

Modeling and forecasting persistent financial durations

Filip Žikeš, Jozef Baruník & Nikhil Shenai

To cite this article: Filip Žikeš, Jozef Baruník & Nikhil Shenai (2017) Modeling and forecasting persistent financial durations, *Econometric Reviews*, 36:10, 1081-1110, DOI: [10.1080/07474938.2014.977057](https://doi.org/10.1080/07474938.2014.977057)

To link to this article: <http://dx.doi.org/10.1080/07474938.2014.977057>



View supplementary material [↗](#)



Accepted author version posted online: 07
Nov 2014.
Published online: 07 Nov 2014.



Submit your article to this journal [↗](#)



Article views: 129



View related articles [↗](#)



View Crossmark data [↗](#)



Modeling and forecasting persistent financial durations

Filip Žikeš^a, Jozef Baruník^b, and Nikhil Shenai^c

^aBank of England, London, UK; ^bInstitute of Economic Studies, Charles University, Prague and Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Prague, Czech Republic; ^cImperial College London, Business School, London, UK

ABSTRACT

This article introduces the Markov-Switching Multifractal Duration (MSMD) model by adapting the MSM stochastic volatility model of Calvet and Fisher (2004) to the duration setting. Although the MSMD process is exponential β -mixing as we show in the article, it is capable of generating highly persistent autocorrelation. We study, analytically and by simulation, how this feature of durations generated by the MSMD process propagates to counts and realized volatility. We employ a quasi-maximum likelihood estimator of the MSMD parameters based on the Whittle approximation and establish its strong consistency and asymptotic normality for general MSMD specifications. We show that the Whittle estimation is a computationally simple and fast alternative to maximum likelihood. Finally, we compare the performance of the MSMD model with competing short- and long-memory duration models in an out-of-sample forecasting exercise based on price durations of three major foreign exchange futures contracts. The results of the comparison show that the MSMD and the Long Memory Stochastic Duration model perform similarly and are superior to the short-memory Autoregressive Conditional Duration models.

KEYWORDS

Long memory; multifractal models; price durations; realized volatility; whittle estimation

JEL CLASSIFICATION

C13; C58; G17

1. Introduction

Financial durations measure the time elapsed between various financial market events related to transactions arrivals, price fluctuations, or trading volumes. Modeling durations may be useful for measuring and predicting instantaneous volatility and integrated variance and so may aid high-frequency volatility trading and risk management. Exploiting the intimate relationship between durations and volatility, Tse and Yang (2010) employ parametric duration models to measure daily volatility using high-frequency data. Andersen et al. (2008) propose a nonparametric duration-based approach to measuring volatility by relying on the properties of Brownian motion. More generally though, durations are useful for gaining more insight into any information events or variables which change values at each tick, as implied by the theory of market microstructure, and thus may be useful for examining a number of interesting economic hypotheses related to trading and price discovery; see Engle (2000) for an excellent discussion.

A key stylized fact noted in the empirical irregularly-spaced event literature is persistence in financial durations. Ever since the seminal contribution of Engle and Russell (1998), who introduced the first time-series model for financial durations, a number of studies have documented the slowly decaying autocorrelation function of transaction, price and volume durations; see Pacurar (2008) for a detailed literature review. Deo et al. (2010) recently test for long memory in durations and the associated counts and find significant evidence to support the presence of long memory in durations. Despite this empirical

regularity, there is currently no article that explores the alternative approaches to capturing the persistent autocorrelations of durations and its implications for forecasting. We aim to fill this gap.

Inspired by the success of the Markov-Switching Multifractal (MSM) stochastic volatility model of Calvet and Fisher (2004) in forecasting persistent volatility of financial returns, we start by adapting the MSM model to the duration setting, calling the new model the Markov-Switching Multifractal Duration (MSMD) model. This model adds to the class of stochastic durations models of Bauwens and Veredas (2004) and Deo et al. (2010), which also evolved from the stochastic volatility literature, though the latent process driving the dynamics of durations in an MSMD is a hidden Markov chain rather than a linear Gaussian process. The properties of the model are therefore quite distinct from its competitors and exhibit a number of attractive features.

First, the process possesses a very flexible autocovariance structure, which is capable of generating a wide range of persistence behaviors, including a slowly decaying autocorrelation. This long-memory feature of the MSMD process is induced by regime switching of heterogeneous persistence: the hidden Markov chain is driven by k independent Markov-switching processes with different transition probabilities. Second, despite being able to exhibit high persistence, we show formally that the MSMD process is exponential β -mixing and short-memory. This property of the model is unique: all the existing duration models, whether observation-driven or stochastic, exhibit either genuine exponential or hyperbolic autocovariance functions (ACF). The MSMD can produce an ACF that is nearly hyperbolic over a range of lags before smoothly transitioning to exponential decay. Finally, the flexibility of MSMD model is not traded off against parsimony. Although the hidden chain consists of k independent components with distinct transitional probabilities, these probabilities are tightly parametrized, and the number of parameters of the model does not depend on k . Standard methods and tools for Markov switching models therefore readily apply.

Since durations and volatility are directly linked, we next explore the properties of volatility in a pure jump model of Oomen (2006), where the inter-jump durations follow MSMD. Relying on the recent results of Deo et al. (2009) on the propagation of memory from durations to counts and realized volatility, we establish formally that the short memory of MSMD durations translates into short memory in counts and realized volatility. However, simulation reveals that the shape of the ACF of MSMD durations—hyperbolic over a range of lags but eventually exponential—is inherited by realized volatility. Again this feature cannot be generated by the existing durations models.

Having studied the properties of the MSMD process, we next propose quasi-maximum likelihood estimation of the MSMD parameters based on the Whittle approximation. The main motivation for exploring this estimation method as an alternative to exact maximum likelihood is computational burden associated with the latter in large samples, and its limitation to the case of an MSM specification with a finite number of states. Contrary to this, the Whittle estimator works in either case and is computationally simple and fast. Relying on results from the statistics literature, we formally establish strong consistency and asymptotic normality of the Whittle estimator under fairly mild assumptions for a wide range of MSMD specifications.

Note that computational speed is not a mere convenience in our context: given the increasing importance of algorithmic and high-frequency traders, who are capable of generating tens of thousands of limit and market orders in a single day, the amount of data usable for estimation has grown enormously in many markets (Hasbrouck and Saar, 2010). For such environments, fast estimation methods simply become a necessity, even with ever-faster modern computers. Last but not least, the Whittle estimator can be easily adapted to the original MSM stochastic volatility model of Calvet and Fisher (2004), and thus represents a contribution to the MSM literature that goes beyond the context of financial durations.

Finally, we compare our estimation and forecasting results with those possible from established duration models. As noted by Pacurar (2008), there is a scarcity of comparisons of duration models, and ideally one would like to undertake a comparison of all the models she has detailed. However, as noted above, only long memory models are able to account for the key stylised fact of long-range dependence in durations. We therefore restrict attention to the Long Memory Stochastic Duration (LMSD) model of Deo et al. (2010). To investigate the benefits of the relatively complicated MSMD and LMSD models over

their simple and easy to estimate short-memory counterparts, we also compare our results with those from the widely used Autoregressive Conditional Duration (ACD) model introduced in the seminal paper by Engle and Russell (1998).

We implement the models on price durations of three major foreign exchange futures contracts traded on the Chicago Mercantile Exchange (CME) in the period between November 9, 2009 and January 29, 2010: Euro, Japanese Yen, and Swiss Franc. We find that the LMSD and MSMD models generally perform on par, although MSMD sometimes produces better results. The forecast combinations of the LMSD and MSMD models almost always significantly outperform those of the short-memory ACD model. Given that we provide a simple and computationally cheap method for estimating the MSMD model parameters from large samples of data, this results is potentially important for practitioners. Needless to say, our results are limited to three FX futures contracts and a more comprehensive empirical exercise is needed to confirm a general validity of our results. Chen et al. (2012) (henceforth CDS) confirm the superiority of MSMD over ACD for a sample of U.S. equity transactions durations, though they do not include the LMSD model in the comparison.

The Markov switching multifractal duration model has been proposed independently and in parallel to our work in a recent paper by CDS articles. The main differences between the two articles can be summarized as follows. We are not restricting attention to the binomial MSMD model with exponentially distributed innovations, but consider more general versions of the model. Allowing for a wider class of distributions is made possible in practice by employing the Whittle estimator, and it turns out to be empirically beneficial. In terms of empirical application, we differ from CDS by modeling and forecasting price durations as opposed to transactions durations, and focus on foreign exchange futures prices in 2009/2010 rather than individual equities in 1993. Finally, given the high persistence of the durations in our sample, the natural competitor of the MSMD model is the LMSD model rather than the short-memory ACD, and hence, unlike CDS, we include the LMSD model in our forecasting exercise as well.

The rest of the article is organized as follows. Section 2 introduces the MSMD model and discusses its properties. Section 3 discusses estimation and forecasting for the MSMD model. Section 4 reviews the competing duration models and Section 5 looks at the link between durations, counts, and realized volatility. In Section 6, we describe the data, and in Section 7, we present estimation and forecasting results. Section 8 concludes. Mathematical proofs are provided in the Appendix. Some auxiliary results and derivations are collected in the Supplemental Appendix.

2. The MSMD model

Let $X_i = t_i - t_{i-1}$ denote the duration between two event arrival times. The three most common events studied in the literature relate to transaction arrivals, price changes, and transaction volumes. The MSMD model is defined by

$$X_i = \psi_i \epsilon_i, \quad i \in \mathbb{Z}, \quad (2.1)$$

where ψ_i is the Markov-switching multifractal process of Calvet and Fisher (2004);

$$\psi_i = \bar{\psi} \prod_{j=1}^k M_{j,i}, \quad (2.2)$$

and ϵ_i is a sequence of independent unit-mean innovations identically distributed according to some parametric distribution. The latent process in (2.2) is determined by k independent unit-mean multipliers, $M_{j,i}$, $j = 1, \dots, k$, and a scaling constant, $\bar{\psi}$. At every point in time i , each multiplier $M_{j,i}$ takes, with probability γ_j , a new value M drawn from a common distribution F_M , and remains unchanged with probability $1 - \gamma_j$:

$$M_{j,i} = \begin{cases} M & \text{where } M \text{ is drawn from } F_M \text{ with probability } \gamma_j \\ M_{j,i-1} & \text{with probability } 1 - \gamma_j \end{cases}.$$

The transition probabilities are parsimoniously parametrized by

$$\gamma_j = 1 - (1 - \gamma_k)^{(b_j - k)}, \quad j = 1, \dots, k, \tag{2.3}$$

where $\gamma_k \in (0, 1)$ and $b \in (1, \infty)$. Two specifications for the distribution of the multipliers F_M have been proposed by Calvet and Fisher (2004): binomial and log-normal. In the binomial specification, each multiplier, if at all, is renewed by drawing the values m_0 and $2 - m_0$; $m_0 \in (1, 2)$, with equal probability, ensuring that the mean is equal to one. The transition matrix associated with each multiplier is thus given by

$$P_j = \begin{pmatrix} 1 - \frac{1}{2}\gamma_j & \frac{1}{2}\gamma_j \\ \frac{1}{2}\gamma_j & 1 - \frac{1}{2}\gamma_j \end{pmatrix}.$$

Since the multipliers are independent, the transition matrix of the state vector $\mathbf{M}_i = (M_{1,i}, \dots, M_{k,i})$ is simply $\mathbf{P} = \mathbf{P}_1 \otimes \mathbf{P}_2 \otimes \dots \otimes \mathbf{P}_k$, where “ \otimes ” denotes the Kronecker product. The dimension of the transition matrix is $2^k \times 2^k$ and the state vector takes values in the finite state space $\Omega_M = \{m_0, 2 - m_0\}^k$.

In general, any distribution with positive support can be used to model the multipliers. For example, the log-normal specification of Calvet and Fisher (2004) replaces the Bernoulli distribution by a log-normal one, i.e., upon switching, the new value of the log multiplier is drawn from $N(-\lambda, 2\lambda)$, where the parameter restriction again imposes unit means for the multipliers. When drawn from a continuous distribution (with respect to Lebesgue measure on \mathbb{R}_+), each multiplier assumes a new value with probability one, and the transition kernel of the multiplier is given by

$$\mathbb{P}(M_{j,i+h} \in B_j | M_{j,i} = x_j) = (1 - (1 - \gamma_j)^h) \mathbb{P}(M \in B_j) + (1 - \gamma_j)^h \mathbf{1}_{\{x_j \in B_j\}}, \quad j = 1, \dots, k, \tag{2.4}$$

for any $B_j \in \mathcal{B}(\mathbb{R}_+)$ and $x_j \in \mathbb{R}_+$, where $\mathcal{B}(\mathbb{R}_+)$ is the Borel σ -algebra on \mathbb{R}_+ . Since the multipliers are independent, the transition kernel of the chain \mathbf{M}_i reads $\mathbb{P}(\mathbf{M}_{i+h} \in B | \mathbf{M}_i = \mathbf{x}) = \prod_{j=1}^k \mathbb{P}(M_{j,i+h} \in B_j | M_{j,i} = x_j)$ for any $\mathbf{x} = (x_1, x_2, \dots, x_k)'$ and any $B \in \mathcal{B}(\mathbb{R}_+^k)$, a Borel σ -algebra on \mathbb{R}_+^k , where $B = B_1 \times B_2 \times \dots \times B_k$, $B_j \in \mathcal{B}(\mathbb{R}_+)$, $j = 1, \dots, k$. The chain takes values in a state space $\Omega_M \subseteq \mathbb{R}_+^k$.

Having specified the law governing the multipliers, it remains to choose a distribution for the innovations, ϵ_i . As is common in the literature, we consider here the exponential and Weibull distributions (Supplemental Appendix A2). With these specifications of ϵ_i , the law governing the durations, x_i , is a mixture of exponentials and a mixture of Weibull distributions, respectively. Other, more flexible multiparameter alternatives have been proposed in the context of modeling financial durations: the Burr distribution (Grammig and Maurer, 2000) and the generalized gamma distribution (Lunde, 1999), both of which encompass the exponential and Weibull cases. As we are primarily interested in point forecasts in this article, for the sake of parsimony we confine our attention to the latter two distributions.

To illustrate the behavior of the multipliers and durations in the MSMD model, we plot in Fig. 1 simulated samples from the binomial and log-normal MSMD processes with $k = 6$ multipliers and parameters $b = 3$, $\gamma_k = 0.5$, $m_0 = 1.4$, and $\lambda = 0.15$. In this MSMD specification, $\gamma_1 = 0.0028$ implies that the most persistent multiplier, $(M_{1,i})$, switches, on average, around 3 times in a sample of 1,000 observations if it is drawn from the log-normal distribution, and 1.5 times if it is drawn from the Bernoulli distribution. The least persistent multiplier, $(M_{6,i})$, switches with probability 0.5 and 0.25 in the log-normal and binomial MSMD specifications, respectively. Clearly, both specifications can produce rich dynamics: the duration process is highly persistent but can exhibit sudden erratic movements as observed in empirical data.

2.1. Stationarity, ergodicity and strong mixing

It is relatively easy to establish that the Markov chain \mathbf{M}_i driving the MSMD process is geometrically ergodic as long as the conditions $b > 1$ and $0 < \gamma_k < 1$ are satisfied. Starting with the binomial MSMD specification, we see that under these conditions all elements of the transition matrix of the chain \mathbf{P}

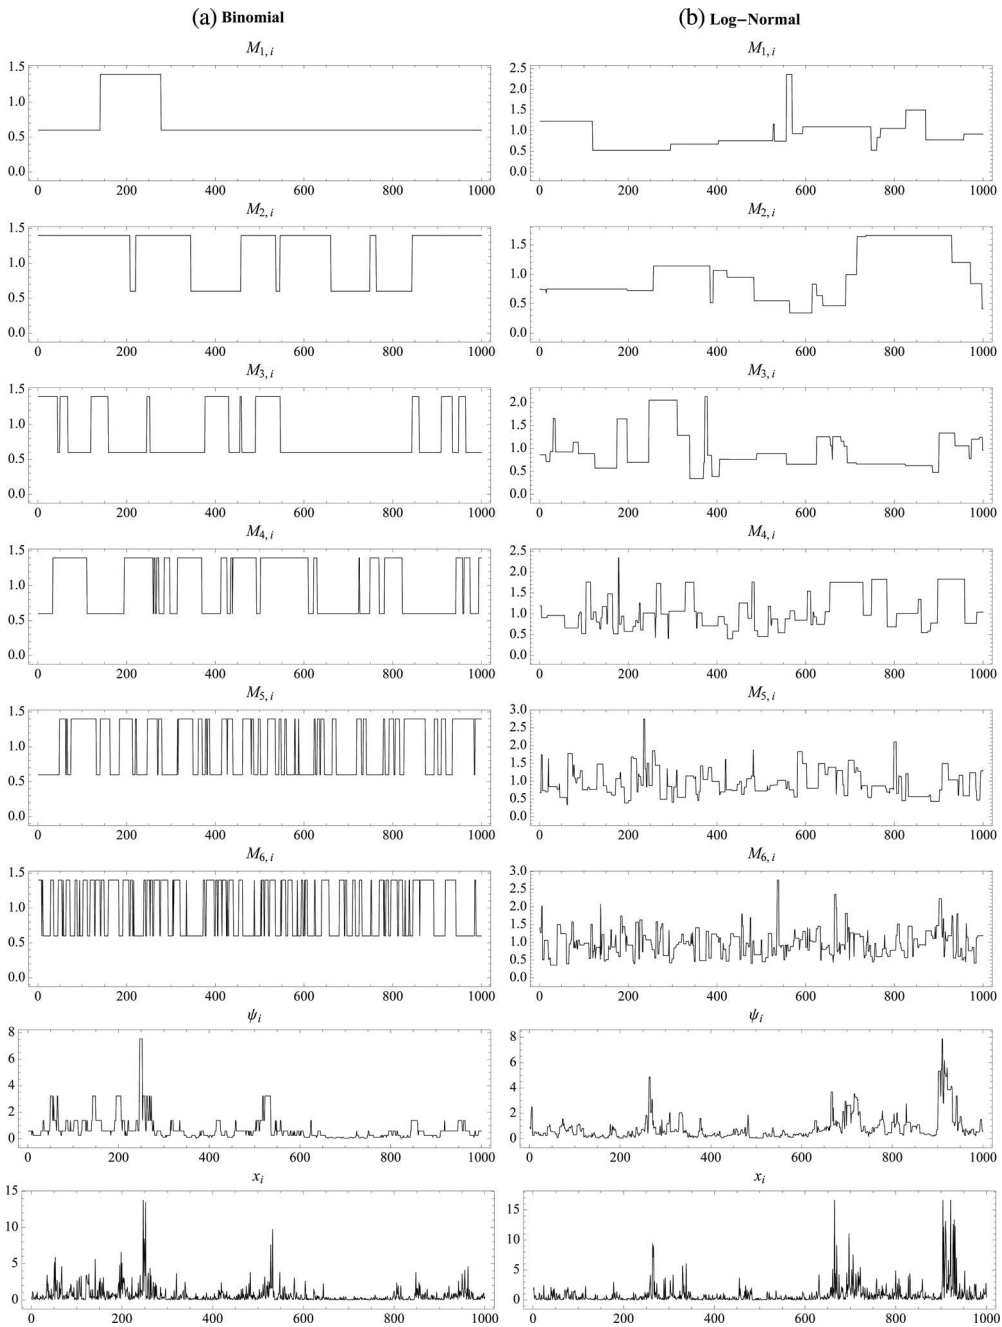


Figure 1. Simulated binomial and log-normal MSMD processes with six multipliers and exponentially distributed innovations. The parameters of the processes are $b = 3$, $\gamma_k = 0.5$, $m_0 = 1.4$, and $\lambda = 0.15$.

are strictly positive since $0 < \gamma_j < 1$ for all j , and it follows directly from the proof of Theorem 1 in (Shiryayev, 1995, Chapter 1, Section 12) that the chain is geometrically ergodic. The ergodic distribution is given by $\pi_l = 1/2^k, l = 1, \dots, 2^k$.

If upon switching the multipliers, $M_{j,i}, j = 1, \dots, k$, are drawn randomly from a continuous distribution, F_M , with support \mathbb{R}_+ , the transition kernel associated with the j th multiplier is given in Eq. (2.4) and the ergodic distribution of the multiplier reads $\pi(B_j) := \lim_{h \rightarrow \infty} \mathbb{P}(M_{j,i+h} \in B_j | M_{j,i} = X_j) = \mathbb{P}(M \in B_j), B_j \in \mathcal{B}(\mathbb{R}_+)$. Then for any $X_j \in \mathbb{R}_+, j = 1, \dots, k$ and $h \in \mathbb{N}, \sup_{B_j \in \mathcal{B}(\mathbb{R}_+)} |\mathbb{P}(M_{j,i+h} \in B_j | M_{j,i} = X_j) - \pi(B_j)| \leq (1 - \underline{\gamma})^h$, where $0 < \underline{\gamma} := \min\{\gamma_1, \dots, \gamma_k\} < 1$. Since the multipliers $M_{j,i}, j = 1, \dots, k$, are independent it follows that the chain M_i is geometrically ergodic.

Geometric ergodicity of the Markov chain M_i in turn implies that the duration process $\{X_i\}$ is strictly stationary β -mixing with an exponential rate of decay, provided that the chain is initialized from the ergodic distribution. To see this, observe that the duration process belongs to the class of generalized hidden Markov models in the sense of Definition 3 in Carrasco and Chen (2002): the hidden Markov chain M_i is strictly stationary and, conditionally on M_i , the durations X_i are independently distributed where the conditional distribution only depends on M_i and not on i . Given geometric ergodicity of the hidden chain, Proposition 4 of Carrasco and Chen (2002) then implies that the duration process is exponential β -mixing.

2.2. Moments, autocovariance function, and spectral density

In the Supplemental Appendix A.1, we show that the first two moments of the MSMD process are given by

$$E(X_i) = \bar{\psi}, \tag{2.5}$$

$$\text{Var}(X_i) = \bar{\psi}^2 [E(M^2)^k E(\epsilon_1^2) - 1]. \tag{2.6}$$

The model can exhibit both under- and overdispersion depending on the distributional assumptions about M and ϵ_i , since the ratio of the variance to the squared mean, $E(M^2)^k E(\epsilon_1^2) - 1$, can in general be smaller or larger than one. An MSMD process with exponential innovations, however, always exhibits over-dispersion since for an exponentially distributed ϵ_i , we have $E(\epsilon_1^2) = 2$, and by construction $E(M^2) > 1$.

An attractive property of the MSMD model is that it possesses a very flexible ACF that can exhibit behavior similar to long-memory. In the Supplemental Appendix A.1, we show that for a general MSMD process with finite $E(M^2)$ and $E(\epsilon_1^2)$, we have

$$\text{Cov}(X_i, X_{i-h}) = \bar{\psi}^2 \left(\prod_{j=1}^k [1 + \text{Var}(M)(1 - \gamma_j)^h] - 1 \right). \tag{2.7}$$

Although (2.7) implies that the MSMD process is short-memory as the autocovariance function declines exponentially fast and the spectral density is bounded at origin, it is capable of mimicking hyperbolic decay over a wide range lags. More specifically, it follows directly from Proposition 1 in Calvet and Fisher (2004) that the autocorrelation function of the MSMD durations decays hyperbolically over a large range of lags before transitioning smoothly into exponential decay as the number of multipliers, k , grows without bound. So despite being a short-memory process, the MSMD model can mimic the persistence of a genuine long-memory process with a hyperbolically decaying autocorrelation function.

For illustration purposes, Fig. 2 plots the autocorrelation function of a binomial MSMD process with exponential innovations and various sets of parameter values. We take the case of $k = 8$ multipliers and parameters $b = 2, \gamma_k = 0.5$, and $m_0 = 1.4$, as a benchmark and vary each parameter separately to study how it affects the shape of the autocorrelation function. Increasing b or decreasing γ_k both increase the persistence of the process since the switching probabilities of the multipliers decrease (panels (a) and (b)). In the former case, the increase is more pronounced at the long end of the ACF, while in the latter case it affects the short lags of the ACF more. This is due to the different impact of a change in b and γ_k on

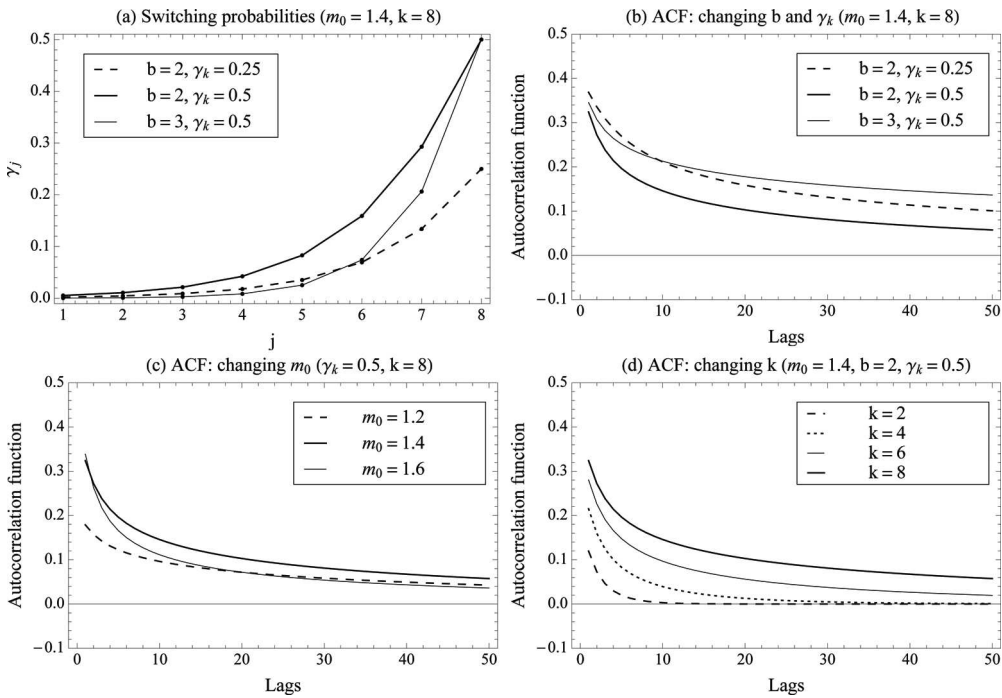


Figure 2. Transition probabilities and the autocorrelation function of a binomial MSMD process with exponentially distributed innovations.

the various switching probabilities as illustrated in panel (a). Increasing the volatility of the multipliers by reducing m_0 lowers the multipliers' persistence and thus the persistence of the MSMD process (panel (c)). Finally, increasing the number of multipliers (k) while keeping the parameters of the model fixed increases persistence (panel (d)).

2.3. Exogenous and predetermined variables

Exogenous or predetermined variables can be easily incorporated into the model by setting $\bar{\psi} = \bar{\psi}_i = \exp(\beta_0 + \beta'z_i)$, for some vector of variables z_i . This is useful for several reasons. First, to incorporate the deterministic intraday duration pattern observed in most durations data (Engle and Russell, 1998; Bauwens and Veredas, 2004; Fernandes and Grammig, 2006; Deo et al., 2010, among many others). Due to the deterministically varying trading activity during the day, the durations tend to be shorter during the early and late trading hours, and relatively longer over lunchtime. Second, one may wish to include additional predictive variables to enhance the forecasting power of the model. A natural candidate when forecasting price durations may be option-implied volatility for which high-frequency data is either available readily (e.g., VIX) or can be constructed from high-frequency options data. Finally, it may be interesting to include some predetermined variables related to market microstructure as in Engle and Russell (1998), Bauwens and Veredas (2004), and others.

3. Estimation, inference and forecasting

3.1. Maximum likelihood and optimal forecasting

The binomial MSM with finite k implies a finite number of states of the hidden Markov process and hence can be estimated by exact maximum likelihood (MLE) via Bayesian updating. This approach has been advocated by Calvet and Fisher (2004) for the binomial MSM model of stochastic volatility, and has been shown to work well for sample sizes typically used for estimating models of time-varying volatility.

Moreover, the Bayesian filter allows for estimation of the unobserved state probabilities, which in turn permits optimal forecasting. To save space, we omit the details here and refer the reader to Calvet and Fisher (2004).

A disadvantage of the exact maximum likelihood estimator is that it becomes computationally demanding for $k \geq 10$, since the dimension of the transition matrix grows at a rate of 2^k . Also, it is not applicable to the log-normal MSM process, where the state space of the hidden Markov chain is infinite. These issues have motivated Lux (2008) to develop a generalized method of moments (GMM) approach, which works for a wide range of MSM specifications and requires only moderate computational resources.¹ The drawback of the GMM estimator of Lux (2008) is that it is applied to the first differences rather than levels of the process and this makes the identification of the parameters b and γ_k difficult even when the sample size is very large. Lux (2008) circumvents this problem by setting these parameters to some prespecified values that seem to work well for a number of data sets, and estimates by GMM the remaining two parameters only. This may be quite restrictive, however, especially in our context where no previous evidence exists to suggest reasonable values of b and γ_k for modeling and forecasting financial durations.

3.2. Whittle estimation

We propose an alternative autocovariance-based estimator of the MSMD parameters. In contrast to Lux (2008) and Bacry et al. (2008, 2012), we work in the frequency domain and employ the Whittle quasi-likelihood. An advantage of the Whittle estimation compared to GMM is that it takes into account the entire autocovariance function rather than just a finite subset of lags, and thus avoiding the problem of which autocovariances to match. To obtain better finite-sample properties, we implement the Whittle estimator on logarithmic durations, as the logarithmic durations are much closer to being Gaussian than the durations themselves (see Section 6 for some empirical evidence). Defining $x_i := \log X_i$, $i \in \mathbb{Z}$, and taking logs of both sides of Eq. (2.1), we have

$$x_i = \log \bar{\psi} + \sum_{j=1}^k m_{j,i} + e_i, \quad i \in \mathbb{Z},$$

where $m_{j,i} := \log M_{j,i}$ and $e_i := \log \epsilon_i$. We further define $\sigma_m^2 := \text{Var}(m)$ and $\sigma_e^2 := \text{Var}(e_1)$.

It is well-known that for a stationary Gaussian process maximizing the frequency domain representation of the log-likelihood turns out to be asymptotically equivalent to the usual maximum likelihood estimator (Whittle, 1962). The so-called negative Whittle log-likelihood is given by

$$Q_n(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^{n-1} \left(\log f(\omega_i; \boldsymbol{\theta}) + \frac{I_n(\omega_i)}{f(\omega_i; \boldsymbol{\theta})} \right), \quad (3.1)$$

where $f(\omega_i; \boldsymbol{\theta})$ is the spectral density of the logarithmic MSMD process with parameter $\boldsymbol{\theta}$, i.e., the spectral density associated to $\{x_j\}$ and $I_n(\omega_i) = \frac{1}{2\pi n} \left| \sum_{j=1}^n x_j e^{-i\omega_j} \right|^2$ is the periodogram of the observations x_1, x_2, \dots, x_n , both evaluated at the i th Fourier frequency, $\omega_i = 2\pi i/n$. The Whittle estimator of $\boldsymbol{\theta}$ is obtained by minimizing $Q_n(\boldsymbol{\theta})$:

$$\hat{\boldsymbol{\theta}}_n = \arg \min_{\boldsymbol{\theta} \in \Theta} Q_n(\boldsymbol{\theta}),$$

Now if the process is not Gaussian, which is our case, minimizing the negative Whittle log-likelihood still works but the resulting estimator is no longer asymptotically equivalent to MLE. The intuition for $\hat{\boldsymbol{\theta}}_n$ in the non-Gaussian case is straightforward: under a mixing assumption, the periodogram $I_n(\omega_i)$ is asymptotically distributed as an exponential random variable with parameter $f(\omega_i)$, and for any two

¹ Similarly, Bacry et al. (2008, 2012) use a GMM approach to estimate parameters of the Multifractal Random Walk (MRW).

Fourier frequencies, ω_i and ω_j , $i \neq j$, $I_n(\omega_i)$ and $I_n(\omega_j)$ are asymptotically independent (Rice, 1973). Hence (3.1) has a quasi-likelihood interpretation and $\hat{\theta}_n$ has been shown to be consistent for θ and asymptotically normally distributed under appropriate regularity conditions.

Implementing the Whittle estimator for the MSMD model is easy since the spectral density is available in closed form. In the Supplemental Appendix A.1, we show that provided the logarithmic MSMD durations possess finite second moments, the spectral density reads

$$f(\omega) = \frac{\sigma_m^2}{2\pi} \left(\sum_{j=1}^k \frac{1 - (1 - \gamma_j)^2}{1 + (1 - \gamma_j)^2 - 2(1 - \gamma_j) \cos \omega} \right) + \frac{\sigma_e^2}{2\pi}, \quad (3.2)$$

for $\omega \in [-\pi, \pi]$. In the rest of the article, we will always assume that $E(m^2) < \infty$ and $E(e_1^2) < \infty$ so that (3.2) are well defined.

We see from Eq. (3.2) that the logarithmic MSMD process $\{x_i\}$ is a signal-plus-noise process, where the signal is given by a sum of k independent Markov chains and the noise is an independently and identically distributed (*iid*) process independent of the signal. Whittle estimation of signal-plus-noise models has been studied by Hosoya and Taniguchi (1982) and Zaffaroni (2009). Compared to Whittle estimation of linear processes, a complication arises here from the fact that the spectral density of a signal-plus-noise model cannot be easily factored in the sense that the Whittle log-likelihood cannot be expressed as a sum of two components that depend on disjoint parameter sets. In general, this gives rise to a more complicated limiting distribution of the Whittle estimator.

The asymptotic results obtained by Hosoya and Taniguchi (1982) and Zaffaroni (2009) can be applied in our context despite the fact the both the signal and the noise processes have different specifications in these papers. In case of Hosoya and Taniguchi (1982), the signal is an AR(1) process with *iid* innovations, uncorrelated with, though not necessarily independent of, the *iid* noise process. Zaffaroni (2009) considers a class of models where the signal is an MA(∞) process with *iid* innovations and potentially hyperbolically declining MA coefficients (long memory) and allows for correlation between the signal and the *iid* noise. In case of the logarithmic MSMD, while the signal is independent of the noise, it is not a linear process. Given strict stationarity and ergodicity, which was established in Section 2.1, we can, however, invoke the results of Hannan (1973) and establish strong consistency of $\hat{\theta}_n$ for an MSMD model.

Proposition 1. *Let $\{X_i\}$ be an MSMD process with parameter $\theta_0 \in \Theta$, where Θ is a compact subset of the parameter space such that for all $\theta_1, \theta_2 \in \Theta$, $\theta_1 \neq \theta_2$ implies $f(\omega; \theta_1) \neq f(\omega; \theta_2)$ on a set of positive Lebesgue measure. Then $\hat{\theta}_n \xrightarrow{a.s.} \theta_0$ as $n \rightarrow \infty$.*

The proof is given in section A.1 of the Appendix. The only assumption in Proposition 1 is an identification assumption. Clearly, the Whittle estimator cannot in general work with multiparameter distributions for the multipliers M and innovations ϵ ; it is easy to see from (3.2) that the Whittle estimator can only identify σ_m^2 and σ_e^2 . The functions mapping the parameters of the distribution of M and ϵ into σ_m^2 and σ_e^2 have to be continuous, differentiable, one-to-one and onto. This is clearly satisfied for the Bernoulli, log-normal, and Weibull distributions. In addition, the Whittle estimator cannot identify the mean of the duration process, $\bar{\psi}$, but this is of lesser concern in our application since durations are typically seasonally pre-adjusted and the model is estimated using the seasonally adjusted durations that have unit mean by construction. Nonetheless, the sample mean can be always used to consistently estimate $\bar{\psi}$ if needed.

Turning to the central limit of $\hat{\theta}_n$, we exploit the fact that despite nonlinearity, $\{x_i\}$ has a simple vector MA(∞) representation (see Eq. (A.1) in the Appendix), which allows us to utilize the general results of Hosoya and Taniguchi (1982) provided we verify the relevant regularity conditions. This is done in Section A.2 of the Appendix and proves the following proposition.

Proposition 2. *Let the assumption of Proposition 1 hold with $E(m^{4+r}) < \infty$, $r > 0$ and $E(e_1^4) < \infty$. Then as $n \rightarrow \infty$,*

$$\sqrt{n}(\hat{\theta}_n - \theta_0) \xrightarrow{d} N(0, \mathbf{M}(\theta_0)^{-1} \mathbf{V}(\theta_0) \mathbf{M}(\theta_0)^{-1}), \tag{3.3}$$

where

$$\mathbf{M}(\theta_0) = \frac{1}{2\pi} \int_{-\pi}^{\pi} [\mathbf{g}(\omega; \theta) \mathbf{g}(\omega; \theta)']_{\theta=\theta_0} d\omega, \tag{3.4}$$

$$\begin{aligned} \mathbf{V}(\theta_0) = & \frac{1}{\pi} \int_{-\pi}^{\pi} [\mathbf{g}(\omega; \theta) \mathbf{g}(\omega; \theta)']_{\theta=\theta_0} d\omega \\ & + \frac{1}{2\pi} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} \left[\frac{\mathbf{g}(\omega_1; \theta) \mathbf{g}(\omega_2; \theta)'}{f(\omega_1; \theta) f(\omega_2; \theta)} S(-\omega_1, \omega_2, -\omega_2; \theta) \right]_{\theta=\theta_0} d\omega_1 d\omega_2, \end{aligned} \tag{3.5}$$

$\mathbf{g}(\omega; \theta) = \frac{\partial \log f(\omega; \theta)}{\partial \theta}$, and $S(\omega_1, \omega_2, \omega_3; \theta)$ denotes the model trispectrum.

The trispectrum entering the limiting variance through (3.5) is defined as the Fourier transform of the fourth-order cumulants of x_i (see, e.g., Mendel, 1991, for details). It is very difficult to obtain the trispectrum in closed form, except for some special cases. The most simple case arises when both m_i and ϵ_i are log-normally distributed, since then m_i and e_i , and thus x_i , are Gaussian implying that the fourth-order cumulants of x_i are identically zero and $S(\omega_1, \omega_2, \omega_3; \theta) \equiv 0$. The limiting variance of the Whittle estimator in (3.3) then reduces to $4\pi \mathbf{M}(\theta_0)^{-1}$.

Relaxing the Gaussianity of e_i while maintaining Gaussianity of m_i leads to a limiting variance matrix that is no longer robust to fourth-order cumulants, but is still available in closed form. Due to the independence of the multipliers and e_i , the cumulants and hence the trispectrum are additive, and since e_i is *iid* with finite fourth moment, the fourth-order cumulants satisfy $\text{cum}(e_i, e_{i+h_1}, e_{i+h_2}, e_{i+h_3}) = E(e_1^4)$ if $h_1 = h_2 = h_3 = 0$, and equal zero otherwise. Thus, $S(\omega_1, \omega_2, \omega_3; \theta) = E(e_1^4)/(2\pi)^3$ (e.g., Mendel, 1991). From this point of view, the log-normal specification of the MSM multipliers appears to be particularly attractive in practice, as the limiting variance of the Whittle estimator takes a manageable form and can be easily estimated by the plug-in estimators provided below.

Before we turn to the estimation of the asymptotic variance, we remark that the requirement in Proposition 2 that $4 + r$ moment of m_i exist for some $r > 0$, rather than for $r = 0$, is dictated precisely by the fact that we are unable to derive the trispectrum in closed form and verify directly that it is well-defined for a general MSMD process. Instead, we have to rely on a mixing inequality to establish that the fourth-order cumulants are absolutely summable, and this requires $r > 0$. Given strict stationarity and exponential strong mixing of x_i we nonetheless conjecture that Proposition 2 holds with $r = 0$ as well.

To estimate the asymptotic variance, we can use the plug-in estimators $\mathbf{M}(\hat{\theta})$ and $\mathbf{V}(\hat{\theta})$ for $\mathbf{M}(\theta_0)$ and $\mathbf{V}(\theta_0)$, respectively, provided in the Supplemental Appendix A3. Consistency of $\mathbf{M}(\hat{\theta})$ follows from the consistency of $\hat{\theta}_n$ and stochastic equicontinuity of $\mathbf{M}(\hat{\theta})$ where the latter is implied by the smoothness of the third-derivatives of the model spectral density on Θ , see Zaffaroni (2009) for details. For $\mathbf{V}(\hat{\theta})$ the consistency cannot be in general established unless one knows the trispectrum in closed form. When this is not the case, we propose a Newey–West estimator $\widehat{\mathbf{V}}(\hat{\theta}_0)$, see the Supplemental Appendix A3 for details. Alternatively, one can use a similar estimator proposed by Taniguchi (1982). A rigorous proof of consistency of $\widehat{\mathbf{V}}(\hat{\theta}_0)$ is beyond the scope of this article and is left for future work.

3.3. Linear forecasting

When optimal forecasting discussed in the previous subsection is not feasible due to the dimensionality of the state space, Lux (2008) suggests using best linear forecasts (e.g., Brockwell and Davis, 1991). This forecasting rule only requires the knowledge of the autocovariance function of the model and thus works

as long as one has a set of consistent parameter estimates at hand, regardless of the estimation method used to obtain them. Formally, an h -step ahead forecast based on the most recent n observations, denoted by $\hat{X}_{n+h|n}$, is obtained from $\hat{X}_{n+h|n} = \sum_{j=1}^n \phi_{nj}^{(h)} X_{n+1-j} = \phi_n^{(h)} X_n$, where the vector of weights $\phi_n^{(h)}$ is a solution to $\Gamma_n \phi_n^{(h)} = c_n^{(h)}$, in which $c_n^{(h)} = (c(h), c(h+1), \dots, c(n+h-1))'$ denotes the vector of autocovariances of the true process from lag h to lag $n+h-1$, and $\Gamma_n = \{c(i-j)\}_{i,j=1, \dots, n}$ is the variance-covariance matrix of $X_n = (X_1, X_2, \dots, X_n)'$. The autocovariance function of the MSMD process is provided in (2.7) and the weights $\phi_n^{(h)}$ can be efficiently calculated using the generalized Levinson-Durbin algorithm developed by Brockwell and Dahlhaus (2004).

3.4. Specification testing

To test the goodness of fit of the MSMD model, we employ the specification test of Chen and Deo (2004). The idea of the test is to compare the estimated model's spectral density with the smoothed periodogram of the data. Under the null hypothesis of correct model specification, the two should be close. The main advantage of this approach is that the test statistic does not require residuals, which makes it particularly suitable for specification testing of stochastic durations models.

The test statistic is given by

$$T_n = \left(\frac{2\pi}{n} \sum_{i=0}^{n-1} \tilde{f}(\omega_i) \right)^{-2} \left(\frac{2\pi}{n} \sum_{l=0}^{n-1} \tilde{f}^2(\omega_l) \right),$$

where

$$\tilde{f}(\omega) = \frac{2\pi}{n} \sum_{i=0}^{n-1} \frac{W(\omega - \omega_i) I_n(\omega_i)}{f(\omega_i; \hat{\theta})}, \quad W(\omega) = \frac{1}{2\pi} \sum_{|h| < n} k(h/p_n) e^{-ih\omega},$$

k is a symmetric kernel function with $k(0) = 1$, and p_n is a bandwidth parameter. Provided that (i) $\hat{\theta}$ is \sqrt{n} -consistent, (ii) the underlying process $\{x_i\}$ can be written as $x_i = \sum_{l=0}^{\infty} \psi_l \epsilon_{i-l}$, where ϵ_i is iid with zero mean, constant variance, and finite eighth moment, and $\sum_{l=0}^{\infty} |\psi_l| l^{1/2} < \infty$, (iii) the model spectral density is bounded away from zero on $[-\pi, \pi]$, (iv) the bandwidth satisfies $\log^6 n/p_n \rightarrow 0$ and $p_n^{3/2}/n \rightarrow 0$, and (v) the kernel satisfies certain regularity conditions, Chen and Deo (2004) show that $n(T_n - C_n(k))/D_n(k) \xrightarrow{d} N(0, 1)$, where the centering and scaling terms, $C_n(k)$ and $D_n(k)$, are provided in the Supplemental Appendix A4.

The assumptions underlying this result are clearly not satisfied for the logarithmic MSMD as the process is not linear and cannot be written in the form required by (ii) above. The process nonetheless possess a vector MA(∞) representation (A.1) with geometrically declining coefficients and martingale-difference innovations, which leads us to conjecture that the asymptotic normality of the test statistics still holds, though it remains unclear whether the limiting variance involves fourth-order cumulants. To shed some light on this issue, we examine the distribution of the test statistic T_n for a variety of MSMD specifications by Monte Carlo simulation, leaving the development of a rigorous limit theory for future work. The results are reported at the end of the next section.

3.5. Simulations

Before taking the model to the data, it is worthwhile exploring the finite-sample properties of the maximum likelihood and Whittle estimators. To do that, we run a simple Monte Carlo experiment for the binomial and log-normal MSMD models with $k = 8$ multipliers and either exponential or Weibull

innovations.² Following Lux (2008), we set the parameters of the MSM process as $b = 2$, $\gamma_k = 0.5$, $m_0 = 1.4$ (binomial), and $\lambda = 0.15$ (log-normal), and the parameter in the Weibull distribution of innovations as $\kappa = 1.45$.

Due to the computational burden associated with the exact maximum likelihood estimator, the number of Monte Carlo replications for MLE is limited to 500, 250, and 100 replications for $n = 1,000$, 2,500, and 5,000, respectively. All simulation results for the Whittle estimator are based on 1,000 replications, and we also consider very large samples of 10,000 observations, as the application of the Whittle estimator to the MSMD model is new and the large-sample properties have not been investigated by simulation before.

Table 1 summarizes the simulation results. Starting with the maximum likelihood estimator in the binomial MSMD model, we find that MLE delivers accurate and almost unbiased estimates for both exponential and Weibull specifications; the simulated standard errors scale with \sqrt{n} as dictated by asymptotic theory. As expected, the Whittle estimator is less precise than MLE, and it also entails a significant bias in samples smaller than 5,000 observations, particularly for the parameter b . The bias, however, disappears in large samples, and the standard errors also scale with \sqrt{n} as claimed in Proposition 2.

Table 2 reports the simulated size of the goodness-of-fit test discussed in the previous section. The test is implemented using the Bartlett kernel and setting the bandwidth according to $p_n = 3n^{0.4}$ as in

Table 1. Monte carlo simulation of the maximum likelihood (MLE) and whittle estimators of the parameters of the MSMD model with $k = 8$ binomial or log-normal multipliers and exponential or weibull innovations. We report average parameter estimates obtained in the simulation together with standard errors in parentheses. The true parameters used in the simulations are $b = 2$, $\gamma_k = 0.5$, $m_0 = 1.4$, $\lambda = 0.15$, and $\kappa = 1.45$. The results for mle are based on 500, 200, and 100 replications for the samples of $n = 1,000, 2,000$, and 5,000 observations, respectively. All simulations of the whittle estimator are based on 1,000 replications.

	<i>Exp</i>				$\mathcal{W}(\kappa)$			
	1,000	2,500	5,000	10,000	1,000	2,500	5,000	10,000
A. Binomial multipliers: MLE								
m_0	1.382 (0.035)	1.393 (0.021)	1.395 (0.016)	— (—)	1.393 (0.036)	1.3993 (0.022)	1.400 (0.015)	— (—)
b	1.832 (0.346)	1.936 (0.250)	1.949 (0.162)	— (—)	1.982 (0.438)	1.998 (0.258)	2.022 (0.180)	— (—)
γ_k	0.501 (0.164)	0.499 (0.107)	0.494 (0.069)	— (—)	0.506 (0.160)	0.502 (0.101)	0.509 (0.064)	— (—)
κ	— (—)	— (—)	— (—)	— (—)	1.465 (0.098)	1.458 (0.064)	1.453 (0.037)	— (—)
B. Binomial multipliers: Whittle								
m_0	1.425 (0.080)	1.409 (0.047)	1.400 (0.018)	1.400 (0.007)	1.437 (0.088)	1.411 (0.045)	1.403 (0.019)	1.401 (0.013)
b	2.608 (1.871)	2.231 (1.190)	1.996 (0.214)	1.999 (0.131)	2.883 (2.263)	2.222 (1.145)	2.022 (0.247)	2.012 (0.152)
γ_k	0.467 (0.221)	0.497 (0.146)	0.499 (0.101)	0.502 (0.075)	0.514 (0.273)	0.524 (0.208)	0.523 (0.158)	0.514 (0.104)
κ	— (—)	— (—)	— (—)	— (—)	1.617 (0.415)	1.523 (0.244)	1.491 (0.161)	1.466 (0.098)
C. Log-normal multipliers: Whittle								
λ	0.206 (0.143)	0.167 (0.076)	0.153 (0.031)	0.150 (0.015)	0.214 (0.137)	0.168 (0.068)	0.155 (0.030)	0.151 (0.015)
b	2.803 (2.089)	2.279 (1.282)	2.034 (0.483)	1.994 (0.182)	2.924 (2.267)	2.257 (1.216)	2.046 (0.455)	2.008 (0.195)
γ_k	0.474 (0.247)	0.498 (0.173)	0.501 (0.124)	0.499 (0.086)	0.525 (0.296)	0.531 (0.225)	0.528 (0.176)	0.516 (0.128)
κ	— (—)	— (—)	— (—)	— (—)	1.583 (0.315)	1.515 (0.192)	1.487 (0.137)	1.465 (0.075)

²In an earlier version of the article we also reported MLE simulation results for the MSMD model with Burr and generalized gamma distributions of the innovations. The results are qualitatively similar to the exponential and Weibull cases and show that the ML estimator works well even when the innovations are drawn from multiparameter distributions.

Table 2. Monte carlo simulation of the goodness-of-fit test for the MSMD model with $k = 8$ binomial or log-normal multipliers and exponential, weibull, or log-normal innovations. We report the simulated size for 5 and 10% nominal level. The true parameters used in the simulations are $b = 2$, $\gamma_k = 0.5$, $m_0 = 1.4$, $\lambda = 0.15$, and $\kappa = 1.45$. The sample size used in the simulations is $n = 10,000$ observations and the simulations are based on 1,000 replications.

	Distribution of ϵ		
	<i>exp</i>	$\mathcal{W}(\kappa)$	LN(λ)
A. 10% nominal			
Binomial multipliers	11.3	9.2	9.9
Lognormal multipliers	11.0	11.7	13.6
B. 5% nominal			
Binomial multipliers	6.2	6.0	5.6
Lognormal multipliers	6.2	6.4	7.0

Chen and Deo (2004). We use the same MSMD specifications as in the previous simulations and report results for samples of size $n = 10,000$. We find that the simulated size is relatively close to the nominal levels across the different MSMD specifications.

4. Competing duration models

Models in the duration literature mimic those in the stochastic volatility literature, and might be similarly divided into observable or GARCH-type models and latent factor or Stochastic Volatility (SV)-type models. The ACD model of Engle and Russell (1998) is a member of the former class, and was extended by Jasiak (1998) to the Fractionally Integrated ACD (FIACD) model to incorporate long memory. Bauwens and Veredas' 2004 Stochastic Conditional Duration (SCD) model is a latent factor model, and was modified by Deo et al. (2006) to create the Long-memory Stochastic Duration model (LMSD) by letting the latent factor follow a long-memory process.

It is beyond the scope of this article to review and compare all existing durations models; we refer the reader to a survey by Pacurar (2008). Since we are interested in modeling and forecasting persistent durations, we focus here on those models that can capture slowly decaying autocorrelations. As noted by Deo et al. (2010), the FIACD model is not a long-memory model in the usual sense, as it has infinite mean and hence the autocorrelation function does not exist. We are therefore left with the LMSD model as the only genuine long-memory duration model with well-behaved moments. To assess the benefits of using the relatively more complicated MSMD and LMSD models in practice, we also compare their performance with the short-memory ACD model of Engle and Russell (1998). The Supplemental Appendix A5 provides a brief description of these models.

5. Relation to counts and realized volatility

Deo et al. (2009) and Deo et al. (2010) recently investigate the propagation of memory of durations to counts and thereby realized volatility.³ They show that if durations have long (short) memory, then under certain conditions the counts have long (short) memory as well. They also note that, alternatively, long memory in realized volatility can be generated by *iid* infinite-variance durations, as originally modeled by Liu (2000), where the memory parameter associated to realized volatility is inversely proportional to the tail index of the distribution of durations. These are, of course, two fundamentally different approaches to generating long memory in volatility, and we naturally focus on the former here, not only because the MSMD durations are not *iid*, but also because we find no empirical evidence supporting the infinite variance assumption required by Liu (2000).

To fix notation, recall that t_i denotes the time of the i th event, $X_i = t_i - t_{i-1}$ is the duration between two consecutive events, and let $N(t)$ denote the counting process that counts the number of events that

³See McAleer and Medeiros (2008) for a review of the literature on realized volatility.

have occurred up to time t . In more detail, counts and durations are stationary under different measures, since they define the irregularly-spaced event process (a point process) in terms of different sets of events. We refer to these measures as \mathbb{P}^N and \mathbb{P} , respectively. As illuminated by Deo et al. (2009), the relevant measure depends on how $N(t)$ is calculated: if it is calculated from the opening of the market on a given day, the relevant measure is \mathbb{P}^N , while if from the first event on that day, the relevant measure is \mathbb{P} . Since most assets tend to be heavily traded after market opening, the difference may be empirically small.

By making use of equivalence theorems (e.g., Nieuwenhuis, 1989), Deo et al. (2009) establish conditions under which memory propagates from durations to counts, then to squared returns and realized volatility. In particular, they show that under certain conditions, the short memory of durations generated by a stationary ACD model implies short memory in the associated counts and realized volatility, while the long memory of durations in the LMSD model implies long memory in counts and realized volatility. With respect to the MSMD process now, the following proposition establishes the conditions under which the short-memory feature of the MSMD (for finite k) translates into short memory in the induced counts.

Proposition 3. *Let $\{X_i\}$ be an MSMD process with finite k , $E(M^{3+r}) < \infty$ and $\mathbb{E}(\epsilon_1^{3+r}) < \infty$ for some $r > 0$. Then the induced counting process $N(t)$ satisfies $\text{Var}_N(N(t)) \sim ct$ for some $c < \infty$, where Var_N denotes the variance under \mathbb{P}^N .*

The proof is provided in Section A.3 of the Appendix. To link the counts and realized volatility, we follow Deo et al. (2009) and employ the simple continuous-time pure-jump model of Oomen (2006). The logarithmic price process, $p(t)$, is assumed to have the following dynamics:

$$p(t) = p(0) + \sum_{j=1}^{N(t)} \xi_j, \quad \xi_j \stackrel{iid}{\sim} N(0, \sigma_\xi^2), \quad (5.1)$$

where $N(t)$ is the counting process defined above and ξ_j is the size of the j th jump, which is assumed to be independent from the counting process. This assumption significantly simplifies the analysis, but it may not be appropriate for all asset classes: a recent study by Renault and Werker (2011) shows that it is indeed violated in the case of selected individual stocks traded on the New York Stock Exchange.

A natural measure of variation in the model in Eq. (5.1) is the quadratic variation given by

$$\langle p \rangle_t = \sum_{j=1}^{N(t)} \xi_j^2.$$

The quadratic variation can be estimated consistently by realized variance. Dividing the time interval $[0, t]$ into n non-overlapping intervals of length $\Delta t = t/n$, the realized variance is defined as

$$RV_{t,n} = \sum_{i=1}^n (p(i\Delta t) - p((i-1)\Delta t))^2. \quad (5.2)$$

It follows from Deo et al. (2009) that for the MSMD process satisfying the assumptions of Proposition 1, the realized volatility is a short-memory process.

It is difficult to derive analytically the autocorrelation functions of counts and realized volatility induced by the MSMD process and its competitors. We therefore proceed by simulation. For each duration model, we simulate a trajectory of the induced counting process $N(t)$ and via (5.1) a trajectory of the associated logarithmic price process $p(t)$. From the simulated price process we then calculate a time series of *daily* realized variance according to (5.2), where we define one day to have 6.5 hours, or 23,400 seconds. For all duration models, we set the unconditional mean of durations equal to 2 minutes, so that there are around 195 price changes on a typical day in the simulation. The price innovations, ξ_j , are drawn randomly from the normal distribution with zero mean and variance $\sigma_\xi^2 = 1/195$, implying that the




Figure 3. (a) Theoretical autocorrelation functions of durations from i) the ACD model with parameters $\alpha = 0.24$, $\beta = 0.69$, ii) the binomial MSMD(4) model with parameters $m_0 = 1.84$, $b = 3.30$, $\gamma_k = 0.047$, iii) the binomial MSMD(8) model with parameters $m_0 = 1.55$, $b = 3.00$, $\gamma_k = 0.076$, and iv) the LMSD model with $\omega = 1.028$, $\beta = 0.73$, $d = 0.47$, $\sigma_u^2 = 0.029$. (b) Simulated autocorrelation functions of daily realized volatility generated by the corresponding duration models i)–iv).

average daily realized variance is around 1%. Finally, to facilitate comparison, we calibrate the parameters of the duration models so that they share the same first-order autocorrelation coefficient, which we set equal to 0.45; see the caption of Fig. 3 for the exact parameters values used in the simulations.

Figure 3 plots the theoretical autocorrelation functions of the ACD, MSMD and LMSD durations and their corresponding simulated autocorrelation functions for realized variance as implied by model (5.1). The figure clearly illustrates how memory propagates from durations to realized volatility. The short-memory ACD model generates realized variance with little persistence, while the long-memory LMSD generates a highly persistent realized variance. The MSMD model is capable of generating both: when the number of multipliers is small ($k = 4$), the autocorrelation function of realized variance decays very quickly despite the ACF of durations being quite persistent. Increasing the number of multipliers to 8, the persistence of realized volatility increases dramatically and its ACF now clearly exhibits long-memory features. Thus, despite being short-memory, the MSMD model is capable of generating both highly persistent durations as well as highly persistent realized volatility in the pure jump model (5.1).

6. Data description

We now apply the MSMD model and its competitors to price durations of three major foreign exchange (FX) futures contracts traded on the Chicago Mercantile Exchange (CME). Our dataset includes all transactions for the Swiss Franc (CHF), Euro (EUR) and Japanese Yen (JPY) futures contracts between November 9, 2009 and January 29, 2010. The data is supplied by TickData, Inc. We focus on the most liquid (front) contracts and restrict attention to the main CME trading hours of 7:20 - 14:00 Chicago time. U.S. and U.K. Bank holidays are discarded.

Price durations are defined as the minimum time it takes for the price to move by a certain amount. We construct them from the transactions durations, which are simply the durations between successive trades, by a process called thinning. Due to microstructure frictions, such as bid-ask bounce, the price durations may be more informative about the underlying prices process and its volatility as thinning reduces the distortions due to microstructure noise and eliminates duplicate prices, that is transactions with zero price changes. Also, Engle and Russell (1998) show that price durations are closely related to the instantaneous volatility: low price durations imply high instantaneous volatility of the underlying price process, and vice versa.

Correspondingly, we construct the price durations by successively measuring the minimum time required for the futures price to move by at least c , starting from the first transaction on each day and discarding overnight durations. The FX futures contracts are highly liquid and usually trade with a tight bid-ask spread of 1-2 ticks, where the tick size equals 0.0001 for CHF and EUR, and 0.01 for JPY.

To eliminate spurious price changes due to the bid-ask bounce, we set $c = 0.0003$ for CHF and EUR, and $c = 0.03$ for JPY. To facilitate comparison across the different currencies, we work with the first 12,000 prices durations for each FX futures contract available in our sample period. The sample size is, therefore, kept fixed at 12,000 but the sample period varies across the three data sets, though they all start on the November 9, 2009.

Table 3 reports the descriptive statistics for the FX futures price durations data. The mean of the price durations is 118s, 106s and 90s for CHF, EUR, and JPY, respectively, while the median is around half the mean at 66s, 59s, and 43s, respectively, indicating that the distributions of the price durations are heavily positively skewed. The minimum price duration equals 1s for all currencies, while the maximum price duration reaches 42 minutes, 1 hour, and 53 minutes, respectively. Consistent with previous empirical evidence, we find that the distribution of price durations exhibits overdispersion, i.e., the standard deviation of the price durations significantly exceeds the mean by a factor of 1.351, 1.337, and 1.512 for CHF, EUR, and JPY, respectively.

It is well-known that the trading activity in most financial markets varies considerably over the course of the day, see, e.g., Engle and Russell (1998) who note a hump-shaped pattern for transaction and price durations of individual stocks traded on the New York Stock Exchange (NYSE), with relatively shorter durations at the start and end of the trading day, and longer durations during lunchtime. Consequently, the duration process contains a significant seasonal component that has to be accounted for when estimating a duration model.

There are in principle two ways to do that. First, by incorporating seasonality into the duration models directly and estimating the seasonal parameters jointly with the dynamic parameters of the duration process (Rodríguez-Poo et al., 2007). Alternatively, one can first estimate the seasonal component semi- or nonparametrically and fit the duration model to the seasonally-adjusted durations (e.g., Engle and Russell, 1998, and Fernandes and Grammig, 2006, among many others). Engle (2000) notes that the large sample sizes typically available in empirical work make the loss of efficiency of the two-step procedure relatively small. Given the complexity of the duration models we are considering in this article, we opt for the two-stage approach and employ nonparametric regression (the Nadaraya–Watson estimator) to estimate the seasonal component of the price durations, separately for each day of the week as in Bauwens and Veredas (2004).

The estimated intraday seasonal patterns are reported in the top panel of Fig. 4. The diurnal pattern is relatively stable across the days of the week and currencies up to around 11:00 Chicago time. During

Table 3. Descriptive statistics for Swiss franc (CHF), Euro (EUR), and Japanese Yen (JPY) futures price durations. The sample period runs between November 9, 2009 and January 29, 2010, and the sample size is $n = 12,000$ for all datasets. The columns labeled “raw” report descriptive statistics for the raw price durations, the columns “adj” give descriptive statistics for the seasonally adjusted durations and the columns “log-adj” report descriptive statistics for the logarithmic seasonally adjusted durations. The rows labeled “left tail” and “right tail” report the estimated tail indexes (with standard error in parentheses) for the left and right tail, respectively.

	CHF			EUR			JPY		
	raw	adj	log-adj	raw	adj	log-adj	raw	adj	log-adj
Mean	117.8	1.000	−0.602	105.5	1.002	−0.580	89.50	1.002	−0.720
Median	66.00	0.618	−0.481	59.00	0.619	−0.479	43.00	0.536	−0.623
Minimum	1.000	0.004	−5.613	1.000	0.004	−5.482	1.000	0.007	−5.028
Maximum	2498	15.78	2.759	3574	20.55	3.023	3160	20.31	3.011
Std.dev.	159.1	1.194	1.204	141.1	1.205	1.167	135.3	1.386	1.296
Dispersion	1.351	1.194	—	1.337	1.204	—	1.512	1.383	—
Skewness	3.775	3.307	−0.497	4.410	3.731	−0.439	4.586	4.053	−0.322
Kurtosis	26.43	22.10	3.187	49.99	29.71	3.140	46.06	30.60	2.817
Left tail	—	—	10.16	—	—	9.409	—	—	12.88
	—	—	(0.002)	—	—	(0.002)	—	—	(0.002)
Right tail	—	3.976	9.703	—	3.501	7.949	—	3.206	7.955
	—	(0.003)	(0.002)	—	(0.003)	(0.003)	—	(0.003)	(0.003)

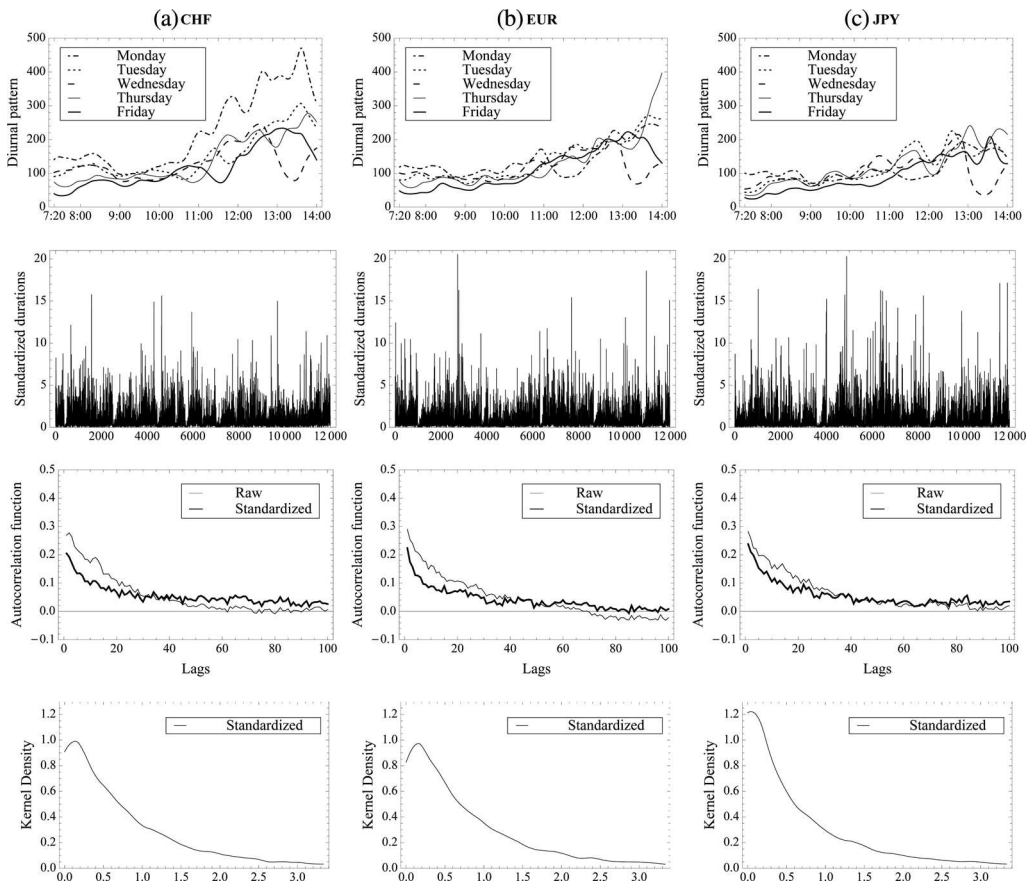


Figure 4. Foreign exchange price durations data. The top row shows the diurnal pattern estimated by kernel regression separately for each day of the week. The second row shows the time series of standardized durations and the third row reports the autocorrelation functions of raw and standardized durations. The bottom row plots the empirical density of standardized durations obtained by a boundary-corrected kernel estimator.

this period the U.S. and European trading hours overlap and trading activity in the market is at its peak. After 11:00, trading in London, where a large proportion of global FX trading takes place (King et al., 2011), gradually ceases and the average price durations become progressively longer. The exception is Wednesdays, for which we observe a significant dip in the average price durations around 13:30, most likely due to elevated volatility surrounding macroeconomic announcements.

Figure 4 plots the autocorrelation function of the adjusted durations obtained by dividing the raw durations by the estimated intraday component. Clearly, the persistence in the price durations is not induced by the seasonal component. The descriptive statistics for the adjusted durations are reported in Table 3. The mean is, by construction, close to one, the median remains significantly lower than the mean, and over-dispersion is slightly attenuated by the adjustment. The empirical densities of the standardized durations, estimated in Fig. 4 by a boundary-corrected kernel estimator, are non-monotonic and heavily positively skewed. Finally, we examine the descriptive statistics for the logarithmic standardized durations. We find that the logarithmic durations exhibit negative skewness but almost no excess kurtosis. The tail index estimates obtained by the method of Huisman et al. (2001) indicate that the first 8-9 moments exist, which is in stark contrast to the non-logarithmic durations that only seems to possess the first three moments. The asymptotics for the Whittle estimator discussed in Section 3.2 therefore applies to our logarithmic durations data.

Table 4. MSMD parameter estimates for de-seasonalised CHF price durations. (A) maximum likelihood estimates (MLE) of the binomial MSMD models with exponential and weibull innovations. (B) Whittle estimates of the binomial MSMD models with exponential and weibull innovations. (C) Whittle estimates of the log-normal MSMD models with exponential and weibull innovations. Standard errors are reported in parentheses. The specification test T_n is reported with p -values in parentheses.

	<i>Exp</i>			$\mathcal{W}(\kappa)$		
	$k = 4$	$k = 6$	$k = 8$	$k = 4$	$k = 6$	$k = 8$
A. Binomial multipliers: MLE						
m_0	1.356 (0.009)	1.328 (0.008)	1.316 (0.010)	1.439 (0.008)	1.394 (0.009)	1.355 (0.007)
b	2.339 (0.551)	1.962 (0.311)	2.045 (0.308)	7.059 (0.779)	6.075 (0.621)	4.188 (0.336)
γ_k	0.067 (0.014)	0.059 (0.013)	0.064 (0.014)	0.533 (0.062)	0.973 (0.031)	0.999 (0.002)
κ	— (—)	— (—)	— (—)	1.346 (0.024)	1.485 (0.041)	1.566 (0.045)
T_n	20.078 (0.000)	22.644 (0.000)	24.322 (0.000)	−1.150 (0.875)	−1.252 (0.895)	−1.295 (0.902)
log L	−8881.967	−8859.848	−8863.670	−8661.562	−8614.860	−8594.576
B. Binomial multipliers: Whittle						
m_0	1.293 (0.014)	1.243 (0.012)	1.212 (0.010)	1.369 (0.017)	1.308 (0.017)	1.269 (0.015)
b	1.876 (0.222)	1.477 (0.121)	1.326 (0.083)	4.150 (0.836)	2.600 (0.449)	2.045 (0.250)
γ_k	0.058 (0.015)	0.058 (0.016)	0.058 (0.016)	0.476 (0.213)	0.568 (0.299)	0.608 (0.287)
κ	— (—)	— (—)	— (—)	1.376 (0.047)	1.380 (0.059)	1.379 (0.054)
T_n	35.836 (0.000)	36.270 (0.000)	36.480 (0.000)	−1.084 (0.861)	−1.079 (0.860)	−1.077 (0.859)
q-log L	−2923.704	−2924.308	−2924.557	−3381.884	−3381.712	−3381.651
C. Log-normal multipliers: Whittle						
λ	0.096 (0.010)	0.064 (0.007)	0.048 (0.005)	0.162 (0.017)	0.107 (0.014)	0.079 (0.009)
b	1.876 (0.222)	1.477 (0.121)	1.326 (0.083)	4.150 (0.836)	2.600 (0.449)	2.045 (0.250)
γ_k	0.058 (0.015)	0.058 (0.016)	0.058 (0.016)	0.476 (0.213)	0.568 (0.299)	0.608 (0.287)
κ	— (—)	— (—)	— (—)	1.376 (0.047)	1.380 (0.059)	1.379 (0.054)
T_n	35.835 (0.000)	36.270 (0.000)	36.479 (0.000)	−1.084 (0.861)	−1.078 (0.860)	−1.077 (0.859)
q-log L	−2923.704	−2924.308	−2924.557	−3381.884	−3381.712	−3381.651

7. Empirical results

The following section compares the estimation and forecasting performance of our MSMD model to the competing ACD and LMSD models. We use the first 10,000 observations for estimation and in-sample specification tests and reserve the remaining 2,000 observations for evaluating out-of-sample forecasting performance. As is common in the durations literature, in the rest of the article we work exclusively with the seasonally-adjusted durations. Since the mean of the standardized durations is, by construction, close to one, we impose this restriction in all models and do not report the (restricted) estimates of the various constant terms (ψ in the MSMD model and ω in the ACD and LMSD models).

7.1. Estimation results

We start by describing the in-sample estimates of MSMD for the three currencies in our sample. We estimate the MSMD model with $k = 4, 6,$ and 8 multipliers; increasing the number of multipliers beyond 8 does not improve the in-sample and out-of-sample performance of the model. We use exact maximum likelihood to estimate the MSMD model with binomial multipliers and the Whittle estimator for both

the binomial and log-normal multipliers. All models are estimated with either the exponential or the Weibull distribution of innovations. Since the MSMD parameter space is not compact ($0 < m_0 < 1$, $\lambda > 0$, $b > 1$, and $0 < \gamma_k < 1$), some constraints are generally required to achieve numerical stability of the optimization routines. For both MLE and Whittle estimation, we use the `MaxSQP` function in the Ox language of Doornik (2006) to maximize the respective objective functions and search over the following parameter space: $m_0 \in [1.001, 1.999]$, $\lambda \in [0.001, 10]$, $b \in [1.001, 10]$, and $\gamma_k \in [0.001, 0.999]$.

Tables 4, 5, and 6 show the estimation results for the CHF, EUR, and JPY, respectively. All estimated parameters have reasonable standard errors. The goodness-of-fit test of Chen and Deo (2004) strongly rejects the null hypothesis of correct model specification for all MSMD models with exponentially distributed innovations. This is generally true for all k 's, and across currencies. On the contrary, both the binomial and the log-normal MSMD models with Weibull innovations seem to be correctly specified as we can not reject the null hypothesis at the 5% level for any of the estimated models. In addition, the log-likelihood is uniformly higher for the binomial MSMD models with Weibull innovations. Thus, the additional flexibility of the Weibull distribution seems to improve the in-sample fit of the MSMD models significantly.

Table 5. MSMD parameter estimates for de-seasonalised Euro futures price durations. (A) maximum likelihood estimates (MLE) of the binomial MSMD models with exponential and weibull innovations. (B) whittle estimates of the binomial MSMD models with exponential and weibull innovations. (C) whittle estimates of the log-normal MSMD models with exponential and weibull innovations. Standard errors are reported in parentheses. The specification test T_n is reported with p -values in parentheses.

	<i>Exp</i>			$\mathcal{W}(\kappa)$		
	<i>k</i> = 4	<i>k</i> = 6	<i>k</i> = 8	<i>k</i> = 4	<i>k</i> = 6	<i>k</i> = 8
A. Binomial multipliers - MLE						
m_0	1.349 (0.011)	1.281 (0.011)	1.331 (0.020)	1.452 (0.007)	1.397 (0.008)	1.342 (0.008)
b	3.001 (0.636)	1.932 (0.299)	3.000 (0.525)	7.241 (0.737)	5.491 (1.032)	3.289 (0.563)
γ_k	0.086 (0.015)	0.088 (0.022)	0.089 (0.017)	0.881 (0.032)	0.999 (0.004)	0.999 (0.004)
κ	— (—)	— (—)	— (—)	1.554 (0.029)	1.675 (0.068)	1.670 (0.067)
T_n	33.655 (0.000)	36.987 (0.000)	36.923 (0.000)	−0.935 (0.825)	−0.975 (0.835)	−1.041 (0.851)
log L	−9007.010	−8998.909	−9001.426	−8706.568	−8652.807	−8649.950
B. Binomial multipliers - Whittle						
m_0	1.272 (0.013)	1.224 (0.011)	1.195 (0.010)	1.402 (0.062)	1.354 (0.023)	1.308 (0.017)
b	1.643 (0.343)	1.353 (0.197)	1.242 (0.139)	5.060 (2.837)	3.391 (0.585)	2.459 (0.260)
γ_k	0.058 (0.017)	0.058 (0.019)	0.058 (0.021)	0.901 (0.438)	0.999 (0.008)	0.999 (0.005)
κ	— (—)	— (—)	— (—)	1.588 (0.302)	1.719 (0.173)	1.701 (0.136)
T_n	53.333 (0.000)	53.720 (0.000)	53.887 (0.000)	−1.614 (0.947)	−1.619 (0.947)	−1.619 (0.947)
q-log L	−3071.748	−3071.930	−3072.001	−3648.922	−3648.827	−3648.782
C. Log-normal multipliers - Whittle						
λ	0.081 (0.008)	0.053 (0.005)	0.040 (0.004)	0.199 (0.076)	0.147 (0.023)	0.106 (0.014)
b	1.643 (0.343)	1.352 (0.197)	1.242 (0.139)	5.060 (2.835)	3.391 (0.585)	2.459 (0.263)
γ_k	0.058 (0.017)	0.058 (0.019)	0.058 (0.021)	0.901 (0.438)	0.999 (0.008)	0.999 (0.005)
κ	— (—)	— (—)	— (—)	1.588 (0.301)	1.719 (0.173)	1.701 (0.139)
T_n	53.333 (0.000)	53.720 (0.000)	53.888 (0.000)	−1.614 (0.947)	−1.619 (0.947)	−1.619 (0.947)
q-log L	−3071.748	−3071.930	−3072.001	−3648.922	−3648.827	−3648.782

Table 6. MSMD parameter estimates for de-seasonalised Japanese Yen futures price durations. (A) maximum likelihood estimates (MLE) of the binomial MSMD models with exponential and weibull innovations. (B) whittle estimates of the binomial MSMD models with exponential and weibull innovations. (C) whittle estimates of the log-normal MSMD models with exponential and weibull innovations. standard errors are reported in parentheses. The specification test T_n is reported with p -values in parentheses.

	Exp			$\mathcal{W}(k)$		
	$k = 4$	$k = 6$	$k = 8$	$k = 4$	$k = 6$	$k = 8$
A. Binomial multipliers - MLE						
m_0	1.415 (0.010)	1.343 (0.014)	1.343 (0.013)	1.481 (0.007)	1.424 (0.007)	1.376 (0.008)
b	3.348 (0.552)	2.188 (0.285)	2.454 (0.263)	4.848 (0.366)	4.006 (0.340)	2.933 (0.253)
γ_k	0.152 (0.018)	0.164 (0.025)	0.180 (0.025)	0.826 (0.047)	0.985 (0.012)	0.999 (0.002)
κ	— (—)	— (—)	— (—)	1.481 (0.035)	1.590 (0.040)	1.672 (0.063)
T_n	26.401 (0.000)	28.896 (0.000)	30.695 (0.000)	1.006 (0.157)	1.508 (0.066)	1.459 (0.072)
$\log L$	-8255.173	-8252.226	-8256.733	-8053.231	-8008.807	-8000.278
B. Binomial multipliers - Whittle						
m_0	1.342 (0.013)	1.283 (0.011)	1.247 (0.010)	1.434 (0.012)	1.367 (0.015)	1.323 (0.015)
b	1.991 (0.315)	1.560 (0.157)	1.391 (0.103)	3.954 (0.422)	2.574 (0.285)	2.046 (0.185)
γ_k	0.084 (0.014)	0.091 (0.016)	0.095 (0.018)	0.569 (0.079)	0.723 (0.170)	0.784 (0.183)
κ	— (—)	— (—)	— (—)	1.396 (0.031)	1.422 (0.063)	1.430 (0.073)
T_n	49.062 (0.000)	49.129 (0.000)	49.138 (0.000)	0.757 (0.225)	0.748 (0.227)	0.745 (0.228)
$q\text{-log } L$	-2501.523	-2501.468	-2501.450	-2883.415	-2883.466	-2883.507
C. Log-normal multipliers - Whittle						
λ	0.136 (0.012)	0.088 (0.008)	0.065 (0.007)	0.241 (0.017)	0.160 (0.016)	0.119 (0.012)
b	1.991 (0.315)	1.560 (0.157)	1.391 (0.125)	3.954 (0.422)	2.574 (0.284)	2.046 (0.185)
γ_k	0.084 (0.014)	0.091 (0.016)	0.095 (0.011)	0.569 (0.079)	0.723 (0.170)	0.784 (0.183)
κ	— (—)	— (—)	— (—)	1.396 (0.031)	1.422 (0.063)	1.430 (0.073)
T_n	49.061 (0.000)	49.128 (0.000)	49.514 (0.000)	0.757 (0.225)	0.748 (0.227)	0.745 (0.228)
$q\text{-log } L$	-2501.523	-2501.468	-2501.347	-2883.415	-2883.466	-2883.507

Turning to the number of multipliers, we find that the log-likelihood increases with increasing k in all MSMD models with Weibull innovations. In the case of exponential innovations, the models with six multipliers yield the highest log-likelihood. We have initially experimented with a wider range of values of k and found that going beyond 8 multipliers offers little improvement in terms of both in-sample as well as out-of-sample performance, while reducing k below 4 diminishes performance considerably. The results are available upon request.

Comparing the MLE and Whittle parameter estimates for the binomial MSMD specifications, we find that the latter are typically smaller than the former, but generally exhibit a similar pattern. Specifically, both the MLE and Whittle estimates of b tend to decrease with increasing k , while the estimates of γ_k tend to increase. Intuitively, holding all parameters fixed, increasing the number of multipliers increases the persistence of the MSMD process (see Fig. 2(d)), and hence to fit a given persistence in the data the parameters b and γ_k must fall and/or rise, respectively, to compensate (see (b)). Additionally, we observe that the estimates of m_0 fall with increasing k , in order to compensate for the increase in unconditional variance of the MSMD process associated with rising k (see Eq. (2.6)). A similar pattern is found for

the parameter λ in the specification with log-normal multipliers. Note that it is not surprising to find that the Whittle estimates of b , γ_k , and κ (when applicable) are the same across the binomial and log-normal specifications; the two model spectral densities only differ in the parametrization of $\text{Var}(\log M)$, see Eq. (3.2). This does not, however, imply that the linear forecasts of durations obtained from these models will be the same. The linear forecasts of durations depend on $E(M^2)$ and $\text{Var}(M)$ (see Eqs. (2.6) and (2.7)), and the fact that $\text{Var}(\log M)$ is the same across the binomial and log-normal specifications does not imply that $E(M^2)$ and $\text{Var}(M)$ are as well. This will generally be the case whenever the parameter estimates are obtained by implementing the Whittle estimator on nonlinearly transformed durations (logs in the present application), rather than the durations themselves.

Having estimated the MSMD model, we now turn to the competing duration models. Table 7 shows the results from estimating the exponential and Weibull ACD and LMSD models for the three FX futures price durations. All estimated parameters have reasonable standard errors. The ACD model with Weibull innovations achieves higher log-likelihood than the ACD model with exponentially distributed innovations, but none of these models generate higher log-likelihoods than the corresponding binomial MSMD models estimated by maximum likelihood. The ACD parameter estimates are qualitatively

Table 7. Maximum likelihood estimates of exponential and weibull ACD models and whittle estimates of exponential and weibull LMSD models for de-seasonalized (A) Swiss franc, (B) Euro, and (C) Japanese Yen futures price durations. Standard errors are reported in parentheses.

	ACD			LMSD	
	Exp	$\mathcal{W}(\kappa)$		Exp	$\mathcal{W}(\kappa)$
A. Swiss Franc					
α	0.152 (0.005)	0.151 (0.000)		— (—)	— (—)
β	0.830 (0.007)	0.858 (0.000)		0.879 (0.043)	−0.043 (0.099)
d	— (—)	— (—)		0.500 (0.074)	0.477 (0.033)
σ_u^2	— (—)	— (—)		0.005 (0.001)	0.382 (0.143)
κ	— (—)	0.996 (0.002)		— (—)	1.519 (0.125)
log L	−8953.300	−8930.917	q-log L	−2928.500	−3376.700
B. Euro					
α	0.157 (0.006)	0.155 (0.006)		— (—)	— (—)
β	0.811 (0.008)	0.812 (0.007)		0.850 (0.054)	−0.054 (0.046)
d	— (—)	— (—)		0.500 (0.073)	0.387 (0.037)
σ_u^2	— (—)	— (—)		0.006 (0.002)	0.640 (0.186)
κ	— (—)	1.034 (0.007)		— (—)	1.951 (0.369)
log L	−9068.800	−9058.900	q-log L	−3074.600	−3643.700
C. Japanese Yen					
α	0.188 (0.005)	0.193 (0.006)		— (—)	— (—)
β	0.780 (0.007)	0.775 (0.008)		0.848 (0.043)	−0.081 (0.025)
d	— (—)	— (—)		0.367 (0.076)	0.373 (0.034)
σ_u^2	— (—)	— (—)		0.020 (0.004)	1.022 (0.231)
κ	— (—)	0.949 (0.007)		— (—)	2.570 (1.067)
log L	−8500.100	−8473.900	q-log L	−2515.400	−2888.600

similar (relatively high β and small α) and imply very high persistence as $(\alpha + \beta)$ is close to one. High persistence is also implied by the LMSD parameter estimates, where the long memory parameter estimates (d) lie between 0.37 and 0.50. It is difficult to assess the relative in-sample fit of the Weibull and exponential LMSD specifications, since the Whittle quasi-likelihoods are not directly comparable.

7.2. Out-of-sample forecasting performance

Our main interest lies in relative forecasting performance rather than in the in-sample fit of the various duration models. As we experiment with alternative estimation methods (MLE vs. Whittle) and forecasting schemes (optimal vs. linear), we are really going to be comparing alternative forecasting methods rather than models (Giacomini and White, 2006). The goal is to shed light not only on the relative ability of the alternative models to capture persistence in the data, but also on the impact of parameter uncertainty and the choice of forecasting rule on relative predictive performance. Specifically, we compare the following methods: (a) optimal forecasts from binomial MSMD(6) or MSMD(8) models estimated by maximum likelihood; (b) linear forecasts from binomial and log-normal MSMD(6) or MSMD(8) models estimated by the Whittle estimator; (c) Kalman filter-based forecasts from the LMSD model estimated by the Whittle estimator; and (d) ARMA representation-based forecasts from the ACD model estimated by maximum likelihood. We also experiment with equally-weighted combinations of (a) and (c), and (b) and (c), as model averaging may help reduce model uncertainty.

We compute and evaluate one step ahead and cumulative 5, 10, and 20 step ahead forecasts of price durations. The cumulative h -step ahead forecast, which we denote by $x_{n,h}$, are obtained from the usual multistep ahead forecast by $x_{n,h} = \sum_{j=1}^h x_{n+j|n}$. Thus, $x_{n,h}$ forecasts the time it takes for h price changes to occur, as opposed to $x_{n+h|n}$, which forecasts the time elapsed between the $(h - 1)$ th and h th price changes. We focus on the cumulative forecasts as they are more interesting in applications, for example in predicting realized variance. We evaluate the accuracy of the forecasts using two common loss functions, the mean square error (MSE) and the mean absolute deviation (MAD), and assess the differences between models statistically by the Diebold and Mariano (1995) test for equality of forecast accuracy; the Newey–West estimator is used in the denominator of the Diebold–Mariano test statistic to account for autocorrelation in the multistep forecasts. Our benchmark against which we assess the MSMD and LMSD models is the short-memory ACD, and we compare models with exponential and Weibull innovations separately.

Tables 8 and 9 report the results of the forecasting performance of the different methodologies. Although there is no uniform ranking across the currencies, forecast horizons, and loss functions, a few clear patterns emerge from the exercise. Both the LMSD and MSMD forecasts generally outperform the ACD forecasts in terms of both the MSE and MAD. The gains in forecasting performance increase with the forecast horizon and are generally statistically significant at the 5% level. The MSMD model performs better when the parameters are estimated by maximum likelihood and the optimal forecasting rule is used, but the linear forecasting scheme coupled with parameter estimates obtained by Whittle estimation also deliver better performance than the ACD, although the difference is not always statistically significant. The superior in-sample fit of the models with Weibull innovations that we documented in the previous section does not necessarily translate into better out-of-sample performance. Similarly, while the MSMD(8) has a higher log-likelihood in-sample, it does not always outperform the MSMD(6) specification.

The LMSD and MSMD models generally perform on par if optimal forecasting and MLE estimates are used for the latter model, with the MSMD sometimes producing slightly better results. The forecast combinations of the LMSD and MSMD models almost always significantly outperform the ACD model, and this is generally true regardless of the estimation method and forecasting rule used for the MSMD model. This is a potentially important result for practitioners, for the Whittle estimators of both MSMD and LMSD parameters are very easy to implement regardless of the size of the sample or the various

Table 8. Mean square error (MSE) of out-of-sample forecasts for (A) Swiss franc, (B) Euro, and (C) Japanese Yen futures price durations for 1, 5, 10, and 20 step ahead horizons. We use * and ** to denote that a competing model has significantly lower MSE (is better) in comparison to the ACD model at 95% and 99% significance levels, respectively. We use † and †† to denote that a competing model has significantly higher MSE (is worse) in comparison to the ACD model at 95% and 99% significance levels, respectively.

	h = 1		h = 5		h = 10		h = 20	
	Exp	$\mathcal{W}(\kappa)$	Exp	$\mathcal{W}(\kappa)$	Exp	$\mathcal{W}(\kappa)$	Exp	$\mathcal{W}(\kappa)$
A. Swiss franc								
ACD	1.289	1.243	11.652	9.815	39.372	28.899	159.352	91.237
LMSD	1.213**	1.248	8.900**	9.061	25.760**	25.806*	79.218**	77.916*
MSMD(6) MLE	1.221**	1.206*	8.843**	8.378**	23.981**	23.992*	77.708**	78.975
MSMD(8) MLE	1.219**	1.215	8.781**	8.292**	23.628**	23.109**	75.082**	76.671*
MSMD(6) Bin-Whittle	1.235**	1.269	9.645**	10.088	27.539**	27.989	85.064**	84.511
MSMD(8) Bin-Whittle	1.236**	1.270	9.656**	10.102	27.583**	28.016	85.250**	84.565
MSMD(6) Log-Whittle	1.236**	1.273	9.682**	10.147	27.658**	28.097	85.467**	84.721
MSMD(8) Log-Whittle	1.236**	1.273	9.684**	10.147	27.674**	28.098	85.560**	84.720
MSMD(6)+MSMD(8)+LMSD MLE	1.211**	1.213**	8.708**	8.427**	23.939**	23.202**	74.231**	73.393**
MSMD(6)+MSMD(8)+LMSD Bin-Whitt.	1.215**	1.224**	8.672**	8.503**	26.015**	25.770**	80.372**	78.770**
MSMD(6)+MSMD(8)+LMSD Log-Whitt.	1.216**	1.225**	8.691**	8.346**	26.072**	25.820**	80.561**	78.852**
B. Euro								
ACD	1.636	1.635	13.495	13.476	37.628	37.565	106.023	105.897
LMSD	1.570**	1.577*	12.219**	12.424*	32.266**	32.501**	86.183**	86.461**
MSMD(6) MLE	1.558**	1.572**	12.004**	12.599**	31.103**	33.614*	91.264*	92.473*
MSMD(8) MLE	1.562**	1.557**	12.135**	12.083**	31.678**	31.452**	90.499*	90.801*
MSMD(6) Bin-Whittle	1.594**	1.635	12.972*	13.260	35.308**	35.093*	98.206*	95.891*
MSMD(8) Bin-Whittle	1.594**	1.637	12.985*	13.274	35.357**	35.113*	98.348*	95.914*
MSMD(6) Log-Whittle	1.595**	1.638	13.020*	13.275	35.475**	35.105*	98.655*	95.888*
MSMD(8) Log-Whittle	1.595**	1.638	13.022*	13.286	35.484**	35.123*	98.689**	95.911*
MSMD(6)+MSMD(8)+LMSD MLE	1.556**	1.562**	11.980**	12.306**	31.162**	31.779**	86.533**	87.319**
MSMD(6)+MSMD(8)+LMSD Bin-Whitt.	1.571**	1.573**	12.017**	12.161**	33.276*	32.811**	91.385**	89.567**
MSMD(6)+MSMD(8)+LMSD Log-Whitt.	1.571**	1.574**	11.987**	12.195**	33.355**	32.812**	91.590**	89.556**
C. Japanese Yen								
ACD	1.769	1.771	14.624	14.658	43.349	43.432	133.613	133.654
LMSD	1.714**	1.793	13.149**	13.694**	36.577**	37.600**	107.462**	109.681**
MSMD(6) MLE	1.712**	1.715*	12.984**	12.685**	34.812**	34.280**	103.209**	102.523**
MSMD(8) MLE	1.719*	1.730	13.184**	12.960**	35.509**	34.777**	102.423**	103.065**
MSMD(6) Bin-Whittle	1.756	1.829†	13.947*	14.254	40.057**	39.714**	119.285**	114.268**
MSMD(8) Bin-Whittle	1.757	1.829†	13.968*	14.246	40.119**	39.675**	119.399**	114.092**
MSMD(6) Log-Whittle	1.760	1.838†	14.042	14.317	40.378**	39.815**	120.163**	114.273**
MSMD(8) Log-Whittle	1.760	1.835†	14.288	14.336**	40.336**	39.736**	119.983**	114.055**
MSMD(6)+MSMD(8)+LMSD MLE	1.706**	1.735	12.883**	12.988**	34.832**	34.727**	100.499**	100.807**
MSMD(6)+MSMD(8)+LMSD Bin-Whitt.	1.720**	1.745	12.970**	13.176**	37.390**	36.752**	111.419**	107.634**
MSMD(6)+MSMD(8)+LMSD Log-Whitt.	1.722**	1.749	12.934**	12.950**	37.533**	36.779**	111.809**	107.578**

Table 9. Mean absolute deviations (MAD) of out-of-sample forecasts for (A) Swiss franc, (B) Euro, and (C) Japanese Yen Futures price durations for 1, 5, 10 and 20 step ahead horizons. We use * and ** to denote that a competing model has significantly lower MAD (is better) in comparison to the ACD model at 95% and 99% significance levels, respectively. We use † and †† to denote that a competing model has significantly higher MAD (is worse) in comparison to the ACD model at 95% and 99% significance levels, respectively.

	$h = 1$		$h = 5$		$h = 10$		$h = 20$	
	Exp	$\mathcal{W}(\kappa)$	Exp	$\mathcal{W}(\kappa)$	Exp	$\mathcal{W}(\kappa)$	Exp	$\mathcal{W}(\kappa)$
A. Swiss franc								
ACD	0.752	0.741	2.411	2.227	4.469	3.843	8.976	6.824
LMSD	0.770††	0.784††	2.249**	2.261	3.870**	3.869	6.862**	6.791
MSMD(6) MLE	0.742*	0.724**	2.153**	2.084**	3.608**	3.488**	6.730**	6.567
MSMD(8) MLE	0.741*	0.738	2.148**	2.095**	3.600**	3.496**	6.628**	6.578
MSMD(6) Bin-Whittle	0.741*	0.750†	2.226**	2.270	3.789**	3.828	6.747**	6.767
MSMD(8) Bin-Whittle	0.741*	0.751†	2.227**	2.271	3.791**	3.830	6.751**	6.768
MSMD(6) Log-Whittle	0.741*	0.751†	2.229**	2.275	3.795**	3.834	6.756**	6.771
MSMD(8) Log-Whittle	0.741*	0.751†	2.229**	2.275	3.795**	3.834	6.756**	6.771
MSMD(6)+MSMD(8)+LMSD MLE	0.748	0.746	2.177**	2.134**	3.653**	3.541**	6.605**	6.463**
MSMD(6)+MSMD(8)+LMSD Bin-Whitt.	0.746	0.753	2.174**	2.155**	3.750**	3.743**	6.677**	6.640**
MSMD(6)+MSMD(8)+LMSD Log-Whitt.	0.747	0.753	2.161**	2.117**	3.753**	3.746**	6.681**	6.642**
B. Euro								
ACD	0.789	0.788	2.515	2.513	4.303	4.299	7.462	7.458
LMSD	0.795	0.810††	2.451	2.474	4.090*	4.131	6.808**	6.861**
MSMD(6) MLE	0.777**	0.784	2.381**	2.455	3.986**	4.170	7.096	7.193
MSMD(8) MLE	0.777**	0.780	2.390**	2.390**	3.984**	3.996**	7.029	7.048
MSMD(6) Bin-Whittle	0.779*	0.790	2.506	2.545	4.265	4.280	7.363	7.327
MSMD(8) Bin-Whittle	0.779*	0.790	2.507	2.547	4.267	4.281	7.367	7.328
MSMD(6) Log-Whittle	0.780*	0.790	2.509	2.547	4.272	4.281	7.376	7.327
MSMD(8) Log-Whittle	0.780*	0.790	2.509	2.548	4.272	4.282	7.377	7.328
MSMD(6)+MSMD(8)+LMSD MLE	0.780*	0.789	2.401**	2.442*	3.980**	4.050*	6.872*	6.930*
MSMD(6)+MSMD(8)+LMSD Bin-Whitt.	0.779*	0.785	2.401*	2.422	4.134**	4.131*	7.057**	7.018*
MSMD(6)+MSMD(8)+LMSD Log-Whitt.	0.780*	0.785	2.392*	2.421	4.137**	4.132*	7.063**	7.019*
C. Japanese Yen								
ACD	0.855	0.855	2.801	2.804	4.689	4.694	8.348	8.349
LMSD	0.880††	0.908††	2.723	2.776	4.384**	4.451*	7.581**	7.688**
MSMD(6) MLE	0.844*	0.840*	2.631**	2.598**	4.207**	4.164**	7.409**	7.389**
MSMD(8) MLE	0.846*	0.845	2.648**	2.624**	4.204**	4.204**	7.415**	7.415**
MSMD(6) Bin-Whittle	0.855	0.874†	2.744**	2.772	4.530**	4.513*	7.976**	7.819**
MSMD(8) Bin-Whittle	0.855	0.874†	2.746*	2.771	4.532**	4.511*	7.979**	7.813**
MSMD(6) Log-Whittle	0.856	0.876††	2.751	2.777	4.543**	4.517*	8.001**	7.819**
MSMD(8) Log-Whittle	0.856	0.875††	2.750	2.775	4.541**	4.513*	7.996**	7.812**
MSMD(6)+MSMD(8)+LMSD MLE	0.853	0.861	2.656**	2.667**	4.229**	4.230**	7.328**	7.349**
MSMD(6)+MSMD(8)+LMSD Bin-Whitt.	0.857	0.869†	2.661**	2.684**	4.382**	4.349**	7.677**	7.573**
MSMD(6)+MSMD(8)+LMSD Log-Whitt.	0.858	0.870†	2.646**	2.649**	4.388**	4.350**	7.689**	7.572**

distributional assumptions made. Thus, we conclude that the long-memory duration models do provide better forecasts than the simple short-memory ACD model.

8. Conclusion

This article introduces a new model for financial durations, featuring persistence that translates from durations to realized volatility. We establish the main properties of the model and propose the Whittle estimator of its parameters as an alternative to maximum likelihood. The asymptotics we obtain for the estimator is by no means confined to the MSMD specifications explored in this article, and can be readily adapted to other MSM applications, such as stochastic volatility modeling. In an empirical application, we show that the MSMD model performs well in multistep forecasting.

There are several avenues for future research. It would be worthwhile to experiment with the “enhanced” Whittle estimator proposed by Deo et al. (2006) in order to improve the finite-sample properties. The idea of this approach is to apply the Whittle estimator to durations transformed as $\frac{1}{\nu}X^\nu$, $\nu > 0$, rather than to logarithmic durations as we did in this article, as the distribution of $\frac{1}{\nu}X^\nu$ may be closer to Gaussian for some ν than the distribution of $\log X$. Abadir and Tailman (2005) show that this transformation is smooth in the sense that the autocovariance function of $\frac{1}{\nu}X^\nu$ approaches the autocovariance function of $\log X$ as $\nu \rightarrow 0$. The moments and spectral density of $\frac{1}{\nu}X^\nu$ can be obtained in closed form, which facilitates implementation.

On the empirical side, it would be interesting to use the MSMD model in various risk-management applications. Given the success of the model in multistep forecasting, one may for example explore its ability to forecast realized volatility over short time-horizons, such as 1 hour, and compare the resulting forecasts with those obtained from popular time-series models for realized volatility. Similarly, the model may be fit to volume durations, and used to predict market trading activity with the aim to optimally time trade execution. We will explore these applications in future work.

Appendix: Proofs

A.1. Proof of Proposition 1

We follow (Zaffaroni, 2009, Theorem 1) and rely on the proof of (Hannan, 1973, Lemma 1). In particular, we show that $\sup_{\theta \in \Theta} |Q_n(\theta) - Q(\theta)| \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$, where $Q(\theta) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \log(f(\omega; \theta)) d\omega + \frac{1}{2\pi} \int_{-\pi}^{\pi} f(\omega; \theta_0) f^{-1}(\omega; \theta) d\omega$, and that $Q(\theta) \geq Q(\theta_0)$ for any $\theta \in \Theta$ with equality holding only if $\theta = \theta_0$. The statement in the proposition then follows.

Starting with $\frac{1}{n} \sum_{i=1}^{n-1} I_n(\omega_i) / f(\omega_i; \theta)$, observe that the continuity and boundedness away from zero of $f(\omega; \theta)$ on $[-\pi, \pi]$ implies that the Cesaro sum of the Fourier series of $f^{-1}(\omega; \theta)$ taken to M terms converges to $f^{-1}(\omega; \theta)$ uniformly on $[-\pi, \pi] \times \Theta$ as $M \rightarrow \infty$ (e.g., Brockwell and Davis, 1991, Theorem 2.11.1). Thus the uniform convergence a.s. of $\frac{1}{n} \sum_{i=1}^{n-1} I_n(\omega_i) / f(\omega_i; \theta)$ follows by the same argument as in Hannan (1993, Lemma 1), provided that (a) $\frac{1}{n} \sum_{i=h+1}^{n-1} x_i x_{i-h} \xrightarrow{a.s.} E(x_i x_{i-h})$ as $n \rightarrow \infty$ for all h , $0 \leq h \leq M$, and (b) $x_j x_{n-m+j} / n \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$ for any fixed $j > 0$ and $m > 0$. Since $\{x_t\}$ is strictly stationary, ergodic and has finite and constant second moments, the almost sure convergence in (a) and (b) follows from the Ergodic Theorem (e.g. Davidson, 1994, Theorem 13.12). Uniform convergence of the non-random term $\frac{1}{n} \sum_{i=1}^{n-1} \log f(\omega_i; \theta)$ follows by the same argument as in Zaffaroni (2009, Lemma 6). Finally, for all $\theta \in \Theta$ such that $\theta \neq \theta_0$, the first assumption in the proposition implies that $f(\omega; \theta_0) / f(\omega; \theta) - 1 > \log f(\omega; \theta_0) - \log f(\omega; \theta)$ on a set of positive Lebesgue measure, and hence $Q(\theta) > Q(\theta_0)$ unless $\theta = \theta_0$.

A.2. Proof of Proposition 2

Define $\delta_j = 1 - \gamma_j$, $\eta_{j,i} = (m_{j,i} - E(m)) - \delta_j(m_{j,i-1} - E(m))$, $j = 1, \dots, k$, $\eta_{0,i} = \log \epsilon_i - E(\log \epsilon_1)$, $\psi_h = (\mathbf{1}_{\{h=0\}}, \delta_1^h, \delta_2^h, \dots, \delta_k^h)$, $\boldsymbol{\eta}_i = (\eta_{0,i}, \eta_{1,i}, \dots, \eta_{k,i})'$, and $\mu = \bar{\psi} + E(\log \epsilon_1) + kE(m)$, and observe that x_i can be written as

$$x_i = \mu + \sum_{l=0}^{\infty} \psi_l \boldsymbol{\eta}_{i-l}, \quad i \in \mathbb{Z}. \tag{A.1}$$

Now define $\mathcal{F}_{j,i} = \sigma(\eta_{j,i}, \eta_{j,i-1}, \dots)$, $j = 0, \dots, k$ and observe that $E|\eta_{j,i}| < \infty$, $E(\eta_{j,i}|\mathcal{F}_{j,i-1}) = 0$ and $E(\eta_{j,i}^2) = c < \infty$, so $\{\eta_{j,i}, \mathcal{F}_{j,i}\}$ is a homoskedastic martingale difference sequence, $j = 1, \dots, k$. Now $\epsilon_i - E(\epsilon_1)$ is *iid* and the components of $\boldsymbol{\eta}_i$ are independent, and hence $\{\boldsymbol{\eta}_i, \mathcal{F}_i^\eta\}$, where $\mathcal{F}_i^\eta = \otimes_{j=0}^k \mathcal{F}_{j,i}$, is also homoskedastic martingale difference. Thus the statement in the proposition follows from (Hosoya and Taniguchi, 1982, Theorem 3.1) provided we verify that (a) the model spectral density satisfies the regularity conditions required by the theorem, and (b) the conditions (i)–(v) of the theorem are satisfied. Note that the nonzero constant term μ in (A.1) is inconsequential since the periodogram is evaluated at Fourier frequencies, and hence the constant can be ignored in the sequel. If need be, it can be consistently estimated by the sample mean of x_i .

(a) *Regularity conditions for spectral density.* We need to show that the model spectral density is bounded away from zero on $[-\pi, \pi]$, twice continuously differentiable function in $\boldsymbol{\theta}$ on Θ , and that the second derivatives are continuous in $\omega \in [-\pi, \pi]$. This follows directly from (3.2) and the fact that Θ is a compact subset of the parameter space (i.e., γ_k is bounded away from zero and one and b is bounded away from one).

(b.i) *Conditional second moments.* We need to show that for each j_1, j_2 , and q ,

$$\text{Var}[E(\eta_{j_1,i} \eta_{j_2,i+h} | \mathcal{F}_{i-q}^\eta) - \mathbf{1}_{\{j_1=j_2\}} \sigma_{\eta_{j_1}}^2] = O(q^{-2-r}), \quad r > 0, \tag{A.2}$$

uniformly in i . Since the elements of η_i are independent and second-order stationary, this follows immediately whenever $j_1 \neq j_2$. For $j_1 = j_2$, we simplify notation by writing $j = j_1 = j_2$ and further by writing $\bar{m}_{j,i} = m_{j,i} - E(m)$, $\bar{m}_{j,i} = \bar{m}_i$, $\eta_{j,i} = \eta_i$, and $\delta = \delta_j$, $j = 1, \dots, k$. The case of $j = 0$ follows trivially since $\eta_{0,i}$ is *iid*. Now for all $h \geq 0$ and $p = 1, \dots, 4$,

$$E(\bar{m}_{i+h}^p | \bar{m}_i) = \delta^h \bar{m}_i^p + (1 - \delta^h) E(\bar{m}^p), \tag{A.3}$$

and for all $0 \leq h_1 < h_2$ and $q \geq 1$, by law of iterated expectations (LIE),

$$\begin{aligned} E(\bar{m}_{i+h_1} \bar{m}_{i+h_2} | \bar{m}_{i-q}) &= E(\bar{m}_{i+h_1} E(\bar{m}_{i+h_2} | \bar{m}_{i+h_1}) | \bar{m}_{i-q}) \\ &= \delta^{h_2+q} \bar{m}_{i-q}^2 + \delta^{h_2-h_1} (1 - \delta^{h_1+q}) E(\bar{m}^2). \end{aligned}$$

For $h \geq 1$ and $q \geq 2$, we therefore have

$$\begin{aligned} E(\eta_i \eta_{i+h} | \mathcal{F}_{i-q}^\eta) &= E(\bar{m}_i \bar{m}_{i+h} | \bar{m}_{i-q}) - \delta E(\bar{m}_{i-1} \bar{m}_{i+h} | \bar{m}_{i-q}) \\ &\quad - \delta E(\bar{m}_i \bar{m}_{i+h-1} | \bar{m}_{i-q}) + \delta^2 E(\bar{m}_{i-1} \bar{m}_{i+h-1} | \bar{m}_{i-q}) \\ &= 0. \end{aligned}$$

For $h = 0$ and any $q \geq 2$, we have

$$\begin{aligned} E(\eta_i^2 | \mathcal{F}_{i-q}^\eta) &= E(\bar{m}_i^2 | \bar{m}_{i-q}) - 2\delta E(\bar{m}_i \bar{m}_{i-1} | \bar{m}_{i-q}) + \delta^2 E(\bar{m}_{i-1}^2 | \bar{m}_{i-q}) \\ &= (\delta^q - \delta^{q+1}) \bar{m}_{i-q}^2 + [(1 - \delta^q) - \delta^2(1 - \delta^{q-1})] E(\bar{m}^2), \end{aligned}$$

and by fourth-order stationarity,

$$\begin{aligned} \text{Var}[E(\eta_i^2 | \mathcal{F}_{i-q}^\eta) - E(\eta_i^2)] &= \text{Var}[(\delta^q - \delta^{q+1})\bar{m}_{i-q}^2 + \{(1 - \delta^q) - \delta^2(1 - \delta^{q-1})\}E(\bar{m}^2) \\ &\quad - (1 - \delta^2)E(\bar{m}^2)], \\ &= (\delta^q - \delta^{q+1})^2 \text{Var}(\bar{m}_{i-q}^2 - E(\bar{m}^2)), \\ &= O(\delta^{2q}), \end{aligned}$$

which verifies (A.2) since $O(\delta^{2q})$ decays faster than $O(q^{-2-r})$ for any $r < \infty$.

(b.ii) *Conditional fourth-moments.* We need to show that for each j_1, j_2, j_3, j_4 and h_1, h_2, h_3 ,

$$E|E(\eta_{j_1, i} \eta_{j_2, i+h_1} \eta_{j_3, i+h_2} \eta_{j_4, i+h_3} | \mathcal{F}_{i-q}) - E(\eta_{j_1, i} \eta_{j_2, i+h_1} \eta_{j_3, i+h_2} \eta_{j_4, i+h_3})| = O(q^{-1-r}), \quad r > 0, \tag{A.4}$$

uniformly in i , where $0 \leq h_1 \leq h_2 \leq h_3$. To save space, we establish this for $j_1 = j_2 = j_3 = j_4$, noting that the other cases can be shown analogously, and are easier since the elements of η_i are independent. We again simplify notation by writing $j = j_1 = j_2 = j_3 = j_4$, $\bar{m}_j = \bar{m}_{j,i}$, and $\delta = \delta_j$. Now

$$\begin{aligned} \eta_i \eta_{i+h_1} \eta_{i+h_2} \eta_{i+h_3} &= (\bar{m}_i - \delta \bar{m}_{i-1})(\bar{m}_{i+h_1} - \delta \bar{m}_{i+h_1-1})(\bar{m}_{i+h_2} - \delta \bar{m}_{i+h_2-1}) \\ &\quad (\bar{m}_{i+h_3} - \delta \bar{m}_{i+h_3-1}), \end{aligned} \tag{A.5}$$

where for any $0 \leq h_1 \leq h_2 \leq h_3$ the sum on the right-hand side contains sixteen terms, each being a product of four terms with either zero, one, two, three, or four lags in common depending on the particular choice of h_1, h_2 , and h_3 . Starting with the case when all the lags are distinct, we obtain by iterated conditioning and repeated use of (A.3)

$$\begin{aligned} E(\bar{m}_i \bar{m}_{i+l_1} \bar{m}_{i+l_2} \bar{m}_{i+l_3} | \bar{m}_{i-q}) &= \delta^{l_3} [\delta^q \bar{m}_{i-q}^4 + (1 - \delta^q)E(\bar{m}^4)] + \delta^{l_3-l_1} (1 - \delta^{l_1})E(\bar{m}^3) \delta^q \bar{m}_{i-q} \\ &\quad + \delta^{l_3-l_2+l_1} (1 - \delta^{l_2-l_1})E(\bar{m}^2) [\delta^q \bar{m}_{i-q}^2 + (1 - \delta^q)E(\bar{m}^2)] \\ E(\bar{m}_i \bar{m}_{i+l_1} \bar{m}_{i+l_2} \bar{m}_{i+l_3}) &= \delta^{l_3} E(\bar{m}^4) + \delta^{l_3-l_2+l_1} (1 - \delta^{l_2-l_1})E(\bar{m}^2)^2, \\ 0 &< l_1 < l_2 < l_3, q \geq 2. \end{aligned}$$

Thus by fourth-order stationarity,

$$\begin{aligned} E|E(\bar{m}_i \bar{m}_{i+l_1} \bar{m}_{i+l_2} \bar{m}_{i+l_3} | \bar{m}_{i-q}) - E(\bar{m}_i \bar{m}_{i+l_1} \bar{m}_{i+l_2} \bar{m}_{i+l_3})| \\ \leq \delta^q \{ \delta^{l_3} E|\bar{m}_{i-q}^4 - E(\bar{m}^4)| + \delta^{l_3-l_1} (1 - \delta^{l_1}) |E(\bar{m}^3)| E|\bar{m}_{i-q}| \} \\ + \delta^{l_3-l_2+l_1} (1 - \delta^{l_2-l_1}) E(\bar{m}^2) E|\bar{m}_{i-q}^2 - E(\bar{m}^2)| \\ = O(\delta^q). \end{aligned}$$

Turning to the case when all the lags are equal, i.e., $0 = l_1 = l_2 = l_3$, we have

$$\begin{aligned} E|E(\bar{m}_i \bar{m}_i^4 | \bar{m}_{i-q}) - E(\bar{m}_i^4)| &\leq \delta^q E|\bar{m}_{i-q}^4 - E(\bar{m}^4)|, \\ &= O(\delta^q). \end{aligned}$$

The terms with one, two, and three common lags in (A.5) can be established analogously. (A.4) then follows from the triangle inequality and the fact that $O(\delta^q)$ decays faster than $O(q^{-1-r})$ for any $r < \infty$.

(b.iii) *Square-integrability of model spectral density.* We need to show that the MSMD spectral density is square-integrable. This follows directly from the continuity and boundedness of the spectral density on $[-\pi, \pi]$, uniformly on Θ .

(b.iv) *Absolute summability of fourth-order cumulants.* We need to show that

$$\sum_{h_1=-\infty}^{\infty} \sum_{h_2=-\infty}^{\infty} \sum_{h_3=-\infty}^{\infty} |\text{cum}(\eta_{j_1, h_1}, \eta_{j_2, h_2}, \eta_{j_3, h_3}, \eta_{j_4, h_4})| < \infty \quad (\text{A.6})$$

for any $j_1, j_2, j_3, j_4, j_l = 0, \dots, k$ and $l = 1, 2, 3, 4$. Since the elements of η_i are independent and fourth-order stationary all cross-cumulants are zero, so we only need to focus on the case $j_1 = j_2 = j_3 = j_4$. Now since $\eta_{0,i}$ is *iid* with finite fourth moment, the absolute summability of its fourth-order cumulants follows directly. For an $\eta_{j,i}$, $j = 1, \dots, k$, we use the fact that $\eta_{j,i}$ is fourth-order stationary exponential strong mixing with $4 + r$ moment finite and hence $\sum_{s=0}^{\infty} (s+1)^{p-2} \alpha(s)^{r/(p+r)} < \infty$, $p = 1, \dots, 4$, where $\alpha(s)$ is the strong mixing coefficient associated to $\eta_{j,i}$ satisfying $\alpha(s) = O(\rho^s)$ for some $\rho \in (0, 1)$. Thus by (Doukhan and Leon, 1989, Proposition 2.2), the inequality (A.6) holds.

(b.v) *Lipschitz continuity of $f(\omega; \theta_0)$.* This follows directly for the smoothness of the first derivatives of $f(\omega; \theta)$ uniformly on Θ .

The statement of the proposition now follows by simplifying the formulae for M_f and \tilde{V} in (Hosoya and Taniguchi, 1982, Theorem 3.1) in view of the fact that the true spectral density of x_i is given by $f(\omega; \theta_0)$, together with the fact that the matrix $M(\theta_0)$ is full rank by the identification assumption stated in Proposition 1.

A.3. Proof of Proposition 3

We need to show that the assumptions of Theorem 1 in Deo et al. (2009) are satisfied. Assumption (ii) requires the duration process to be exponential strong mixing, and this was established in Section 2.1; it is well-known that exponential β -mixing implies exponential strong mixing since the β -mixing coefficients dominate the strong mixing coefficients, e.g., Davidson, 1994, Chapter 14.

Turning to Assumption (i) of Theorem 1 in Deo et al. (2009), we first note that the long-run variance of X_i is strictly positive, i.e., $\lim_{n \rightarrow \infty} \frac{1}{n} \text{Var}(\sum_{i=1}^n X_i) = \sigma^2 > 0$ for some constant σ^2 . Since $\{X_i\}$ is stationary with exponentially declining autocovariances, this is equivalent to showing that the spectral density of X_i is strictly positive at the zero frequency. This follows directly from expression (0.2) in the Supplemental Appendix A1 since when $\omega = 0$ each summand in (0.2) is strictly positive implying $2\pi f_X(\omega) > 0$. Write $\sigma^2 = 2\pi f_X(\omega)$ and define $Y_n(s) = n^{-1/2} \sum_{i=1}^{\lfloor ns \rfloor} (X_i - \bar{\psi})$. Since by assumption ϵ_i is *iid* with $3 + r$ moments finite for some $r > 0$ and independent of ψ_i , which also has $3 + r$ moments finite, it follows that X_i has $3 + r$ moments finite. Then by Deo et al. (2009), $Y_n \Rightarrow \sigma W$, as $n \rightarrow \infty$, where W is a standard Brownian motion. This verifies Assumption (i) of Theorem 1 in Deo et al. (2009).

Finally, by the same argument as in the proof of Theorem 3 in Deo et al. (2009), the exponential strong mixing property and the existence of $3 + r$ (where $r > 0$) moments of x_i imply for $y_n = n^{-1/2} \sum_{i=1}^n x_i$, and some constant c that $E(|y_n - E(y_n)|^{3+r}) \leq c < \infty$ as required by Assumption (iii) of Theorem 1 in Deo et al. (2009).

Acknowledgments

We are grateful to the editor, associate editor, referees, and Karim M. Abadir, Laurent E. Calvet, Marcelo Fernandes, Clifford Hurvich, Paolo Zaffaroni, and seminar participants at the 18th Symposium of the Society for Nonlinear Dynamics and Econometrics at Novara (April 2010) for many useful comments and suggestions. Special thanks go to Liudas Giraitis for invaluable discussions about frequency-domain estimation and inference. The views expressed in this paper are those of the authors and not necessarily those of the Bank of England or any of its committees.

Funding

Jozef Barunik gratefully acknowledges financial support from the Czech Science Foundation under project No. 13-32263S and funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement No. FP7-SSH-612955 (FinMaP).

References

- Abadir, K. M., Tailman, G. (2005). Autocovariance functions of series and of their transforms. *Journal of Econometrics* 124:227–252.
- Andersen, T. G., Dobrev, D., Schaumburg, E. (2008). “Duration-Based Volatility Estimation,” Discussion Paper Series gd08-034. Institute of Economic Research, Hitotsubashi University.
- Bacry, E., Kozhemyak, A., Muzy, J. F. (2008). Continuous cascade models for asset returns. *Journal of Economic Dynamics and Control* 32:156–199.
- Bauwens, L., Veredas, D. (2004). The stochastic conditional duration model: a latent variable model for the analysis of financial durations. *Journal of Econometrics* 119(2):381–412.
- Brockwell, J. P., Dahlhaus, R. (2004). Generalized levinson-durbin and burg algorithms. *Journal of Econometrics* 118:129–149.
- Brockwell, P., Davis, R. (1991). *Time Series: Theory and Methods*. Berlin and Heidelberg: Springer-Verlag.
- Calvet, L. E., Fisher, A. (2004). How to forecast long-run volatility: Regime switching and the estimation of multifractal processes. *Journal of Financial Econometrics* 2:49–83.
- Carrasco, M., Chen, X. (2002). Mixing and moment properties of various garch and stochastic volatility models. *Econometric Theory* 18:17–39.
- Chen, F., Diebold, F. X., Schorfheide, F. (2012). A markov-switching multi-fractal inter-trade duration model, with application to U.S. equities. University of Pennsylvania.
- Chen, W. W., Deo, R. S. (2004). A generalized portmanteau goodness-of-fit test for time series models. *Econometric Theory* 20:382–416.
- Davidson, J. (1994). *Stochastic Limit Theory*. Oxford University Press.
- Deo, R., Hsieh, M., Hurvich, C. (2006). The persistence of memory: From durations to realized volatility. Technical report.
- Deo, R., Hsieh, M., Hurvich, C. M. (2010). Long memory in inter trade durations, counts and realized volatility of NYSE stocks. *Journal of Statistical Planning and Inference* 140:3715–3733.
- Deo, R., Hurvich, C., Lu, Y. (2006). Forecasting realized volatility using a long memory stochastic volatility model: Estimation, prediction and seasonal adjustment. *Journal of Econometrics* 131(1–2):29–58.
- Deo, R., Hurvich, C. M., Soulier, P., Wang, Y. (2009). Conditions for the propagation of memory parameter from durations to counts and realized volatility. *Econometric Theory* 25:764–792.
- Diebold, F. X., Mariano, R. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*.
- Doornik, J. A. (2006). *The object-oriented language - Ox 4*. London: Timberlake Consultants Press.
- Doukhan, P., Leon, J. R. (1989). Cumulants for stationary mixing random sequences and applications to empirical spectral density. *Probability and Mathematical Statistics* 10(1):11–26.
- Engle, R. F. (2000). The econometrics of ultra-high-frequency data. *Econometrica* 68(1):1–22.
- Engle, R. F., Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced time series data. *Econometrica* 66(5):1127–1162.
- Fernandes, M., Grammig, J. (2006). A family of autoregressive conditional duration models. *Journal of Econometrics* 130:1–23.
- Giacomini, R., White, H. (2006). Test of conditional predictive ability. *Econometrica* 74:1545–1578.
- Grammig, J., Maurer, K. O. (2000). Non-monotonic hazard functions and the autoregressive conditional duration model. *Econometrics Journal* 3:16–38.
- Hannan, E. J. (1973). The asymptotic theory of linear time-series models. *Journal of Applied Probability* 10:130–145.
- Hasbrouck, J., Saar, G. (n.d.). Low-latency trading. Johnson School Research Paper Series No 35-2010.
- Hosoya, Y., Taniguchi, M. (1982). A central limit theorem for stationary processes and the parameter estimation of linear processes. *The Annals of Statistics* 10(1):132–153.
- Huisman, R., Koedijk, K. G., Kool, C. J. M., Palm, F. (2001). Tail-index estimates in small samples. *Journal of Business and Economic Statistics* 19:208–216.
- Jasiak, J. (1998). Persistence in intertrade durations. *Finance* 19(2):166–195.
- King, M. R., Osler, C., Rime, D. (2011). Foreign exchange market structure, players and evolution. Working Paper 2011/10, Norges Bank.
- Liu, M. (2000). Modeling long memory in stock market volatility. *Journal of Econometrics* 99:139–171.
- Lunde, A. (1999). A generalized gamma autoregressive conditional duration model. Discussion Paper, Aalborg University.
- Lux, T. (2008). The markov-switching multifractal model of asset returns: GMM estimation and linear forecasting of volatility. *Journal of Business and Economic Statistics* 26(2):194–210.
- McAleer, M., Medeiros, M. C. (2008). Realized volatility: A review. *Econometric Reviews* 27(1–3):10–45.
- Mendel, J. M. (1991). Tutorial on higher-order statistics (spectra) in signal processing and spectral theory: theoretical results and some applications. *Proceedings of the IEEE* 79(3):278–305.
- Nieuwenhuis, G. (1989). Equivalence of functional limit theorems for stationary point processes and their palm distributions. *Probability Theory and Related Fields* 81:593–608.
- Oomen, R. C. A. (2006). Properties of realized variance under alternative sampling schemes. *Journal of Business & Economic Statistics* 24:219–237.

- Pacurar, M. (2008). Autoregressive conditional duration models in finance: A survey of the theoretical and empirical literature. *Journal of Economic Surveys* 22(4):711–751.
- Renault, E., Werker, B. J. (2011). Causality effects in return volatility measures with random times. *Journal of Econometrics* 160:272–279.
- Rice, J. (1973). On the estimation of the parameters of the power spectrum. *Journal of Time Series Analysis* 9:378–392.
- Rodríguez-Poo, J. M., Veredas, D., Espasa, A. (2007). Semiparametric estimation for financial durations. In: Luc Bauwens, W. P., Veredas, D., eds. *High Frequency Financial Econometrics. Recent Developments*. Springer Verlag, Physica Verlag Heidelberg.
- Shiryayev, A. (1995). *Probability*. 2 ed. Graduate Texts in Mathematics. Springer Verlag.
- Taniguchi, M. (1982). On estimation of the integrals of the fourth order cumulant spectral density. *Biometrika* (117–122).
- Tse, Y. K., Yang, T. (2010). Estimation of high-frequency volatility: Autoregressive conditional duration models approach. Research Collection School of Economics: Paper 1276.
- Whittle, P. (1962). Gaussian estimation in stationary time series. *Bulletin of the International Statistical Institute* 39:105–129.
- Zaffaroni, P. (2009). Whittle estimation of egarch and other exponential volatility models. *Journal of Econometrics* 151:190–200.