask quotes for overnight observations, although we get lower extreme values. We get the same extreme values for intraday observations; however, the frequencies of no change in the quotes and of more than 1 tick change are inverted. In short, it seems that more complex specifications, as discussed above, need to be considered in order to get a good fit of the model.

	TRUTH	Mean	MCse	Inefficiency	
$\sigma_{\epsilon}\|y$	.003162	.003150	4.921e-007	2.8	
$\sigma_{\varepsilon_1}\|y$	.004472	.004553	2.748e-006	1.7	
$\mu\|y$	-3.715	-3.7147	.0009077	14.6	
$\sigma_{\nu}\|y$	1.0247	1.0276	.0007813	21.5	
$\phi\|y$	.4	.38618	.0007204	12.7	
		Covariance and *Correlation*			
$\sigma_{\epsilon}\|y$	8.663e-010	*-.004*	*.016*	*-.058*	*.047*
$\sigma_{\varepsilon_1}\|y$	-2.740e-011	4.4895e-008	*.008*	*-.025*	*.009*
$\mu\|y$	1.0881e-008	4.2217e-008	.000563	*-.570*	*-.117*
$\sigma_{\nu}\|y$	-2.8732e-008	-9.105e-008	-.0002280	.0002843	*-.308*
$\phi\|y$	2.8146e-008	4.0681e-008	-5.601e-005	-.0001050	.000409

Table 6: Summaries of Figure 11. MCse denotes the Monte Carlo standard error of the simulation estimator of the mean of the posterior density. Throughout, these standard errors are computed using 500 lags and 10,000 iterations. Numbers in italics are correlations rather than covariances.

To sample $\theta_1|s_{1.}$we sample (1) $\sigma_{\varepsilon}^2|s_{1.}$,(2) $\sigma_{\varepsilon_1}^2|s_{1.}$,with

$$\sigma_{\varepsilon}^2|s_{1.} \sim \chi_{n-R}^{-2} \sum_{r=1}^{R} \sum_{t=N_{r-1}+2}^{N_r} \left(s_{1t} - s_{1t-1}\right)^2, \tag{39}$$

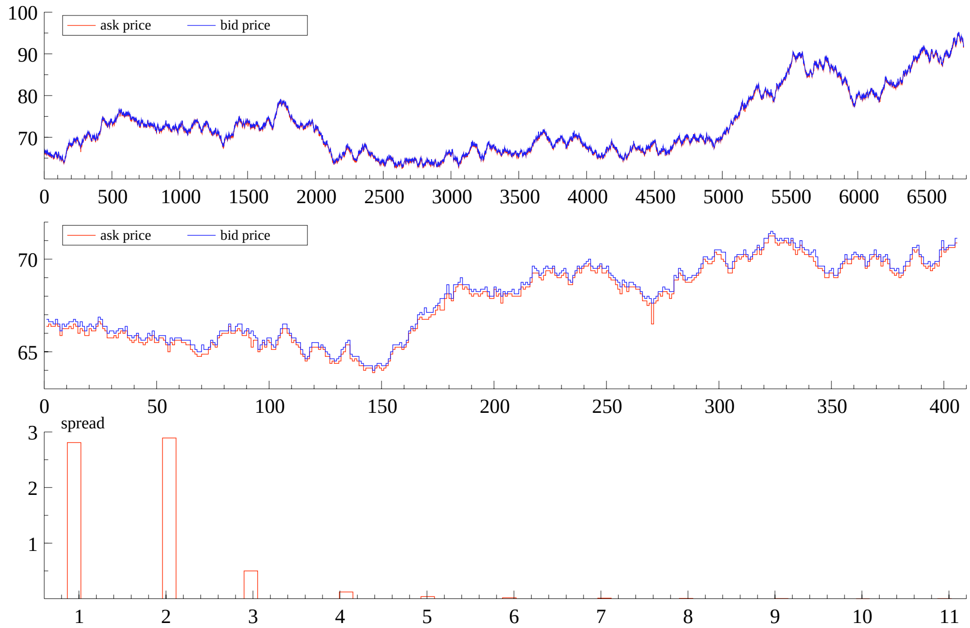

Figure 10: *Simulated data using the new specification of the model:* $\sigma_\varepsilon = 0.00316$, $\sigma_{\varepsilon_1} = 0.00447$, $\mu = -3.715$, $\sigma_\nu = 1.025$, *and* $\phi = 0.4$. *Top graph: bid and ask quotes. Middle graph: bid and ask quotes for the first* 405 *observations. Bottom graph: histograms of the bid/ask spread. The middle graph is simply an enlargement of the first part of the top graph to illustrate how the bid and ask quotes evolve over time.*

and

$$\sigma^2_{\varepsilon_1}|s_{1.} \sim \chi^{-2}_{R-1} \sum_{r=2}^{R} \left(s_{1,N_{r-1}+1} - s_{1,N_{r-1}}\right)^2. \tag{40}$$

Note that we use non-informative prior distributions for $\sigma^2_\varepsilon$and $\sigma^2_{\varepsilon_1}$.

Let $w = \sigma^2_\varepsilon/\sigma^2_{\varepsilon_1}$($w = 0.5$in our experiment). Then

$$s_{1t}|s_{1t-1}, s_{1t+1}; \theta \sim N\left\{\frac{1}{2}\left(s_{1t-1} + s_{1t+1}\right), \sigma^2_\varepsilon/2\right\} \tag{41}$$

if $t$is not the first observation or the last observation of the day,

$$s_{1N_r}|s_{1,N_r-1}, s_{1,N_r+1}; \theta \sim N\left\{\left(s_{1,N_r-1} + w \cdot s_{1,N_r+1}\right)/(1+w), \sigma^2_\varepsilon/(1+w)\right\} \tag{42}$$

if $t$is the last observation in day r,and

$$s_{1,N_r+1}|s_{1,N_r}, s_{1,N_r+2}; \theta \sim N\left[\left\{s_{1,N_r} + (1/w) \cdot s_{1,N_r+2}\right\}/\{1+(1/w)\}, \sigma^2_{\varepsilon_1}/\{1+(1/w)\}\right] \tag{43}$$

if $t$is the first observation in day $r+1$.The end conditions are as in (16).

The sampler is initialized as before, taking the initial value for $\sigma^2_{\varepsilon_1}$to be equal to that for σ^2_ε.

Results are presented in Figure 11 and Table 6. We consider a burn in period of 750iterations. The next 3,000iterations are recorded. The correlograms and posterior covariances and correlations are very similar to those obtained in section 4.2.3 using the simpler specification of the model. The histogram of the estimates for σ_ν has moved slightly towards the right while the

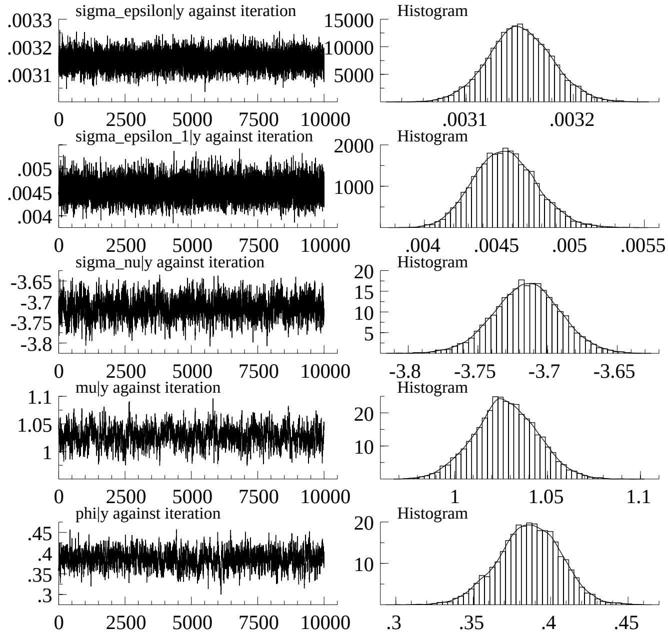

Figure 11: *Metropolis-Hastings sampler for the new specification of the model. Left graphs: simulated values of θ against iteration number. Middle graphs: histograms of the resulting marginal distributions and estimated densities using a non-parametric density estimator. Right graphs: the corresponding correlograms for the iterations. In total,* $10,000$ *iterations were drawn, discarding the first* 750. *The true parameter values are* $\sigma_\varepsilon = 0.00316$, $\sigma_{\varepsilon_1} = 0.00447$, $\mu = -3.715$, $\sigma_\nu = 1.025$, *and* $\phi = 0.4$.

histogram for ϕ has moved slightly towards the left. The inefficiency factor for μ is now larger. However, inefficiency factors of less than 22suggest that these models can be estimated quite precisely with $2,200$iterations of the sampler.

We do not report results on the sampling behaviour of the estimators because this model does not appear to better fit the data than the previous one. A better fit should be expected for models with time-varying variance and cost functions that allow for the observed intraday 'U' shapes.

7 Conclusions

In this paper we have discussed Bayesian estimation of a dynamic model of discrete bid and ask quotes suggested by Hasbrouck (1996). We assume that the observed bid and ask quotes are obtained from underlying continuous variables by some rounding function. Specifically, the bid quote is the efficient price minus the bid cost exposure rounded down to the next tick and the ask quote is the efficient price plus the ask cost exposure rounded up to the next tick. These costs reflect asymmetric information costs and fixed transaction costs. The model can be thought of as a limited dependent process. In this context, the scan sampler can be used inside a Metropolis-Hastings algorithm to sample a single s_t,given the most recent updated values of all the other latent points, the observations and some parameter estimates. We postulate a very simple model where the log of the efficient price follows a random walk and the logs of the ask cost and of the bid cost are stationary AR(1) processes. The three variables are assumed to be independent. Due to the simplicity of the model, we design a simpler MCMC method to simulate a single $s_t$given s_{t-1},s_{t+1},y_t,and a fixed parameter value, inside a Metropolis-Hastings algorithm. The simulations are not continuous with respect to the model parameters as the acceptance probability is less than 1 , and hence, we carry out a Bayesian analysis. Assuming this form for the underlying model, the Metropolis-Hastings sampler is applied to estimate the bid and ask quotes for a stock in the NYSE for all trading days in 1994. This data has previously been analyzed in Hasbrouck (1996) using non-linear filtering techniques as suggested in Kitagawa (1987). Only the estimates of the disturbance variance in the cost equations and the autoregression coefficient can be compared with the estimates in Hasbrouck (1996). Despite the simplicity of our model we get quite similar estimates. The analysis can be extended straightforwardly in different ways to more realistic models which incorporate institutional features exhibited by market data.

8 Acknowledgments

We would like to thank Joel Hasbrouck for allowing us to quote his unpublished paper, as well as providing access to his data. Financial support from the ESRC under the grant "Estimation via Simulation in Econometrics", and from the Europen Union's Human Capital and Mobility Program under the research network grant "Econometric Inference using Simulation Techniques" is gratefully acknowledged.

References

Andrews, D. W. K. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica 59*, 817–858.

Barndorff-Nielsen, O. E. and D. R. Cox (1994). *Inference and Asymptotics.* London: Chapman & Hall.

Bernhardt, D. and E. Hughson (1995). Discrete pricing and the design of dealership markets. *J. Economic Theory 71*, 148–182.

Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity. *J. Econometrics 51*, 307–327.

Bollerslev, T. (1987). A conditional heteroskedastic time series model for speculative prices and rates of return. *Rev. Economics and Statistics 69*, 542–547.

Bollerslev, T., R. F. Engle, and D. B. Nelson (1994). ARCH models. In R. F. Engle and D. McFadden (Eds.), *The Handbook of Econometrics, Volume 4*, pp. 2959–3038. North-Holland.

Bollerslev, T. and M. Melvin (1994). Bid-ask spreads in the foreign exchange market: an empirical analysis. *Journal of International Economics 36*, 355–372.

Chan, K. S. and J. Ledolter (1995). Monte Carlo EM estimation for time series models involving counts. *J. Am. Statist. Assoc. 89*, 242–252.

Chib, S. and E. Greenberg (1994). Bayes inference for regression models with ARMA(p,q) errors. *J. Econometrics 64*, 183–206.

Chordia, T. and A. Subrahmanyan (1995). Market making, the tick size, and payment-for-order flow: theory and evidence. *Journal of Business 68*, 543–575.

Cordella, T. and T. Foucault (1996). Minimum price variation, time priority and quote dynamics. Universitat Pompeu Fabra, working paper.

de Jong, P. (1997). The scan sampler for time series models. *Biometrika 84*. Forthcoming.

de Jong, P. and N. Shephard (1995). The simulation smoother for time series models. *Biometrika 82*, 339–50.

Devroye, L. (1986). *Non-Uniform Random Variate Generation.* New York: Springer-Verlag.

Doornik, J. A. (1996). *Ox: Object Oriented Matrix Programming, 1.10.* London: Chapman & Hall.

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

Feller, W. (1971). *A Course in Probability Theory and its Applications* (2 ed.), Volume 2. New York: Wiley.

Ghysels, E., A. C. Harvey, and E. Renault (1996). Stochastic volatility. In C. R. Rao and G. S. Maddala (Eds.), *Statistical Methods in Finance.* Amsterdam: North-Holland.

Gilks, W. K., S. Richardson, and D. J. Spiegelhalter (1996). *Markov Chain Monte Carlo in Practice.* London: Chapman & Hall.

Glosten, L. R. (1994). Is the electronic open limit order book inevitable? *J. Finance 49*, 1127–1161.

Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financial Economics 14*, 71–100.

Hamilton, J. (1994). *Time Series Analysis.* Princeton: Princeton University Press.

Hamilton, J. D. (1984). State space models. In Z. Griliches and M. Intriligator (Eds.), *The Handbook of Econometrics, Volume 4*, pp. 3041–3080. North-Holland.

Harris, L. E. (1994). Minimum price variations, discre bid-ask spreads and quotation sizes. *Rev. Financial Studies 7*, 149–178.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter.* Cambridge: Cambridge University Press.

Hasbrouck, J. (1996). The dynamics of discrete bid and ask quotes. mimeo, Stern School of Business, New York University.

Hasbrouck, J. and G. Sofianos (1992). The trades of market makers: an empirical examination of NYSE specialist. Working paper 92-03, NYSE.

Hasbrouck, J., G. Sofianos, and D. Sosebee (1993). New York Stock Exchange systems and trading procedures. mimeo, Stern School of Business, New York University.

Hastings, W. K. (1970). Monte-carlo sampling methods using markov chains and their applications. *Biometrika 57*, 97–109.

Hausman, J. A., A. W. Lo, and A. C. MacKinlay (1992). An ordered Probit analysis of transaction stock prices. *J. Financial Economics 31*, 319–330.

Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models (with discussion). *J. Business and Economic Statist. 12*, 371–417.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, Transactions ASMA, Series D 82*, 35–45.

Kim, S., N. Shephard, and S. Chib (1997). Stochastic volatility: Optimal likelihood inference and comparison with ARCH models. *Rev. Economic Studies 65*. Forthcoming.

Kitagawa, G. (1987). Non-Gaussian state space modelling of non-stationary time series. *J. Am. Statist. Assoc. 82*, 503–514.

Lyons, R. K. (1991). Private beliefs and information externalities in the foreign exchange market. Mimeo, Dep. of Finance, Columbia University, New York.

Manrique, A. (1997). Econometric analysis of limited dependent time series. Nuffield College, Oxford University, D. Phil thesis.

Manrique, A. and N. Shephard (1997). Simulation based likelihood inference for limited dependent processes. Discussion paper: Nuffield College, Oxford.

Marriott, J. M. and A. F. M. Smith (1992). Reparameterization aspects of numerical Bayesian methodology for autoregressivce moving-average models. *J. Time Series Analysis 13*, 327–343.

Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller and E. Teller (1953). Equations of state calculations by fast computing machines. *J Chemical Physics 21*, 1087–1092.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset pricing: a new approach. *Econometrica 59*, 347–370.

Priestley, M. B. (1981). *Spectral Analysis and Time Series*. London: Academic Press.

Qian, W. and D. M. Titterington (1991). Estimation of parameters in hidden Markov Chain models. *Phil. Trans. R. Soc. Lond. A. 337*, 407–428.

Ruud, P. (1991). Extensions of estimation methods using the EM algorithm. *J. Econometrics 49*, 305–341.

Shephard, N. (1993). Fitting non-linear time series models, with applications to stochastic variance models. *J. Appl. Econometrics 8*, S135–S152.

Shephard, N. (1996). Statistical aspects of ARCH and stochastic volatility. In D. R. Cox, O. E. Barndorff-Nielson, and D. V. Hinkley (Eds.), *Time Series Models in Econometrics, Finance and Other Fields*, pp. 1–67. London: Chapman & Hall.

Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika 84*, 653–667.

Tanner, M. A. (1996). *Tools for Statistical Inference: methods for exploration of posterior distributions and likelihood functions* (3 ed.). New York: Springer-Verlag.

Taylor, S. J. (1994). Modelling stochastic volatility. *Mathematical Finance 4*, 183–204.

Tierney, L. (1994). Markov Chains for exploring posterior distributions (with discussion). *Ann. Statist. 21*, 1701–1762.

West, M. and J. Harrison (1997). *Bayesian Forecasting and Dynamic Models* (2 ed.). New York: Springer-Verlag.