Market microstructures

Assumptions	continuous prices	$\times$	ticks
	evenly spaced time intervals	$\times$	transaction data
	market clearing	$\times$	market makers

To relax these assumptions, one must investigate microstructure effects. Albeit prices must fully reflect the available information in the market **in the long run**, issues such as market liquidity, competition and collusion among market makers, and asymmetric information are of major importance **in the short run**. For instance, a market maker determines bid and ask prices for securities to protect himself against insider traders and for inventory reasons. Goodhart and O'Hara (1997) provides an excellent survey of the literature.

Accordingly, we examine three topics in microstructure: nonsynchronous trading, the bid-ask spread and price discreteness. We commence with purely statistical models to give an idea of empirical anomalies that can arise in the presence of nontrading and bid-ask bouncing (Campbell, Lo and MacKinlay, 1997, chapter 3). Next, we introduce asymmetric information models to demonstrate the close connection between these two issues (O'Hara, 1995). Lastly, we discuss briefly how to deal with price discreteness using an ordered probit model as proposed by Hausman, Lo and MacKinlay (1992).

Nonsynchronous trading

Before start modeling nontrading, let's consider the seemingly naive practice of taking closing quotes as daily prices. There is an implicit assumption that observations are recorded in equally spaced intervals of 24 hours. As consequence, a **false impression of predictability** in returns likely arise and so the estimation of moments and co-moments of asset returns might carry severe bias.

Example 1. Consider an actively traded asset A and a rarely traded asset B. Now suppose that news arrive near the market close. Then, A's end-of-the-day price is more likely to reflect this information than B's. Eventually, B's price will mirror the news and spurious cross-correlation will appear. Moreover, B's daily returns will also exhibit spurious autocorrelation. To appreciate this, note that B's return is zero in periods of nontrading, whereas in periods of trading the observed return reverts to the cumulative mean return. Thus, negative autocorrelation arises due to the mean reversion.

The bias generated by nonsynchronous trading is particularly relevant when the focus is on testing nonlinear patterns, on predictability, and on risk analysis. There are two strands of the literature. The first considers nontrading as the result of the **strategic interaction of economic forces and private information**. The second views nontrading as the outcome of institutional features such as lagged adjustments and nonsynchronously recorded prices. It is a sort of **measurement error problem** in which spurious serial dependence arises, hence it suffices to model nontrading in a statistical manner.

To capture the effects of nontrading as a **purely statistical artifact**, we follow the nonsynchronous-trading model of Lo and MacKinlay (1990), which associates an unobserved continuously compounded return r_{it} with each security i at time t. **These virtual returns represent changes in the underlying value of the security in the absence of trading frictions and other institutional rigidities**. Suppose that in each period t, there is a fixed probability π_i, independent of r_{it}, that security i does not trade. Thus the nontrading process can be viewed as an iid sequence of coin tosses, with different nontrading probabilities across securities.

Next, define the nontrading indicator δ_{it}, which takes value one if there is no trade, zero otherwise. In other words,

$$\delta_{it} = \begin{cases} 1 & \text{with probability } \pi_i \\ 0 & \text{with probability } 1 - \pi_i, \end{cases}$$

so that $X_{it}(k) = (1 - \delta_{it})\delta_{it-1} \ldots \delta_{it-k}$ corresponds to the indicator function associated with event of trading after $k > 0$ consecutive nontrading periods, viz.

$$X_{it}(k) = \begin{cases} 1 & \text{with probability } (1 - \pi_i)\pi_i^k \\ 0 & \text{with probability } 1 - (1 - \pi_i)\pi_i^k. \end{cases}$$

Further, we assume that δ_{it} and δ_{jt} are independent for every $i \neq j$, and iid over time.

The observed returns are then $r^*_{it} = \sum_{k=0}^{\infty} X_{it}(k) r_{it-k}$, i.e. a stochastic function of all past returns. Put differently, if $K_t = \sum_{j=1}^{\infty} \prod_{k=1}^{j} \delta_{it-k}$ denotes the number of past consecutive nontrading periods, then the observed return is simply a random sum of the past virtual returns (akin to a Poisson process), i.e. $r^*_{it} = \sum_{k=0}^{K_t} r_{it-k}$ for $i = 1, \ldots, N$. Alternatively, one may also write

$$r^*_{it} = \begin{cases} 0 & \text{with probability } \pi_i \\ r_{it} & \text{with probability } (1-\pi_i)^2 \\ r_{it} + r_{it-1} & \text{with probability } (1-\pi_i)^2 \pi_i \\ \vdots & \vdots \\ r_{it} + \ldots + r_{it-k} & \text{with probability } (1-\pi_i)^2 \pi_i^k \\ \vdots & \vdots \end{cases}$$

which illustrates the fact that spurious autocorrelation arises for r^*_{it} is determined by the sum of past k consecutive virtual returns for all k with positive probability.

Now let's impose some structure on the dynamics of the virtual returns and see what we get from the model. Suppose, for instance, that

$$r_{it} = \mu_i + \beta_i f_t + \upsilon_{it}, \qquad i = 1, \ldots, N$$

where μ_i represents the drift, f_t a zero-mean common factor, and υ_{it} a zero-mean idiosyncratic noise. Assume further that f_t is iid, so that autocorrelation in the observed returns will stem exclusively from nontrading, and that f_t and υ_{it-k} are independent for every i, t, and k. Then, Lo and MacKinlay (1990) show that nontrading involves second-order effects, though no mean effect arises, i.e. $E(r^*_{it}) = \mu_i$. More precisely, it is possible to show that

$$\mathrm{Var}(r^*_{it}) = \sigma_i^2 + \frac{2\pi_i}{1-\pi_i}\mu_i^2$$

$$\mathrm{Cov}(r^*_{it}, r^*_{jt+k}) = \begin{cases} -\mu_i^2 \pi_i^k & \text{for } i = j, k > 0 \\ \frac{(1-\pi_i)(1-\pi_j)}{1-\pi_i\pi_j} \beta_i \beta_j \sigma_f^2 \pi_j^k & \text{for } i \neq j, k \geq 0 \end{cases}$$

$$\mathrm{Corr}(r^*_{it}, r^*_{it+k}) = -\frac{\mu_i^2 \pi_i^k}{\sigma_i^2 + 2\mu_i^2 \pi_i/(1-\pi_i)},$$

where $\sigma_i^2 = \mathrm{Var}(r_{it})$ and $\sigma_f^2 = \mathrm{Var}(f_t)$. Note that, if there is a drift ($\mu_i \neq 0$), observed returns will have a higher variance than virtual returns and will exhibit negative serial correlation with geometric decay as during nontrading periods the observed return is zero, whereas during trading periods the observed returns revert back to its cumulative mean return.

Remark 1. The maximal spurious autocorrelation due to nontrading is easily computed as the serial correlation is a nonpositive continuous function of the probability of nontrading that equals zero for $\pi_i = 0$ and approaches zero again as the probability of nontrading goes to one. Indeed, Lo and MacKinlay (1990) show that $\pi_i^* = (1 + \sqrt{2}|\mu_i/\sigma_i|)^{-1}$, which implies that

$$\inf_{\pi_i^*, \mu_i/\sigma_i} \mathrm{Corr}(r_{it}^*, r_{it+k}^*) = -\frac{1}{2}.$$

Notice that this lower bound is never attained as it requires $|\mu_i/\sigma_i| \to \infty$.

Remark 2. At first glance, it seems that doubling the sampling interval increases the spurious negative autocorrelation as $\mu_i^* = 2\mu_i$ and $\sigma_i^* = \sqrt{2}\sigma_i$. However, the nontrading process is not independent of the sampling interval, so that one must accounts for time-aggregation effects.

Remark 3. The sign of the spurious cross autocorrelation depends on the sign of $\beta_i\beta_j$. More importantly, there is an asymmetry with respect to i and j in the general case where $\pi_i \neq \pi_j$. The intuition is simple. If a nontrading event occurs to security j, then even observed returns to a systematically trading security i can forecast j due to the common factor f_t. The converse does not hold, since $r_{it}^* = r_{it}$ (as $\pi_i = 0$) is by assumption an iid process. Therefore, though $\mathrm{Corr}(r_{it}^*, r_{jt+k}^*) \neq 0$ for all $k > 0$, $\mathrm{Corr}(r_{jt}^*, r_{it+k}^*) = 0$ for every $k > 0$.

The model entails some testable implications. Denote by r_t^* the vector of observed returns $\{r_{1t}^*, \ldots, r_{Nt}^*\}$ and let $\mu = \{\mu_1, \ldots, \mu_N\}$. Recall that the covariance matrix of the vector of observed returns is simply

$$\Gamma_k = E(r_t^* - \mu)(r_t^* - \mu)' = [\gamma_{ij}(k)], \qquad i, j = 1, \ldots, N$$

where

$$\gamma_{ij}(k) = \frac{(1-\pi_i)(1-\pi_j)}{1-\pi_i\pi_j}\beta_i\beta_j\sigma_f^2\pi_j^k \qquad \text{for } i \neq j, k \geq 0,$$

which yields $\gamma_{ij}(k)/\gamma_{ji}(k) = (\pi_j/\pi_i)^k$. Therefore, relative nontrading probabilities may be estimated using sample autocovariances, whereas the sample mean and autocorrelation provide estimates for μ and the reference nontrading probability, say π_1.

The same framework can be extended to analyze returns of a portfolio of securities with identical nontrading probabilities. In particular, the observed returns of the portfolio will follow an AR(1) process with autoregressive coefficient equal to the nontrading probability irrespective of the expected return of the portfolio. Again, one may test for the over-identifying restrictions implied by the model, namely

$$\frac{\mathrm{Cov}(r^*_{at}, r^*_{bt+k})}{\mathrm{Cov}(r^*_{bt}, r^*_{at+k})} = \frac{\mathrm{Cov}(r^*_{at}, r^*_{ct+k})\mathrm{Cov}(r^*_{ct}, r^*_{bt+k})}{\mathrm{Cov}(r^*_{ct}, r^*_{at+k})\mathrm{Cov}(r^*_{bt}, r^*_{ct+k})} = \left(\frac{\pi_b}{\pi_a}\right)^k.$$

As before, if the nontrading probability varies across portfolios, then asymmetry will arise in the autocovariance of observed portfolio returns. This might explain the empirical evidence supporting the existence of a lead-lag relation in size-sorted portfolios.

Further, these results are robust to time aggregation in the sense that the spurious autocorrelation remains negative, though it declines monotonically in magnitude as the sampling interval increases. As before, expected returns aggregate linearly, whilst variances do not due to the presence of negative serial correlation (variance of the sum is lesser than the sum of the variances). For calibrated parameters, the impact of nontrading is negligible in the short horizon for individual stocks, but substantial for portfolios.

Remark 4. Suppose that the nontrading indicator δ_{it} follows a Markov chain. Then, individual stock returns may be even positively autocorrelated and portfolio returns may be negatively autocorrelated, though unlikely for empirically relevant parameter values. Moreover, spurious index autocorrelation increases (decreases) as positive (negative, respectively) persistence in nontrading mounts. However, for plausible parameter vectors, this model is unable to help understanding the magnitude of index autocorrelation in recent stock market data.

Market microstructure theory provides several ways to incorporate economic forces into nontrading models, but we will focus exclusively on **information-based models** for they entail a natural link to the bid-ask spread. These models use adverse selection arguments to show how, even in competitive markets without explicit transaction costs, spreads would exist. One major aspect of information-based models is that they allow for examination of market dynamics and hence provide insights into the adjustment process of prices. The main intuition behind these models dates back to Baghehot (1971), who claims that the market maker sets a bid-ask spread to balance losses to insider (or informed) traders with gains from liquidity (or uninformed) traders.

We follow Easley and O'Hara (1992) in considering a **sequential trade model** akin to Easley and O'Hara (1987). To concentrate on the effect of information on prices, we assume a **single market maker, who is risk neutral and acts competitively**. The former rules out direct inventory effects, whilst the latter implies the existence of at least potential competitors. Let V denote the value of a certain asset and define an information event as the occurrence of a signal ψ about V. The signal can take on one of two values, L and H, with probabilities $\delta > 0$ and $1 - \delta > 0$ ($\hat{\delta} \approx 1/2$). The expected value of the asset conditional on the signal is $E(V \,|\, \psi = L) = V_L$ or $E(V \,|\, \psi = H) = V_H$. If no information event occurs ($\psi = 0$), then the expected value of the asset simply remains at its unconditional level $V_* = \delta V_L + (1 - \delta)V_H$.

Information events occur with probability $\alpha \in (0, 1)$ before the start of the current trading day ($\hat{\alpha} \approx 3/4$). This uncertainty reflects that, since uninformed market participants do not receive any signal, they may also not know whether any new information even exists (e.g. Dow-Jones Rumor Wire). Trade can arise from uninformed and/or informed traders. To keep focus on information effects, we assume that the informed traders are risk neutral and price takers so as to rule out strategic behavior. As such, the resulting trading strategy reads: If a high signal occurs, an insider will buy the stock if the current quote is below V_H; if a low signal occurs, then he will sell if the quote is above V_L.

To avoid no-trade equilibrium, we further assume that at least some uninformed market participants trade for nonspeculative reasons such as liquidity needs or portfolio considerations. Further, suppose that there is a fraction γ of potential sellers and a fraction $1-\gamma$ of potential buyers among liquidity traders ($\hat{\gamma} \approx 1/2$). Uninformed buyers trade with probability ϵ_B, whereas an uninformed seller's trading probability is ϵ_S ($\hat{\epsilon}_B \approx \hat{\epsilon}_S \approx 1/3$).

The assumptions of risk neutrality and competitive behavior for the market maker dictate that price quotes yield zero expected profit conditional on a trade at the quote. As insider traders will profit at the market maker's expense, the probability $\mu \in (0,1)$ that a trade is actually information-based is crucial for determining the price quotes ($\hat{\mu} \approx 1/6$). There are several ways to interpret μ, e.g. the fraction of the trader population who observes the signal, the probability of information disclosure to the selected trader, or the exogenous order arrival rate.

Transaction occurs throughout the trading day, which is divided into discrete intervals of time ($t = 1, 2, \ldots$), long enough to accommodate at most one trade. The exact length of a trading interval is arbitrary and could even approach zero (continuous-time) so as to reformulate the statistical model in terms of Poisson arrivals. In a discrete-time context, inspection of transactions data suggests that trades rarely occur more often than every 5 seconds. At each time t, the market maker announces the bid and ask prices at which he is willing to trade one unit of the asset.

Insiders always trade provided that prices are not at their full information value, whereas liquidity traders may buy, sell or not trade according to their type (buyer or seller) and motivation. Therefore, non trading occurs only when an uninformed trader checks the quotes and decides for portfolio reasons (as captured by ϵ_B and ϵ_S) not to trade. This can happen both when there has been an information event and when there has not. Of course, nontrading outcomes are more likely to occur when there is no new information than when new information arrives: $(1-\gamma)(1-\epsilon_B)+\gamma(1-\epsilon_S) > (1-\mu)[(1-\gamma)(1-\epsilon_B)+\gamma(1-\epsilon_S)]$.

The tree diagram in Figure 1 summarizes the trading process.

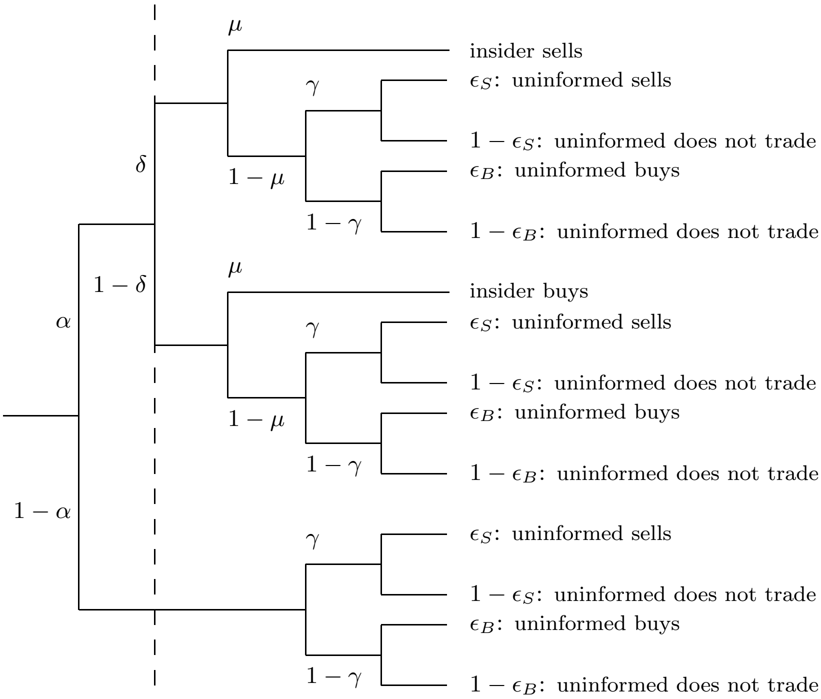

Figure 1 — Tree diagram of the trading process

Notation: α is the probability of an information event, δ is the probability of a low signal, μ is the probability a trade comes from an informed trader, γ is the probability that an uninformed trader is a seller, $1-\gamma$ is the probability that an uninformed trader is a buyer, ϵ_S is the probability that the uninformed trader will sell, and ϵ_B is the probability that the uninformed trader will buy. Nodes to the left of the dotted line occur only at the beginning of the trading day; nodes to the right occur at each trading interval.

The market maker knows the structure of the market and updates his beliefs by Bayes rule. It is precisely this revision that causes quotes, and thus prices, to adjust. The quote-setting process for the first trade of the day looks as follows. To determine the initial bid (ask) price, the market maker must compute the expected value of the asset given that a sale (buy, respectively) takes place. For a trade outcome $Q \in \{B, S, N\}$, the market maker updating formula reads

$$\Pr(V = V_L \,|\, Q) = \Pr(\psi = L \,|\, Q) + \delta \Pr(\psi = 0 \,|\, Q),$$

where by Bayes rule

$$\Pr(\psi = X \,|\, Q) = \frac{\Pr(\psi{=}X)\Pr(Q|\psi{=}X)}{\Pr(\psi{=}L)\Pr(Q|\psi{=}L){+}\Pr(\psi{=}H)\Pr(Q|\psi{=}H){+}\Pr(\psi{=}0)\Pr(Q|\psi{=}0)}.$$

This implies, for instance, that

$$\begin{aligned}
\Pr(\psi = L \,|\, S) &= \frac{\delta[\alpha\mu + \alpha(1-\mu)\gamma\epsilon_S]}{\delta\alpha\mu + (1-\alpha\mu)\gamma\epsilon_S} \\
\Pr(\psi = 0 \,|\, S) &= \frac{(1-\alpha)\gamma\epsilon_S}{\delta\alpha\mu + (1-\alpha\mu)\gamma\epsilon_S},
\end{aligned}$$

hence

$$\Pr(V = V_L \,|\, S) = \frac{\delta[\alpha\mu + (1-\alpha\mu)\gamma\epsilon_S]}{\delta\alpha\mu + (1-\alpha\mu)\gamma\epsilon_S} > \delta.$$

As is apparent, the market maker increases the probability he attaches to V_L once someone wants to sell. Conversely, in the case of a buy the market maker decreases the probability he assigns to V_L.

Given these conditional expectations, the market maker can set initial bid $b_1 \equiv E(V \,|\, S)$ and ask $a_1 \equiv E(V \,|\, B)$ quotes, namely

$$\begin{aligned}
b_1 &= \frac{\delta[\alpha\mu + (1-\alpha\mu)\gamma\epsilon_S]V_L + (1-\delta)(1-\alpha\mu)\gamma\epsilon_S V_H}{\delta\alpha\mu + (1-\alpha\mu)\gamma\epsilon_S} \\
a_1 &= \frac{\delta(1-\alpha\mu)(1-\gamma)\epsilon_B V_L + (1-\delta)[\alpha\mu + (1-\alpha\mu)(1-\gamma)\epsilon_B]V_H}{(1-\delta)\alpha\mu + (1-\alpha\mu)(1-\gamma)\epsilon_B},
\end{aligned}$$

respectively. More generally, the market maker's bid quote at t is the expected value of the asset given the history of trade outcomes $Q^{t-1} = (Q_1, \ldots, Q_{t-1})$ and a sale at t, $Q_t = S$. The bid in period t is then

$$b_t = \Pr\big(\psi = L \,|\, Q^{t-1}, S\big)V_L + \Pr\big(\psi = H \,|\, Q^{t-1}, S\big)V_H + \Pr\big(\psi = 0 \,|\, Q^{t-1}, S\big)V_*.$$

Similarly, the ask at t is the expected value of the asset given the history of trade outcomes $Q^{t-1} = (Q_1, \ldots, Q_{t-1})$ and a buy at t, $Q_t = B$. Put differently, the evolution of prices is determined by the evolution of beliefs.

Property. If there is no trade at time t, then the probability the market maker attaches to no information event rises and the probabilities of a high or a low signal fall, though the relative probability of a low to a high signal remains constant. As such, the market maker learns from the lack of trade as well as from actual transactions. In particular both bid and ask prices move in response to the absence of trade. The bid rises (falls) if inferior (superior, respectively) to V_*, whereas the converse is true for the ask. Moreover, if $a_t > V_* > b_t$ and there is no trade at t, then the spread in period $t+1$ will narrow.

The intuition for this last result is simple. The spread shields the market maker against the risk of trading with an insider. As the lack of trade is more likely to occur when there is no new information in the market, trading is now safer and the market maker reduces the bid-ask spread. What is intriguing about this effect is that it suggests that time between trades may play an important role in the behavior of prices. We will return to this issue later when we discuss conditional duration models.

To determine the evolution of beliefs, notice that there is no need to keep track of the order in which sales, buys and no trades arrive in the market. Suppose, for instance, that in the past t trading intervals there were n_t no trades, β_t buys, and s_t sales. Then, (n_t, β_t, s_t) is a sufficient statistic for Q^t. For instance, the probability the market maker assigns to the absence of new information given this trading history reads

$$\begin{aligned}
\Pr(\psi = 0 \,|\, Q^t) \quad = \quad & (1-\alpha)(\gamma\epsilon_S)^{s_t}[(1-\gamma)\epsilon_B]^{\beta_t} \Big\{(1-\alpha)(\gamma\epsilon_S)^{s_t}\big[(1-\gamma)\epsilon_B\big]^{\beta_t} \\
& + (1-\mu)^{n_t}\Big[\alpha\delta\big(\mu + (1-\mu)\gamma\epsilon_S\big)^{s_t}\big((1-\mu)(1-\gamma)\epsilon_B\big)^{\beta_t} \\
& + \ \alpha(1-\delta)\big((1-\mu)\gamma\epsilon_S\big)^{s_t}\big(\mu + (1-\mu)(1-\gamma)\epsilon_B\big)^{\beta_t}\Big]\Big\}^{-1}.
\end{aligned}$$

More importantly, as beliefs depend on (n_t, β_t, s_t), quotes will also depend on these quantities. In particular, the bid and ask at time $t+1$ can be written as

$$\begin{aligned}
b_{t+1} \quad = \quad & \Pr(\psi = L \,|\, n_t, \beta_t, s_t + 1)V_L + \Pr(\psi = 0 \,|\, n_t, \beta_t, s_t + 1)V_* \\
& + \Pr(\psi = H \,|\, n_t, \beta_t, s_t + 1)V_H \\
a_{t+1} \quad = \quad & \Pr(\psi = L \,|\, n_t, \beta_t + 1, s_t)V_L + \Pr(\psi = 0 \,|\, n_t, \beta_t + 1, s_t)V_* \\
& + \Pr(\psi = H \,|\, n_t, \beta_t + 1, s_t)V_H,
\end{aligned}$$

respectively.

As is apparent, both bid and ask quotes depend not only on the most recent quote, but also on the total numbers of previous buys, sales and no-trade outcomes. This results has two important implications. First, the **bid and ask prices do not satisfies the Markov property**, though the process (n_t, β_t, s_t) does. Second, because volume is related to the number of no-trade outcomes, which is in turn linked with the probability of an information event, **prices depend on the previous volume of trade**. To appreciate this, recall that volume to time t is, by definition, given by $v_t = \beta_t + s_t = t - n_t$. Similarly, the market maker's inventory position at time t is given by $i_t = s_t - \beta_t$. The informational content of (v_t, i_t, t) is thus equivalent to that of (n_t, β_t, s_t).

It seems somewhat counterintuitive that allowing the market maker to learn from the lack of trades as well as from transactions means that volume now matters in the stochastic process governing the dynamics of prices. Yet, it is precisely because no-trade outcomes can occur that the aggregate total of transactions provides information. The probability that the market maker attaches to no new information is increasing in n_t and thus decreasing in v_t. Hence, the market maker views v_t as a signal of the **existence** of information. Similarly, the market maker interprets i_t as a signal of the **direction** of any new information.

To examine the dynamic behavior of the market, we need to define the stochastic process of market prices. Let's first define the stochastic process of conditional expected values $\{p_t^*\}$ by

$$p_t^* = \begin{cases} a_t & \text{if } Q_t = B \\ b_t & \text{if } Q_t = S \\ E(V \mid Q^{t-1}, N) & \text{if } Q_t = N. \end{cases}$$

This process is a martingale because it is a sequence of conditional expectations where $E(p_{t+1}^* \mid Q^t) = E\big[E(V \mid Q^{t+1}) \mid Q^t\big] = E(V \mid Q^t) = p_t^*$. Unfortunately this process is not observable because we do not know the market maker's expectation in the absence of a trade.

A transaction price arises only when a trader decides to buy or sell. From a statistical perspective, we can view the transactions price process as a subordinated process of $\{p_t^*\}$, viz. $p_j = p_{t_j}$, where $t_j = \min\{t : t > t_{j-1} \text{ and } Q_{t_j} \neq N\}$ with $t_0 = 0$. This optional sampling increases substantially the complexity of the statistical problem in view that the sampling times are not iid but are, instead, partially chosen by traders who may be informed about the evolution of the price process. As such, the time between trades will not be independent of the evolution of the price process. Further, the sequence of transaction prices is a martingale not only with respect to the sequence of trades, but also with respect to past prices. This means that the market is weakly efficient, though strong efficiency does not hold as, only in the limit, transaction prices converge to their full information values.

The optional sampling affects the variance structure of the price process as well. Denote by $E\left[\left(p_t^* - p_{t-1}^*\right)^2 \,\middle|\, Q^{t-1}\right]$ the conditional variance of the expected asset value at time t. If there is no trade at t, then the variance falls at time $t+1$. Hence, roughly, the **variance is positively correlated with volume** (Lamoureaux and Lastrapes, 1990). Now as trades are positively serially correlated, **volatility clustering** arises: low (high) variance periods are related to thin (heavy, respectively) trading. The transaction price is a sampling of the expected value process exactly when this process has high variance – at trade times. So, the variance of the transaction price process overestimates the variance of the conditional value of the asset or, equivalently, of the price at which trades could have occurred.

Testable implications. (1) Quotes will change even in the absence of trades, (2) the bid-ask spread will decrease the longer the time between transactions (Hasbrouck, 1991), (3) the time interval between trades is not exogenous to the price process (Hausman and Lo, 1990), (4) durations are autocorrelated (Engle and Russell, 1997 and 1998), (5) price processes are martingale, but not Markovian.

Implication (3) means that models of nonsynchronous trading à la Lo and MacKinlay (1990) are doom to fail as they rely on the independence between nontrading intervals and the virtual return process for the stock. To avoid the optional sampling issue, it seems more convenient to work with data on quote prices rather than transaction prices. Quotes, unlike transaction prices, occur continually and carry more information than transaction prices in Easley and O'Hara (1992). Albeit the process of quotes suffers from history dependence, (n_t, β_t, s_t) stand as a simple sufficient statistic for the history of the process.

Extensions and applications. Easley, Kiefer and O'Hara (1997b) show how to estimate and test the Easley and O'Hara's (1992) model using trade data. Easley, Kiefer and O'Hara (1997a) combine the Easley and O'Hara's (1987 and 1992) models to gauge the information content of trade size. In addition, they permit the probability that the uninformed transact to depend on the previous trade outcome. Easley, Kiefer, O'Hara and Paperman (1996) investigate to what extent differences in spreads for active and infrequently traded stocks are due to adverse selection. Easley, O'Hara and Paperman (1998a) examine the informational role of financial analysts by looking at the probability of information-based trading for a sample of NYSE stocks that differ in analyst coverage. The results indicate that the probability of information-based trade is lower for stocks with many analysts. What happens is that, though stocks with more analysts do have more insider trade, they have even greater rates of uninformed trade. Easley, O'Hara and Srinivas (1998b) incorporate to the baseline model the possibility of trading in the options market to explain how option volumes may affect future stock prices. Lastly, Grammig, Schiereck and Theissen (2000) investigate whether the probability of informed trading vary with the degree of market anonymity. Data stem from the German stock market, where an electronic anonymous market and a nonanonymous floor trading system coexist.

The bid-ask bounce

As random buys and sells arrive, prices bound back and forth between the ask and bid quotes, what may generate spurious volatility and autocorrelation. Of course, it also brings about some other intriguing aspects. For instance, the bid-ask bounce may explain the **January effect**, namely the stylized fact that smaller capitalization stocks outperform larger capitalization stocks over the few days near to the turn of the year. To see why, it suffices to note that agents tend to sell stocks in December and to start buying again in January. Hence, transaction prices are close to the bid quotes in December and then approach the ask quotes in January. Yet, lower capitalization stocks have usually lower prices, and thus higher percentage bid-ask spread. Hence, they will provide larger returns.

In adverse selection model, the bid-ask spread compensates the market maker for engaging in losing trades with informed market participants. More generally, the spread rewards the market maker for providing liquidity to the market. To focus on the effects of the bid-ask bounce, Roll (1984) models a frictionless economy in which the spread s is completely exogenous (and so constant). The model dictates that the market price P_t oscillates about the fundamental value P_t^* of the asset according to whether the trade is at the bid (seller-initiated) or at the ask (buyer-initiated). More precisely, $P_t = P_t^* + \frac{s}{2} I_t$, where

$$I_t = \begin{cases} +1 & \text{if buyer initiative} \\ -1 & \text{if seller initiative} \end{cases}$$

is an iid order-type indicator variable. The assumption that P_t^* is the fundamental value of the security implies that $E(P_t) = E(P_t^*)$, and thus $E(I_t) = 0$ or equivalently $\Pr(I_t = 1) = \Pr(I_t = -1) = 1/2$.

Assume for the moment that the fundamental value is fixed through time, $P_t^* = P^*$. The process for price changes reads then

$$\Delta P_t = \Delta P_t^* + \frac{s}{2}\,\Delta I_t = \frac{s}{2}\,\Delta I_t,$$

which implies that $\mathrm{Var}(\Delta P_t) = s^2/2$ and $\mathrm{Cov}(\Delta P_t, \Delta P_{t-1}) = -s^2/4$. Despite the fact that the fundamental value is constant, price changes exhibit volatility and **negative serial correlation** as the result of the bid-ask bounce. The intuition is simple: If the current price is the ask, then the price change between the current price and the previous price must have been either zero or s, while the price change between the next price and the current price must be either zero or $-s$. Of course, the same argument holds if the current price is the bid. Furthermore, serial correlation of either sign in the order type indicator I_t does not change the fact that bid-ask bouncing induces negative autocorrelation in price changes, although it does affect its magnitude.

Remark 5. Albeit both volatility and first-order autocovariance are increasing in the spread, the first-order autocorrelation remains constant at -1/2. Moreover, the bid-ask spread does not induce any higher-order serial correlation. Now let P_t^* vary through time, but suppose that its increments are serially uncorrelated and independent of I_t. Then, the autocovariance remains the same, but there will be an additional term in the variance of ΔP_t to accommodate the variation in ΔP_t^*. Hence, the first-order autocorrelation will now depend on the spread s, viz.

$$\mathrm{Corr}(\Delta P_t, \Delta P_{t-1}) = -\frac{s^2/4}{s^2/2 + \sigma_*^2} \leq 0,$$

where σ_*^2 denotes the variance of ΔP_t^*.

Roll (1984) takes advantage of the implied relationship between the bid-ask spread and the autocovariance to define the **effective spread**, namely

$$s = 2\sqrt{-\mathrm{Cov}(\Delta P_t, \Delta P_{t-1})}.$$

To recover the effective spread, it suffices then to estimate the sample autocovariance of price changes. In view that the bid and ask quotes are observable, it seems at first glance superfluous to estimate the effective spread. Roll claims however that the quoted spread may differ from the effective spread in many instances, e.g. market makers may wish to rebalance their own inventory. In implementing his model, Roll (1984) argues that it is often more convenient to work with the relative rather than absolute bid-ask spread. The first-order autocovariance of simple returns R_t relates then to the relative spread s_r as follows

$$\mathrm{Cov}(R_t, R_{t-1}) = -\frac{s_r^2}{4} - \frac{s_r^4}{16} \approx -\frac{s_r^2}{4},$$

where the relative spread is measured as the percentage of the geometric average of the bid and ask prices P_a and P_b, i.e. $s_r = s/\sqrt{P_a P_b}$. Hence, it yields

$$s_r \approx 2\sqrt{-\mathrm{Cov}(R_t, R_{t-1})},$$

which is easily estimated from the sample autocovariance of the simple returns.

Price discreteness

In practice, the implicit assumption that prices are continuous does not hold as prices are quoted in ticks. Hausman et al. (1992) attempt to capture price discreteness using an ordered probit model. Let $P(t_0), P(t_1), \ldots, P(t_n)$ denote transaction prices that are possibly irregularly spaced in time, and let $\big\{Y_j = P(t_j) - P(t_{j-1}),\, t_j > 0\big\}_{j=1}^{n}$ denote then the process of price changes. Suppose now that the underlying process of shadow price changes is $Y_j^* = X_j'\beta + \epsilon_j$, where $\epsilon_j \sim IN(0, \sigma_j^2)$ with $\sigma_j^2 = W_j'\gamma$ such that $E(\epsilon_j | X_j, W_j) = 0$ for the vectors X_j and W_j of exogenous variables. The ordered probit model assumes then that

$$Y_j^* = \begin{cases} s_1 & \text{if } Y_j^* \in A_1 \\ s_2 & \text{if } Y_j^* \in A_2 \\ \vdots & \vdots \\ s_m & \text{if } Y_j^* \in A_m, \end{cases}$$

where the sets A_k form a partition of the state space $\mathcal{S}^*$ of Y_j^*, i.e. $\mathcal{S}^* = \cup_{k=1}^{m} A_k$ and $A_k \cap A_{k'} =$ for $k \neq k'$, and s_k are discrete values that comprise the state space $\mathcal{S}$ of Y_j. How to partition the state space is an arbitrary issue that involves a trade-off between parsimony and price resolution. As transactions with returns superior to 4 ticks are pretty rare, Hausman et al. define the s_k's as tick multiples: s_1 corresponds to a price change of -4 ticks or less, s_9 to price change of 4 ticks or more, and s_2 to s_8 to price changes of -3 ticks to 3 ticks, respectively.

The dependence structure of the observed process Y_j is completely characterized by that of Y_j^* and the partition, viz.

$$\Pr\big(Y_j = s_k \,\big|\, Y_{j-1} = s_{k'}\big) = \Pr\big(Y_j^* \in A_k \,\big|\, Y_{j-1}^* \in A_{k'}\big).$$

Time independence of the observed process Y_j depends then on whether X_j and W_j are time independent. Nevertheless, we require only the conditional independence of the error term ϵ_j given $Z_j = (X_j, W_j)$. The partition of the boundaries and the distribution of ϵ_j determine then the conditional distribution of observed price changes Y_j given Z_j. Assuming normality for the ϵ_j's, the conditional distribution reads

$$\begin{aligned}
\Pr\big(Y_j = s_k \,\big|\, Z_j\big) &= \Pr\big(X_j'\beta + \epsilon_j \in A_k \,\big|\, Z_j\big) \\
&= \begin{cases} \Pr\big(X_j'\beta + \epsilon_j \le \alpha_1 \,\big|\, Z_j\big) & \text{if } k = 1 \\ \Pr\big(\alpha_{k-1} < X_j'\beta + \epsilon_j \le \alpha_k \,\big|\, Z_j\big) & \text{if } 1 < k < m \\ \Pr\big(X_j'\beta + \epsilon_j > \alpha_{m-1} \,\big|\, Z_j\big) & \text{if } k = m \end{cases} \\
&= \begin{cases} \Phi\left(\frac{\alpha_1 - X_j'\beta}{\sigma_j}\right) & \text{if } k = 1 \\ \Phi\left(\frac{\alpha_k - X_j'\beta}{\sigma_j}\right) - \Phi\left(\frac{\alpha_{k-1} - X_j'\beta}{\sigma_j}\right) & \text{if } 1 < k < m \\ 1 - \Phi\left(\frac{\alpha_{m-1} - X_j'\beta}{\sigma_j}\right) & \text{if } k = m, \end{cases}
\end{aligned}$$

where $\Phi(\cdot)$ corresponds to the standard normal cumulative distribution function, $A_1 = (-\infty, \alpha_1]$, $A_k = (\alpha_{k-1}, \alpha_k]$ for $1 < k < m$, and $A_m = (\alpha_{m-1}, \infty)$.

The assumption of normality plays no special role in determining the probabilities of states, as the ordered probit can fit any arbitrary multinomial distribution by shifting the boundary appropriately. The logistic distribution could have served equally well to approximate the underlying multinomial distribution, though it would not be straightforward to model conditional heteroskedasticity in the resulting ordered logit model.

Let $I_j(k)$ denote an indicator variable which takes value one if the realization of the jth observation Y_j is the kth state s_k, and zero otherwise. The log-likelihood function $\mathcal{L}$ for the vector of price changes $Y = (Y_1, \ldots, Y_n)$ given $X = (X_1, \ldots, X_n)$ and $Z = (Z_1, \ldots, Z_n)$ reads then

$$\begin{aligned}
\mathcal{L}\big(Y \,|\, X, W\big) &= \sum_{j=1}^{n} I_j(1) \log \Phi\left(\frac{\alpha_1 - X_j'\beta}{\sigma_j}\right) \\
&\quad + \sum_{j=1}^{n} \sum_{k=2}^{m-1} I_j(k) \log\left[\Phi\left(\frac{\alpha_k - X_j'\beta}{\sigma_j}\right) - \Phi\left(\frac{\alpha_{k-1} - X_j'\beta}{\sigma_j}\right)\right] \\
&\quad + \sum_{j=1}^{n} I_j(m) \log\left[1 - \Phi\left(\frac{\alpha_{m-1} - X_j'\beta}{\sigma_j}\right)\right].
\end{aligned}$$

To achieve identification, it is typical to set $\gamma_0 = 1$.

Specification. Marsh and Rosenfeld (1986) show that if Y_j^*'s are increments of an arithmetic Brownian motion with variance proportional to the duration $\Delta t_j = t_j - t_{j-1}$, then the correct specification of the ordered probit model consists of $X_j'\beta = \beta_0 \Delta t_j$ and $W_j'\gamma = \gamma_0^2 \Delta t_j$. More generally, one may control for other microstructure effects including variables such as the bid-ask spread, transaction size, volume intensity, and an order-type indicator.

Price resolution and clustering

There is a palpable tendency for prices to fall more frequently on certain values than on others at the high frequency (Harris, 1991). To appreciate this, it is often convenient to check how uniform is the histogram of price fractions. Alternatively, one could plot phase diagrams and look for radially symmetric patterns. At any rate, round fractions are more frequent than halves, odd quarters are more common than odd eighties, and other fractions are rarely observed. In what follows, we present the framework put forward by Harris (1991, 1994) to investigate how price discreteness relates to price clustering and liquidity issues.

Harris (1991) uses data from high-price stocks to characterize the cross-sectional determinants of these discrete price sets and shows that **stock price clustering increases with price level and volatility and decreases with firm size and transaction frequency**. These results are based on a multivariate regression analysis motivated by the **price resolution hypothesis** advanced by Ball, Torous and Tschoegl (1985). They hypothesize that price clustering depends on how well known is the underlying value of the security. In particular, the degree of price clustering increases with value uncertainty. The reason is simple. Traders use discrete price sets to lower their costs of negotiating. The price resolution, and therefore the extent of clustering, depends on a **trade-off between lower negotiation costs and lost gains-from-trade**. Negotiating costs are low if traders use a coarse set of discrete prices because it reduces the time to strike a bargain by limiting the amount of information that must be exchanged between negotiating parties. On the other hand, gains-from-trade may be lost if the set does not include a price that is acceptable to both parties. Lost gains-from-trade are likely if little dispersion exists among trader reservation prices, such as when underlying security values are well known. Traders will therefore prefer a fine set of prices and little price clustering will be observed.

Harris (1991) gauges the degree of clustering in the price fraction frequency distribution by the sample frequency of whole prices, which is a weighted sum of the summed difference between odd and even eighths, quarters and halves. More precisely, the whole number frequency equals to $f_0^* = 1/8 + \delta_8/8 + \delta_4/4 + \delta_2/2$, where $\delta_8 = \sum_{i=0}^{7}(-1)^i f_i$, $\delta_4 = \sum_{i=0}^{3}(-1)^i f_{2i}$ and $\delta_2 = f_0 - f_4$. As prices are serial dependent, the standard sample frequencies f_i are adjusted to reflect the domain over which prices are observed to wander. If prices do not often visit the region near a given eighth, the frequency for that eighth is adjusted upward, whereas if prices dwell in that region the frequency is corrected downward. Let a domain event over a given eighth occur whenever prices change so that the price path crosses or arrives on that eighth. For example, a price change from 1/4 to 5/8 creates three domain events: a passage over 3/8, a passage over 1/2 and an arrival at 5/8. The domain is then estimated by the sample frequency d_i of these domain events and the adjusted eighth frequency estimates are $\hat{f}_i = f_i + 1/8 - d_i$ for $i = 0, \ldots, 7$. Alternatively, one may think of estimating a cyclic Markov chain to properly account for the time dependence among eighths.

Several variables can proxy for the unobserved degree of reservation price dispersion. Time-series volatility should be correlated with the dispersion of reservation prices because information is not uniformly distributed and interpreted when events cause values to change quickly. Firm size should be inversely correlated with reservation price dispersion because more information is produced about large firms than about small firms and because large firms are generally better diversified (and therefore easier to value). Trade intensity should be inversely correlated with reservation price dispersion because trading tends to reveal stock values by aggregating the information possessed by different traders. In addition, it is convenient to include price level as a regressor because an eighth represents a smaller fraction of price for high-price stocks than for low-price-stocks. As a first approximation, traders are assumed to use discrete price sets based on minimum price variations that are constant fractions of price. This implies larger price variations (and hence more clustering) for high-price stocks than for low-price stocks.

Harris (1994) provides further insight by studying the effects of a $1/16 minimum price variation on quotation sizes and liquidity. The minimum price variation presumably affects trading volume if it forces dealers to establish a larger spread than they would otherwise quote. Large bid-ask spreads make trade expensive, especially for small traders. Further, the minimum price variation will affect displayed market depth (the sizes of the bid and ask quotations) when it is larger than the spread that dealers would otherwise quote. The spread then will equal the minimum price variation, and supplying liquidity could be quite profitable, especially to small orders. A binding minimum price variation (as it usually occurs to low-price or actively traded stocks) may also increase quotation sizes simply because dealers slide up their implicit quotation schedules, which state the prices for given quantities at which they are willing to trade. As quantities increase, dealers typically require greater spreads to cover the risk of losing to large traders who are likely to be better informed than smaller traders (Easley and O'Hara, 1987).

The minimum price variation may also affect displayed market depth even if it is inferior to the quoted spread. **The minimum price variation determines the minimum cost of acquiring order precedence through price priority when time precedence is enforced**. Time precedence protects traders who expose their quotations and limit orders. The so-called quote matchers may quote on the same side of the market when they see large size displayed to take advantage of the revealed information. The success of the quote matchers' strategy hinges on getting their orders filled ahead of the large size. Time precedence and a large minimum price variation then protect traders who display their size by forcing quote matchers to improve price significantly if they wish to acquire precedence. Displayed market depth should be positively related to minimum price variation.

Harris' (1994) discrete model for bid-ask spreads assumes that the discrete quote generation process stems from the rounding of random draws from a continuous distribution. At time t, an unobserved relative spread s_t is drawn from a continuous cumulative distribution function $F(s_t; m_i, \theta)$, where m_i is the mean unrounded relative spread and θ is a distributional shape parameter. The unrounded relative spread is then multiplied by the price level to obtain the unrounded absolute spread, which produces the observed absolute discrete spread after rounding to the nearest 1/8 (in the logarithmic metric). Harris models the mean unrounded relative spread m_i as a function of the price level and volatility, trade intensity, and the market value and volume. Finally, the unrounded relative spread is then assumed to follow a gamma distribution with θ, yielding

$$\begin{aligned}
\beta_1 &= \Pr(s_t = 1/8) = F(\kappa_1/P) \\
\beta_2 &= \Pr(s_t = 1/4) = F(\kappa_2/P) - F(\kappa_1/P) \\
\beta_3 &= \Pr(s_t = 3/8) = F(\kappa_3/P) - F(\kappa_2/P) \\
\beta_4 &= \Pr(s_t \geq 1/2) = 1 - F(\kappa_3/P),
\end{aligned}$$

where P denotes the price level and the geometric rounding midpoints are given by $\kappa_i = \sqrt{i/8(i+1)/8}$. The motivation for aggregating the observed discrete spreads larger than \$1/2 stems from the fact that fewer than 5 percent of the quotes for stocks over \$40 are larger than \$1/2, and far less than 1 percent are larger for the lower-priced stocks.

The resulting pooled time-series cross-sectional model is then estimated by maximum likelihood. The multinomial log-likelihood for a given stock data vector reads $\log L = T \sum_{i=1}^{4} \hat{\beta}_i \log \beta_i$, where T is the number of time-series observations, and $\hat{\beta}_i$ is the observed frequency of spreads at \$i/8. To offer projections for the usage frequency of \$1/16 spreads given a minimum price variation of \$1/16, Harris (1994) computes $\Pr(s_t = 1/16) = F(\kappa_0/P)$ from the estimated model with $\kappa_0 = \sqrt{(1/16)(1/8)}$. As a further correction, the estimated inverse price coefficient that appears in m_i is halved to reflect the smaller tick.

Conditional duration models

Parallel to the research on market microstructure theory, the advent of high frequency databases fostered the development of suitable econometric model and techniques. The first generation of models take advantage of the GARCH and stochastic volatility types of models, where the sampling interval is fixed and observations are equally spaced in time. The second generation recognizes the fact that transaction data are not necessarily equally spaced in time. Engle and Russell (1998) propose the autoregressive conditional duration (ACD, for short) model for durations between two successive market events (e.g. price changes), which Engle (1996) combines with a GARCH process for the returns to model irregularly spaced data.

Let $x_i = \psi_i \epsilon_i$, where the duration $x_i = t_i - t_{i-1}$ denotes the time elapsed between events occurring at time t_i and t_{i-1}, the conditional duration process $\psi_i \equiv E\big(x_i \,|\, I_{i-1}\big)$ is independent of ϵ_i and I_{i-1} is the set including all information available at time t_{i-1}. Engle and Russell (1998) propose a specification for the duration process akin to the GARCH structure, namely

$$\psi_i = \omega + \sum_{j=1}^{m} \alpha_j x_{i-j} + \sum_{j=1}^{q} \beta_j \psi_{i-j} \qquad \sim \mathrm{ACD}(m, q),$$

where $\alpha_j, \beta_j \geq 0$ and $\omega > 0$. One implication is that the unconditional expected mean is simply

$$E(x_i) = \frac{\omega}{1 - \sum_{j=1}^{m} \alpha_j + \sum_{j=1}^{q} \beta_j}.$$

To invoke the maximum likelihood framework, we suppose further that ϵ_i is iid with Burr density

$$f_B\,(\epsilon_i,\, \theta_B) = \frac{\kappa\, \xi_B^{\kappa}\, \epsilon_i^{\kappa-1}}{(1 + \sigma^2\, \xi_B^{\kappa}\, \epsilon_i^{\kappa})^{1+1/\sigma^2}},$$

with $\kappa > \sigma^2 > 0$ and mean

$$\xi_B \equiv \frac{\Gamma\big(1+1/\kappa\big)\,\Gamma\big(1/\sigma^2 - 1/\kappa\big)}{\sigma^{2(1+1/\kappa)}\,\Gamma\big(1+1/\sigma^2\big)}.$$

It is readily seen that the conditional density of x_i is also Burr with parameter vector $\big(\xi_B^\kappa \psi_i^{-\kappa}, \kappa, \sigma^2\big)$. Accordingly, the conditional hazard rate function reads

$$\Gamma_B\left(x_i \,\big|\, I_{i-1};\, \theta_B\right) = \frac{\kappa\,\xi_B^\kappa\,\psi_i^{-\kappa}\,x_i^{\kappa-1}}{1+\sigma^2\,\xi_B^\kappa\,\psi_i^{-\kappa}\,x_i^\kappa},$$

which is nonmonotonic with respect to the standardized duration if $\kappa > 1$.

When σ^2 shrinks to zero, the Burr reduces to a Weibull distribution, viz.

$$f_W\left(\epsilon_i,\, \theta_W\right) = \kappa\,\xi_W^\kappa\,\epsilon_i^{\kappa-1}\exp\left(-\xi_W^\kappa\,\epsilon_i^\kappa\right),$$

where $\xi_W = \Gamma(1+1/\kappa)$. Accordingly, the conditional distribution of the duration process is also Weibull and the conditional hazard rate function reads $\Gamma_W\left(x_i \,\big|\, I_{i-1};\, \theta_W\right) = \kappa\,\xi_W^\kappa\;\psi_i^{-\kappa}x_i^{\kappa-1}$. In contrast to the Burr case, the conditional hazard rate implied by the Weibull distribution is monotonic. It decreases with the standardized duration for $0 < \kappa < 1$, increases for $\kappa > 1$ and remains constant for $\kappa = 1$. In the latter case, the Weibull coincide with the exponential distribution and the conditional hazard rate function of the duration process is simply $\Gamma_E\left(x_i \,\big|\, I_{i-1};\, \theta_E\right) = \psi_i^{-1}$.

As an alternative, Lunde (1999) employs the generalized gamma ACD model in which ϵ_i is iid with density

$$f_G(\epsilon_i, \theta_G) = \frac{\xi_G^{\kappa\gamma}\kappa\epsilon_i^{\kappa\gamma-1}}{\Gamma(\gamma)}\exp\left(-\xi_G^\kappa\epsilon_i^\kappa\right)$$

where $\xi_G \equiv \Gamma(\gamma+1/\kappa)/\Gamma(\gamma)$. The generalized gamma distribution nests both the exponential ($\kappa = \gamma = 1$) and the Weibull ($\gamma = 1$) distributions, though it is nonnested with respect to the Burr distribution. The baseline hazard rate has no closed-form solution because it depends on the incomplete gamma integral $I(\epsilon_i, \gamma) \equiv \int_0^{\epsilon_i} u^{\gamma-1}\exp(-u)\,\mathrm{d}u$. Nonetheless, it is possible to derive its shape properties according to the parameter values (Glaser, 1980). If $\kappa\gamma < 1$, the hazard rate is decreasing for $\kappa \leq 1$, and U-shaped for $\kappa > 1$. Conversely, if $\kappa\gamma > 1$, the hazard rate is increasing for $\kappa \geq 1$, and inverted U-shaped for

$\kappa < 1$. Lastly, if $\kappa\gamma = 1$, the hazard rate is decreasing for $\kappa < 1$, constant for $\kappa = 1$ (exponential case), and increasing for $\kappa > 1$.

Albeit Engle and Russell (1998) suggest the use of exponential and Weibull distributions, the Burr and the generalized gamma ACD models seem to deliver better results for both price and transaction durations (Bauwens, Giot, Grammig and Veredas, 2000; Lunde, 1999; Zhang, Russell and Tsay, 2001).

Extensions and applications. To test for microstructure effects, one could include exogenous variables into the duration process. For that purpose, it is more convenient to work with the log-ACD specification advanced by Bauwens and Giot (2000) so as to avoid the nonnegativeness issue. Ghysels, Gouriéroux and Jasiak (1997) introduce a stochastic volatility duration model to cope with higher order dynamics in the duration process. Ghysels and Jasiak (1998) investigate the persistence of intra-trade durations using a fractionally integrated ACD model, whilst Zhang et al. (2001) advocate for a nonlinear version of the ACD model rooted in a self exciting threshold autoregressive framework. Bauwens and Veredas (1999), Grammig and Maurer (2000), Lunde (1999), and Hamilton and Jorda (1999) argue for conditional duration models that accommodate more flexible hazard rate functions. Bauwens and Giot (1998) and Russell and Engle (1998) propose extensions to deal with competing risks, whereas Russell (1998) and Engle and Lunde (1998) consider bivariate models for trade and quote processes. As price durations are closely linked to the instantaneous volatility of the mid-quote price process (Engle and Russell, 1997), they play an interesting role in option pricing (Pringent, Renault and Scaillet, 1999). Trade and volume durations mirror in turn features such as market liquidity and the information arrival rate (Gouriéroux, Jasiak and Le Fol, 1999).

References

Baghehot, W., 1971, The only game in town, Financial Analysts Journal 27, 12–14.

Ball, C. A., Torous, W. A., Tschoegl, A. E., 1985, The degree of price resolution: The case of the gold market, Journal of Futures Markets 5, 29–43.

Bauwens, L., Giot, P., 1998, Asymmetric ACD models: Introducing price information in ACD models with a two-state transition model, Université Catholique de Louvain.

Bauwens, L., Giot, P., 2000, The logarithmic ACD model: An application to the bid-ask quote process of three NYSE stocks, Annales d'Economie et de Statistique 60, 117–150.

Bauwens, L., Giot, P., Grammig, J., Veredas, D., 2000, A comparison of financial duration models via density forecasts, Université Catholique de Louvain and University of Frankfurt.

Bauwens, L., Veredas, D., 1999, The stochastic conditional duration model: A latent factor model for the analysis of financial durations, Université Catholique de Louvain.

Campbell, J. Y., Lo, A. W., MacKinlay, A. C., 1997, The Econometrics of Financial Markets, Princeton University Press, Princeton.

Easley, D., Kiefer, N. M., O'Hara, M., 1997a, The information content of the trading process, Journal of Empirical Finance 4, 159–186.

Easley, D., Kiefer, N. M., O'Hara, M., 1997b, One day in the life of a very common stock, Review of Financial Studies 10, 805–835.

Easley, D., Kiefer, N. M., O'Hara, M., Paperman, J. B., 1996, Liquidity, information, and infrequently traded stocks, Journal of Finance 51, 1405–1436.

Easley, D., O'Hara, M., 1987, Price, trade size, and information in security markets, Journal of Financial Economics 19, 69–90.

Easley, D., O'Hara, M., 1992, Time and the process of security price adjustment, Journal of Finance 47, 577–605.

Easley, D., O'Hara, M., Paperman, J. B., 1998a, Financial analysts and information-based trade, Journal of Financial Markets 1, 175–201.

Easley, D., O'Hara, M., Srinivas, P. S., 1998b, Option volume and stock prices: Evidence on where informed traders trade, Journal of Finance 53, 431–465.

Engle, R. F., 1996, The econometrics of ultra-high frequency data, Econometrica 68, 1–22.

Engle, R. F., Lunde, A., 1998, Trades and quotes: A bivariate point process, Economics Discussion Paper 9807, University of Califonia at San Diego.

Engle, R. F., Russell, J. R., 1997, Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model, Journal of Empirical Finance 4, 187–212.

Engle, R. F., Russell, J. R., 1998, Autoregressive conditional duration: A new model for irregularly-spaced transaction data, Econometrica 66, 1127–1162.

Ghysels, E., Gouriéroux, C., Jasiak, J., 1997, Stochastic volatility duration models, CREST Working Paper 9746, INSEE.

Ghysels, E., Jasiak, J., 1998, Long-term dependence in trading, Penn State University and York University.

Glaser, R. E., 1980, Bathtub and related failure rate characterizations, Journal of the American Statistical Association 75, 667–672.

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

Gouriéroux, C., Jasiak, J., Le Fol, G., 1999, Intra-day market activity, Journal of Financial Markets 2, 193–226.

Grammig, J., Maurer, K.-O., 2000, Non-monotonic hazard functions and the autoregressive conditional duration model, Econometrics Journal.

Grammig, J., Schiereck, D., Theissen, E., 2000, Knowing me, knowing you: Trader anonymity and informed trading in parallel markets, Journal of Financial Markets. Forthcoming.

Hamilton, J. D., Jorda, O., 1999, A model for the federal funds rate target, University of California at San Diego and University of Califonia at Davis.

Harris, L., 1991, Stock price clustering and discreteness, Review of Financial Studies 4, 389–415.

Harris, L., 1994, Minimum price variations, discrete bid-ask spreads, and quotation size, Review of Financial Studies 7, 149–178.

Hasbrouck, J., 1991, Measuring the information content of stock trades, Journal of Finance 46, 179–207.

Hausman, J. A., Lo, A. W., 1990, A continuous-time discrete-state stochastic process for transaction stock prices: Theory and empirical estimates, MIT.

Hausman, J. A., Lo, A. W., MacKinlay, A. C., 1992, An ordered probit analysis of transaction stock prices, Journal of Financial Economics 31, 319–379.

Lamoureaux, C., Lastrapes, W., 1990, Heteroscedasticity in stock return data: Volume versus GARCH effects, Journal of Finance 45, 221–229.

Lo, A. W., MacKinlay, A. C., 1990, An econometric analysis of nonsynchronous trading, Journal of Econometrics 45, 181–212.

Lunde, A., 1999, A generalized gamma autoregressive conditional duration model, University of Aarhus.

Marsh, T., Rosenfeld, E., 1986, Non-trading, market making, and estimates of stock price volatility, Journal of Financial Economics 15, 359–372.

O'Hara, M., 1995, Market Microstructure Theory, Blackwell Publishers, Cambridge.

Pringent, J.-L., Renault, O., Scaillet, O., 1999, An autoregressive conditional binomial option pricing model, Université de Cergy-Pontoise, CREST, and Université Catholique de Louvain.

Roll, R., 1984, A simple implicit measure of bid-ask spread in an efficient market, Journal of Finance 39, 1127–1140.

Russell, J. R., 1998, Econometric modeling of multivariate irregularly-spaced high-frequency data, Graduate School of Business, University of Chicago.

Russell, J. R., Engle, R. F., 1998, Econometric analysis of discrete-valued irregularly-spaced financial transaction data using a new autoregressive conditional multinomial model, Discussion Paper 9810, University of Califonia at San Diego.

Zhang, M. Y., Russell, J. R., Tsay, R. S., 2001, A nonlinear autoregressive conditional duration model with applications to financial transaction data, Journal of Econometrics 104, 179–207.