



City Research Online

City St George's, University of London

Citation: Corsi, F. & Audrino, F. (2012). Realized covariance tick-by-tick in presence of rounded time stamps and general microstructure effects. *Journal of Financial Econometrics*, 10(4), pp. 591-616. doi: 10.1093/jjfinec/nbs007

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/4437/>

Link to published version: <https://doi.org/10.1093/jjfinec/nbs007>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Realized Covariance Tick-by-Tick in Presence of Rounded Time Stamps and General Microstructure Effects

Fulvio Corsi^{a*} and Francesco Audrino^a

^aUniversity of St. Gallen

Revised: January 2012

*Address for correspondence: Bodanstrasse 6, St. Gallen, Switzerland. Tel: +41712242431 Fax: +41712242894. Email: fulvio.corsi@usi.ch. The authors would like to thank Tim Bollerslev, the editor Eric Renault, an anonymous associate editor, and two anonymous referees for helpful comments that helped us to improve the presentation of the results.

Abstract

This paper presents two classes of tick-by-tick covariance estimators adapted to the case of rounding in the price time stamps to a frequency lower than the typical arrival rate of tick prices. Through Monte Carlo simulations we investigate the behavior of such estimators under realistic market microstructure conditions analogous to those of the financial data examined in this paper's empirical section, that is, non-synchronous trading, general ARMA structure for microstructure noise, and true lead-lag cross-covariance. Simulation results show the robustness of the proposed tick-by-tick covariance estimators to time stamp rounding, and their overall performance is superior to competing covariance estimators under empirically realistic microstructure conditions. These results are confirmed in the empirical application where the economic benefits of the proposed estimators are evaluated with volatility timing strategies applied to a bivariate portfolio of S&P 500 futures and 30-year US treasury bond futures.

JEL classification: C13; C22; C51; C53

Keywords: High frequency data; Realized covariance; Market microstructure; Bias correction; Portfolio selection; Volatility timing.

1 Introduction

Asset return covariance plays a prominent role in many important theoretical as well as practical problems in finance. Analogous to the realized volatility approach (Andersen, Bollerslev, Diebold, and Labys, 2001, 2003,; Barndorff-Nielsen and Shephard, 2001, 2002a, 2002b, 2005; Comte and Renault, 2001), the idea of employing high frequency data in the computation of daily (or lower frequency) covariance between two assets leads to the concept of *realized covariance* (or covariation).

The standard way to compute realized covariance is to choose a time interval, construct an artificially regularly spaced time series by using an interpolation scheme, and then take the contemporaneous sample covariance of those regularly spaced returns. However, simulations and empirical studies indicate that such a covariance measure is biased toward zero and rapidly increases with the reduction of the time length of the fixed interval chosen. Similar to the construction of realized volatility, the presence of market microstructure can induce significant bias in the standard realized covariance measure, but the microstructure effects responsible for this bias are different. Bid-ask bouncing, the major source of bias for the realized volatility, will merely increase the variance of the covariance estimator, but it will not induce any bias. By contrast, the so-called non-synchronous trading effect (Lo and MacKinlay, 1990) strongly affects the estimation of the realized covariance and correlation. Because the sampling from the underlying stochastic process is different for different assets, assuming that two time series are sampled simultaneously (when the sampling is non-synchronous) gives rise to the non-synchronous trading effect. As a result, covariances and correlations measured with high frequency data will possess a bias toward zero that increases as the sampling frequency increases. This dramatic drop of the absolute value of correlations among stocks when increasing the sampling frequency was first reported by Epps (1979). Since then, this effect has been confirmed by many other authors using real data and simulations, such as Dacorogna and Lundin (1999), Renó (2003), and Martens (2004). Various estimators have been proposed that try to correct for this bias using regularly interpolated returns: see, for example, Scholes and William (1977), Cohen et al. (1983) and Lo and MacKinlay (1990).

Instead, an unbiased tick-by-tick realized covariance estimator that does not rely on any

interpolation scheme, has been proposed and formally analyzed under the assumption of no microstructure noise by Hayashi and Yoshida (2005).¹ This estimator was recently investigated by Palandri (2006), Voev and Lunde (2007), Hautsch et al. (2009), and Griffin and Oomen (2011).

The application of the Hayashi and Yoshida estimator requires precise knowledge of the time stamps of every tick. Unfortunately, the time stamps of tick-by-tick data are often rounded at some minimum time frequency (one second or one minute). Depending on the liquidity of the asset, it is possible (and very likely, especially for the one-minute rounding) that a sequence of successive prices will be recorded with identical time stamps. For example, important databases that use this type of format include the index, bond gold and oil futures from Price-data.com (an example of which is shown in Table 1) and Tickdata.com and individual stocks, index futures, and options from the Nikkei Needs tick database.

Insert Table 1 about here

For tick-by-tick realized volatility computation, this time stamp rounding does not pose any problems as long as the true order of the prices is preserved. However, for tick-by-tick realized covariance, even if the order of the prices is kept for each asset, the rounding of the time stamps precludes the knowledge of the correct time ordering among the ticks of the two series inside the minimum time interval, which is a necessary condition for the application of the Hayashi and Yoshida tick-by-tick estimator.

This paper proposes two tick-by-tick covariance estimators adapted to the case of rounding in the price time stamps: the First–Last and the Needlework estimators. Although the time stamp rounding induces a natural calendar time grid, these methods may be seen as a general way of computing Hayashi and Yoshida type estimators at any given lower frequency, thus considerably extending the applicability of the proposed estimators and improving accuracy in presence of strong market microstructure noise.

We investigate, through Monte Carlo simulations, the behavior of the thus modified tick-by-tick estimators under a variety of market microstructure structures and under conditions

¹This also independently appeared (in a much less formalized version) in Martens (2004) as a simple, more efficient version of the De Jong and Nijman estimator in the absence of true lead-lag cross-covariance.

analogous to those of the financial data studied in the empirical section of this paper. In particular, we consider non-synchronous trading, a wide range of different levels of noise to signal ratio, general ARMA structure for the microstructure noise, and true lead-lag cross-covariance. Our main finding is that the proposed tick-by-tick estimators clearly outperform several different alternatives proposed in the literature.

This finding is confirmed in the empirical application where, in the spirit of the methodology suggested by West et al. (1993) and Fleming et al. (2001, 2003), we compare the utility level obtained by different variance-covariance estimators in the context of portfolio allocation. We find that a risk-averse investor who uses a conditional mean-variance optimization rule to reallocate funds daily across stocks (S&P 500) and bonds (30-year US Treasury bond) always reaps an economic gain by switching from standard covariance estimators to the two tick-by-tick covariance estimators proposed. Often such economic gain is substantial.

The remainder of the paper is organized as follows: Section 2 describes the classical tick-by-tick realized covariance estimator introduced in the literature. Section 3 defines the modified tick-by-tick estimators adapted to the case of rounded time stamps. Section 4 shows the results of realistic Monte Carlo simulations, and Section 5 presents the empirical application where the economic benefits of the proposed estimators are evaluated with volatility timing strategies applied to the bivariate portfolio of S&P 500 futures and 30-year US treasury bond futures. Section 6 summarizes and concludes.

2 Realized covariance tick-by-tick

Contrary to standard approaches, the tick-by-tick realized covariance estimator does not rely on the construction of a regular grid since it is based on the whole tick-by-tick raw data series. This approach has the twofold advantage of exploiting all the information available in the data and the ability to avoid the bias toward zero of the realized covariance. The non-synchronous trading effect produces a bias in the usual covariance measure as a consequence of the synchronization of the two series, that is, as a consequence of the construction of a regular grid in physical time.

The bias of the covariance estimator based on fixed interval returns can be seen as arising

from two distinct effects. First, the absence of trading on one asset in a certain interval produces a zero return for that interval and then artificially imposes a zero value on the cross-product of returns producing a bias toward zero in the realized covariance (which, in its standard version, is simply the sum of those cross-products). Second, the construction of a regular grid, depending on the frequency of tick arrivals, affects the computation of the realized covariance. For more liquid assets with higher average arrival rates, the last tick to fall in a certain grid interval is typically much closer to the end point of the grid compared to that of a less liquid asset. Any difference in the time stamps between these last ticks in the grid for the two assets will correspond to a portion of the cross-product returns, which will not be accounted for in the computation of the covariance. This occurs because, for the more liquid asset, the (unobserved) returns corresponding to this time difference will be imputed to the current grid interval; while for the less liquid asset, this portion of returns will be ascribed to the next grid interval, so that the two will no longer be matched and their contribution to the cross-product's sum will be lost (see Figure 1). This lost portion of covariance in each interval also produces a bias towards zero in the realized covariance computed with a regular grid, a bias that will also increase with the number of intervals and hence with the frequency.

Insert Figure 1 about here

Under the assumption of efficient markets (that is, markets that have no leads and lags cross-covariance between the true latent efficient price of the two assets) and no microstructure effects, Hayashi and Yoshida (2005) formally proved that an unbiased and consistent covariance estimator can be computed simply by summing all the cross-products of returns that have a non-zero overlapping of their respective time span. In other words, a given tick-by-tick return on one asset is multiplied with any other tick-by-tick return of the other asset that has a non-zero overlap in time, that is, which shares (even for a very small fraction) the same time interval.

Analytically, given a standard continuous time process for asset j

$$dp_j(t) = \mu_j(t)dt + \sigma_j(t)dW_j(t), \tag{1}$$

where $dW_j(t)$ is a Wiener process, and discrete price observations $\{p_j(n_{j,q})\}_{q=0,1,2,\dots,M}$ at

times $n_{j,q}$ with associated tick returns

$$r_{j,q} = p_j(n_{j,q}) - p_j(n_{j,q-1}) \quad (2)$$

the Hayashi and Yoshida realized covariance estimator for two assets i and j and a given time interval t (for example one day), is defined as:

$$RC_t = \sum_{s=1}^{M_{i,t}} \sum_{q=1}^{M_{j,t}} r_{i,s} r_{j,q} I(\delta_{q,s} > 0) \quad (3)$$

with $M_{i,t} + 1$, and $M_{j,t} + 1$ the total number of ticks on time interval t for asset i and j respectively, $I(\cdot)$ the indicator function, and

$$\delta_{q,s} = \max(0, \min(n_{i,s}, n_{j,q}) - \max(n_{i,s-1}, n_{j,q-1})) \quad (4)$$

the overlap in time between any two tick returns $r_{i,s}$ and $r_{j,q}$.

The Hayashi and Yoshida estimator is unbiased because no portion of covariance will be lost; while the portion of cross-product that does not overlap will have zero mean. Avoiding the noise and the discarding of price observations caused by the regular grid interpolation will considerably reduce the variance of the estimator. Nevertheless, in the presence of a fixed amount of market microstructure noise and under the assumption of no cross-correlated noise structures, the estimator in this form will not be consistent because, although unbiased, its variance will diverge as the number of observations tends to infinity.

The Hayashi and Yoshida estimator, as any other covariance estimator applied bivariately to a large dataset, will not provide a variance-covariance matrix which is guaranteed to be positive definite. In fact, estimation error on the elements of the covariance matrix implies that the largest sample eigenvalues of the matrix are biased upwards, while the smallest ones are biased downwards. Then, in some cases, the smallest eigenvalues can become negative so that the matrix will no longer be definite positive. To recover the positive definiteness one could apply a spectral decomposition and impose positivity to all the eigenvalues. Another general and convenient way to address this problem would be through the so called shrinkage estimator proposed by Frost and Savarino (1986) and Ledoit and Wolf (2003, 2004a, 2004b) which also reduces the variance-covariance matrix estimation error (see Corsi, 2005, for the application of shrinkage to tick-by-tick realized covariance matrices with a flexible target matrix and an asset dependent shrinkage intensity depending on the number of ticks).

3 Realized covariance with rounded time stamps

We propose simple modified tick-by-tick estimators designed to overcome or alleviate the imprecise time stamps problem discussed above that renders the application of the Hayashi and Yoshida realized covariance estimator unfeasible. In order to define them, we first establish the following notation. Time stamps are discretized on a regular grid $\tau = \{h\Delta\}_{h=1,2,\dots,H}$ where Δ is the rounding frequency (one minute on our data):

$$n_{j,\tau}^{(F)} \equiv \min_q \{n_{j,q} : n_{j,q} \geq \tau - 1\} \quad \text{and} \quad n_{j,\tau}^{(L)} \equiv \max_q \{n_{j,q} : n_{j,q} \leq \tau\} \quad (5)$$

$$p_{j,\tau}^{(F)} \equiv p_j \left(n_{j,\tau}^{(F)} \right) \quad \text{and} \quad p_{j,\tau}^{(L)} \equiv p_j \left(n_{j,\tau}^{(L)} \right) \quad (6)$$

so that $p_{j,\tau}^{(F)}$ is the first and $p_{j,\tau}^{(L)}$ is the last price inside the time interval τ with identical time stamps.

For each asset we also define two tick-types of return series:

$$r_{j,\tau}^{(F)} = p_{j,\tau}^{(F)} - p_{j,\tau-1}^{(F)} \quad \text{and} \quad r_{j,\tau}^{(L)} = p_{j,\tau}^{(L)} - p_{j,\tau-1}^{(L)} \quad (7)$$

one is constructed only from $p_\tau^{(F)}$ (for any τ) and the other one is constructed using only $p_\tau^{(L)}$. These two returns series can be seen as the returns of two different regular grid interpolation schemes: one using next-tick interpolation $r_{j,\tau}^{(F)}$ and the other employing previous-tick interpolation $r_{j,\tau}^{(L)}$.

3.1 The First–Last estimator

The first estimator we propose to overcome the problem due to rounded time stamps in the data combines the two different interpolation schemes (previous and next tick), introduced above, in the following way:

$$FL_t = \frac{1}{2} \left(\sum_{s=1}^{M_{i,t}^{(F)}} \sum_{q=1}^{M_{j,t}^{(L)}} r_{i,s}^{(F)} r_{j,q}^{(L)} I(\delta_{q,s} > 0) + \sum_{s=1}^{M_{i,t}^{(L)}} \sum_{q=1}^{M_{j,t}^{(F)}} r_{i,s}^{(L)} r_{j,q}^{(F)} I(\delta_{q,s} > 0) \right). \quad (8)$$

It can be seen as the average of two Hayashi and Yoshida-type estimators: one applied to the return series constructed with next-tick interpolation (that is, taking the first tick) for the first series and previous-tick interpolation (which considers the last tick) for the second series (see Figure 2), and doing exactly the contrary for the second estimator. We call this realized covariance estimator the First–Last tick-by-tick covariance estimator.

Insert Figure 2 about here

The reason we only consider the first $p_\tau^{(F)}$ and last $p_\tau^{(L)}$ price tick inside each bin is that among these type of ticks we can most reasonably ascertain their correct order (by assuming only that $n_{i,\tau}^{(F)} < n_{j,\tau}^{(L)}$ and $n_{j,\tau}^{(F)} < n_{i,\tau}^{(L)}$) and maximize the time overlap between the corresponding return intervals.

As long as $n_{i,\tau}^{(F)} < n_{j,\tau}^{(L)}$ and $n_{j,\tau}^{(F)} < n_{i,\tau}^{(L)}$, and assuming no correlated noise structures, the two Hayashi and Yoshida-type of estimators in the First–Last are unbiased. Since it is the average of two Hayashi and Yoshida estimators, the First–Last estimator clearly has the same statistical properties as the Hayashi and Yoshida one; in particular, it is unbiased and, in the absence of market microstructure noise, consistent as $\Delta \rightarrow 0$ and $H \rightarrow \infty$ (for details see Hayashi and Yoshida, 2005 and Griffin and Oomen, 2011).

To improve the asymptotic properties of the estimators in the presence of market microstructure noise, we apply the methodology introduced by Zhang et al. (2005) based on sub-sampling, that is averaging the estimators obtained by considering several different subgrids (i.e. so-called subsamples) of the original grid for some determined average size $\overline{H}$, with $\overline{H} \leq H$.²

In more details, let us assume that the original grid τ is partitioned in K non-overlapping subgrids $\tau^{(k)}$, $k = 1, \dots, K$, where the natural way to choose the k th subgrid $\tau^{(k)}$ is

$$\tau^{(k)} = \{k\Delta, (k+K)\Delta, (k+2K)\Delta, \dots, (k+(H_k-1)K)\Delta\},$$

where H_k is the integer making $(k+H_kK)\Delta$ the last element of $\tau^{(k)}$. We let H_k be the number of the observations in the subgrid $\tau^{(k)}$. In general H_k need not to be the same across k . Let us denote the sub-sampled version of the FL_t estimator as

$$FL_t^{(\text{sub})} = \frac{1}{K} \sum_{k=1}^K FL_t^{(k)} \quad (9)$$

where $FL_t^{(k)}$ is the First–Last estimator constructed on the subgrid $\tau^{(k)}$.

By combining tick–by–tick features of Hayashi and Yoshida with features of calendar time sampling, the First–Last estimator provides a way to sub-sampling the Hayashi and Yoshida

²For a general theory and application of the sub-sampling scheme see also Kalnina (2010).

estimator in presence of rounded time stamps. This mixed tick-time/calendar time feature permits to extend the applicability of the proposed estimator since it could be applied not only to the calendar time grid naturally induced by rounding, but to any selected time grid. For example, by setting $K = 30$ we would estimate realized covariance using data on a 30-minute grid.

Under standard regularity assumptions on the price processes (1) and on the multivariate volatility process, the proposed $FL_t^{(\text{sub})}$ estimator achieves consistency even in the presence of microstructure noise, that is the $FL_t^{(\text{sub})}$ estimator converges in probability to the true unknown integrated covariance as $\Delta \rightarrow 0$, $H \rightarrow \infty$, and $\frac{H}{K} \rightarrow \infty$. In fact, using a different aggregation scheme, Palandri (2006) and Voev and Lunde (2007) show that under the assumption of i.i.d. microstructure noise the sub-sampling of Hayashi and Yoshida over a number of subgrids that grows proportionally with the number of observations achieves consistency of the estimator by balancing between the discretization and noise errors (as in Zhang et al. 2005). Sub-sampling has also the practical consequence of improving on the estimation accuracy when market microstructure noise becomes large relative to the volatility signal of the true price (i.e. for large values of noise to signal ratios).

Although no problems are apparent in the empirical application, in theory it can occur that $n_{i,\tau}^{(F)} > n_{j,\tau}^{(L)}$ or $n_{j,\tau}^{(F)} > n_{i,\tau}^{(L)}$. In this case, the First-Last estimator may still suffer an attenuation bias. A simple solution would be to skip one time stamp. Consequently, however, we would further reduce the frequency and the number of employed returns, diminishing the precision of the estimator in presence of low noise to signal levels. An alternative approach to overcome this problem is proposed in the next section.

3.2 The Needlework estimator

Consider the “cross-bins” tick returns:

$$r_{j,\tau+1}^{(LF)} = p_{j,\tau+1}^{(F)} - p_{j,\tau}^{(L)} \quad (10)$$

and the standard last-tick interpolation returns $r_{j,\tau}^{(L)}$ to compute:

$$\text{Needlework}_t = \sum_{\tau=1}^M r_{i,\tau}^{(L)} r_{j,\tau}^{(L)} + r_{i,\tau}^{(L)} r_{j,\tau+1}^{(LF)} + r_{j,\tau}^{(L)} r_{i,\tau+1}^{(LF)} \quad (11)$$

We call this realized covariance estimator the Needlework tick-by-tick covariance estimator. Exactly the same idea can be applied to the first-tick interpolation.

The intuition behind this construction is (see Figure 3):

Insert Figure 3 about here

The lost portion of covariance (given by the time difference of the last ticks in the two series) is considered in one of the two cross-products with the cross-bin returns. The other cross-product will be ineffective, but since we do not know the time order between the last two returns of the two series, both cross-bin returns must be included to ensure coverage of the lost portion of covariance induced by the interpolation. This method may be seen as a general way of correcting standard interpolation schemes and could then be used whenever one wants to compute standard realized covariance at any given time interval. In fact, although the time stamp rounding induces a natural calendar time grid (for instance at the one minute interval), any lower frequency can be considered and the estimator computed with the sub-sampling method previously described. As already mentioned, this ensures consistency and extends the applicability of the estimator.

Using a sub-sample of the total number of ticks employed by the Hayashi and Yoshida estimator, we can expect the Needlework and the First-Last tick-by-tick estimators to be less efficient in the absence of microstructure noise or with low level of noise to signal ratio. Nevertheless, these forms of sub-sampling of the Hayashi and Yoshida estimator, could help in the presence of a significant level of market microstructure noise (as suggested by Palandri, 2006, and Voev and Lunde, 2007). Moreover, as the simulation results summarized in the next section show, it could also help to correct for the significant bias towards zero empirically found in the Hayashi and Yoshida estimator for highly liquid assets (see Griffin and Oomen, 2011) produced by the empirical presence of lead-lag cross-covariance (assumed to be zero in Hayashi and Yoshida, 2005).

The efficiency of the proposed estimators in the presence of microstructure noise will depend on the characteristics of the data. Thus, their asymptotic and finite sample properties will be hard to compute analytically. In order to assess the efficiency of the modified tick-by-tick estimators on empirical data, we perform a simulation study in which the data generating

process (DGP) mimics (as closely as possible) the econometric properties of the two empirical series we investigate in our real data application. The parameters are chosen to match (as closely as possible) the empirical observation frequencies, level of volatilities, noise structures, and intensities.

4 Monte Carlo simulations

In this section we evaluate the performance of different covariance estimators in three different simulation environments. In the first, we consider i.i.d. microstructure noise (standard setting) over a wide range of values of the noise to signal level, which is the key parameter affecting simulation results (as also reported in Griffin and Oomen, 2011); in the second environment we closely reproduce the complex dependence structure of the microstructure noise observed in the empirical data; in the third one, we generalize the simulation conditions to consider the empirically relevant effect of significant lead-lag cross-covariance.

4.1 Standard setting

The data generating process we consider here is a Lo and MacKinlay (1990) non-synchronous trading model with a heteroskedastic factor and microstructure noise, calibrated on our empirical data set. Our data is more than 18 years' tick-by-tick bivariate series of S&P 500 and 30-year US treasury bond futures with time stamps rounded to one minute.

The Lo and MacKinlay model defines the true return of an asset as given by a single factor model. Hence, considering two assets, the virtual continuously compounded return (here we consider return intervals of one second) $r_{i,t}$ is given by:

$$r_{i,t} = \mu_i + \beta_i f_t + \epsilon_{i,t} \quad i = 1, 2 \quad (12)$$

where β_i is the factor loading of asset i , $\epsilon_{i,t}$ represents the idiosyncratic noise of asset i , and f_t is the zero mean common factor.

Assuming that the idiosyncratic noises $\epsilon_{1,t}$ and $\epsilon_{2,t}$ are mutually uncorrelated, and that both are uncorrelated with the common factor f_t , the true covariance between the two assets

is:

$$\sigma_{1,2,t} = \beta_1 \beta_2 \sigma_{f,t}^2 \quad (13)$$

where $\sigma_{f,t}^2$ is the variance of the common factor f_t .

In the Lo and MacKinlay model, the common factor f is assumed to be a simple homoskedastic process; hence the variance of f is a constant σ_f^2 . Consequently, the true covariance also remains constant. In the version adopted here, however, in order to increase the realism of the DGP and to allow the true covariance to be time-varying, the dynamics of the common factor f are assumed to follow a discretized version (at the Euler clock of one second) of the stochastic volatility model proposed by Heston (1993)

$$df_t = \left(\mu - \frac{v_t}{2} \right) dt + \sigma_{f,t} dB_t \quad (14)$$

$$dv_t = k(\alpha - v_t)dt + \gamma v_t^{1/2} dW_t \quad (15)$$

where $v_t = \sigma_{f,t}^2$ and the initial value v_0 is drawn from the unconditional gamma distribution of v .

In the Lo and MacKinlay model the prices are assumed to be observed with a certain probability $1 - \pi_i$, where π_i is the so-called non-trading probability. We found it more convenient to express the frequency of the price observations in terms of the corresponding average intertrade duration between ticks τ_i .³

The values of the average intertrade durations and volatilities are chosen to match the statistical properties observed in the empirical data. Therefore, with asset 1 mimicking the S&P and asset 2 the US bond, the following configuration of the parameters is chosen: $\tau_1 = 8$ seconds, $\tau_2 = 18$ seconds, an average annualized volatility of about 20 percent for asset 1 and 10 percent for asset 2, and a correlation of 30 percent between the two assets. Time stamps of the observed prices are rounded at the one-minute level.

Each time a price is observed we simulate market microstructure effects by adding a stationary noise component independent of the price process. In this section we consider standard i.i.d. market microstructure noise whose variance is varied so to obtain a wide range of noise

³For example, a non-trading probability of 90 percent corresponds to an exponential distribution of the intertrade duration with a mean value of 10 seconds. In our data the average intertrade duration is eight seconds for the S&P return series and 18 seconds for the US bond return series.

to signal ratios for the observed returns process⁴ (which, in this analysis, are set to be equal for both assets). We then investigate the performances of realized covariance estimators for noise to signal values ranging from zero (no microstructure noise) to three (extremely high level of microstructure noise).

In our simulation study we compare the following realized covariance measures:

- The proposed First–Last estimator and its sub–sampled version. For each noise–to–signal ratio we select, through a grid search over integer numbers, the degree of sub–sampling K which minimize the root mean square error of the estimator.
- The proposed Needlework estimator and its sub–sampled version with optimized sub–sampling K .
- The Hayashi and Yoshida estimator described in Section 2. Being unfeasible on rounded time stamps, in our simulations setting this tick–by–tick estimator is computed before performing the time stamp rounding, that is, differently from all the other estimators, it is always computed on the exact time stamps.
- The standard realized covariance computed with an interpolated regular grid of one–minute returns.
- The standard realized covariance computed with a fixed return time interval of five minutes.
- The Scholes and William (1977) covariance estimator, which adds to the contemporaneous sample covariance of fixed interval returns, one lead–lag cross–covariance. To improve the performance of this estimator we choose the frequency of the fixed interval returns that provides the best results in terms of the root mean squared errors (RMSE). In our simulation set–up, such an optimal frequency is approximately one minute.
- The estimator proposed by Cohen et al. (1983), which is a simple generalization of the Scholes and Williams estimator where more than one lead and lag are considered. As in Bollerslev and Zhang (2003), we compute the Cohen et al. estimators with 12 leads

⁴As defined in Oomen (2006).

and lags at the 10-second frequency, which was the better performing choice given our simulation set-up.

- The Lo and MacKinlay estimator, given by:

$$\hat{\sigma}_{1,2} = \frac{1 - \hat{\pi}_1 \hat{\pi}_2}{(1 - \hat{\pi}_1)(1 - \hat{\pi}_2)} \text{Cov} [r_{1,t}^s, r_{2,t}^s] \quad (16)$$

where $\text{Cov} [r_{1,t}^s, r_{2,t}^s]$ is the covariance between the observed one-second returns $r_{i,t}^s$. Contrary to the highly noisy non-trading probability estimation proposed by Lo and MacKinlay, where $\hat{\pi}_1 = \text{Cov} [r_{1,t}^s, r_{2,t+1}^s] / \text{Cov} [r_{1,t}^s, r_{2,t}^s]$ ($\hat{\pi}_2$ is defined in an analogous way), we estimate those probabilities by counting the observed number of ticks in each day and dividing that by the total number of seconds in the day.

- The Two-Scale covariance estimator recently proposed by Zhang (2011) which is designed to simultaneously correct for the Epps effect and the microstructure noise. We follow Zhang (2011) recommendation of selecting a sub-sampling frequency for the slower time scale “large enough”, by choosing it equal to 10.

Insert Table 2 about here

Table 2 reports the performance results, measured in terms of root mean square error (RMSE), of the different estimators. The simulated data are generated with standard i.i.d. microstructure noise and noise-to-signal ratios of 0, 0.5, 1, ..., 3. For noise-to-signal ratios lower than one, the best estimator is the unfeasible Hayashi and Yoshida, followed by the First-Last and Needlework estimators with no sub-sampling (i.e. $K = 1$). We then have the Scholes and William and Cohen et al. lead lag estimators, the one minute and five minute standard realized covariance estimators and finally, with a much larger RMSE, the Two-Scale covariance estimator and the Lo and MacKinlay one.

For noise-to-signal ratio between one and two, however, the proposed First-Last and Needlework estimators show RMSE smaller than that of the Hayashi and Yoshida estimator which starts to show an increasing bias. The rank of the other estimators remain the same beside that the five minute covariance now outperforms the one minute estimator. The optimal degree of sub-sampling K remains equal to 1 (i.e. no sub-sampling) up to a level of

noise-to-signal equal to one for the Needlework and equal to two for the First-Last. It thus seems that the Needlework estimator requires a higher degree of sub-sampling than the First-Last, probably because it employs a larger number of cross-returns (each cross-return being contaminated by microstructure noise). However, the final performance with the optimal degree of sub-sampling is similar between the two estimators, although the First-Last seems to perform generally slightly better.

For large level of microstructure noise, i.e. noise-to-signal level larger than two, the Hayashi and Yoshida starts to be very biased and imprecise, and again the best estimators are the sub-sampled version of the First-Last and Needlework, although with different degree of optimal sub-sampling between the two (the Needlework being about double that of the First-Last). It is interesting to note that, contrary to the other estimators, the RMSE of the Two-Scale covariance tends to increase very little with the increase of noise-to-signal ratio, becoming competitive for very large noise-to-signal levels.

Summarizing the results of this first simulation analysis, we found that the proposed estimators outperform all the other feasible estimators for any level of noise-to-signal ratio, and they even outperform the unfeasible Hayashi and Yoshida for noise-to-signal levels larger than one. Moreover, we found that sub-sampling becomes increasingly useful for noise-to-signal ratio of 1.5 and larger, and, as a rule of thumb, the First-Last has an optimal degree of sub-sampling about equal to the noise-to-signal ratio while the one of the Needlework is about two times the noise-to-signal ratio.

4.2 Empirically realistic microstructure noise

From the empirical study of the tick-by-tick series of those assets, we found significant departure from the standard independent and identically distributed random variables (i.i.d.) assumption on the structure of the market microstructure noise. Through study of the autocorrelation of the tick returns of those series, more complex structures than those of a simple MA(1) expected under the standard i.i.d. assumption were found.

Insert Figure 4 about here

We suggest that such autocorrelation patterns of the tick returns could be explained by assuming a more complex ARMA structure for the microstructure noise. The noise structure is closely reproduced (see Figure 4) by introducing an MA(2) for the asset 1, mimicking the S&P 500 with $\theta_1 = 0.85$, $\theta_2 = 0.25$ and a noise-to-signal ratio of 0.45, and a strong oscillatory AR(1) with $\phi_1 = -0.65$ and noise-to-signal ratio of 0.6 for the asset 2 corresponding to the US bond.

Figure 5 and Table 3 report the results of the 25,000 simulations.

Insert Table 3 about here

Insert Figure 5 about here

With these observation frequencies, the one-minute realized covariance is slightly biased. By contrast, the five-minute realized covariance is unbiased, but has a larger variance. Same performances are observed for the Two-Scale covariance estimator. Despite the direct estimation of the non-trading probabilities, the Lo and MacKinlay estimator (though unbiased) is extremely inaccurate, also with this type of market microstructure noise. With chosen frequencies, both the Scholes and Williams and the Cohen et al. estimators are almost unbiased and reasonably accurate. However, the better ones are the tick-by-tick covariance estimators having no bias and the smallest dispersion among the estimators considered. With this moderate level of microstructure noise, the First-Last and Needlework estimators turn out to be less precise than the Hayashi and Yoshida one, which is unfeasible in the presence of rounded time stamps. This loss of efficiency is a direct consequence of the lower number of ticks employed in constructing the estimator. Nevertheless, differences from the Hayashi and Yoshida estimator are small. The First-Last and the Needlework estimators remain superior to all other feasible classical covariance estimators.

4.3 Lead-lag cross-covariances

An important empirical feature recently found by Griffin and Oomen (2011) is that financial high frequency data (especially more liquid data) show the presence of positive and significant lead and lag cross-covariance. As suggested by these authors, cross-dependence between non-overlapping returns can be due to non-instantaneous price adjustment, meaning that more

trades are necessary before prices fully incorporate the new information available. In such cases, the Hayashi and Yoshida estimator becomes biased because it neglects portions of cross-dependence that extend beyond the overlapping interval (dependence is assumed to be zero in the Hayashi and Yoshida derivation of their covariance estimator). The non-instantaneous price adjustment interpretation suggests that lead and lag cross-covariance might be better modeled by introducing delay adjustment in the true underlying price process than by introducing cross-dependence in the microstructure noise. (The latter will also suffer from the problem of linking a cross-dependence that arguably exists in physical time with a microstructure noise but is instead observed under trading time).

In order to reproduce it in our data-generating process, we generate lead and lag cross-dependence between the true return processes by introducing strong persistence in the dynamics of the common factor. To this end, the discretized version of the Heston model for f_t is replaced by a simple discretized Ornstein-Uhlenbeck process with slow mean reversion. In this simple way, our DGP can now produce a lead lag cross-covariance structure that closely mimics those structures empirically found in Griffin and Oomen (2011) (see Figure 6).

Insert Figure 6 about here

Table 4 and Figure 7 report the estimation results of the different covariance estimators.

Insert Table 4 about here

Insert Figure 7 about here

In the presence of lead lag cross-covariance, the Lo and MacKinlay estimator, the standard one-minute estimator, and the Hayashi and Yoshida estimator are all severely biased, while the five-minute, Scholes and Williams, Cohen et al., and the Two-Scale remain unbiased but with large variances. The Hayashi and Yoshida estimator has the smallest dispersion, but the presence of such significant bias substantially increases its RMSE. By contrast, the modified tick-by-tick estimators proposed remain both virtually unbiased and achieve the lowest RMSE values among all the competing estimators (also including the Hayashi and Yoshida one). The intuition for this result is that both the First-Last and the Needlework are form of sub-sampling (to the rounding frequency) of the Hayashi and Yoshida, thus reducing the sensitivity of the estimators to the presence of high-frequency lead-lag cross-covariances.

4.4 Summarizing the simulation results

The proposed First–Last tick-by-tick estimator is the best performing estimator (closely followed by the Needlework estimator) among the feasible ones, since the Hayashi and Yoshida estimator is unfeasible on tick-by-tick data with rounded time stamps. It also performs favorably compared to the Scholes and Williams, Zhang Two-Scale covariance, and the Cohen et al. estimators, even if their return frequency was chosen according to the simulation settings to give the best results. When lead-lag cross-covariance is introduced, both proposed estimators remain unbiased and clearly outperform even the Hayashi and Yoshida estimator, which under these conditions becomes severely biased.

5 Empirical application

5.1 Data

We now apply the proposed First–Last covariance estimator to the tick-by-tick bivariate series of S&P 500 futures and 30-year US treasury bond futures. The data are from the Price-data.com database. Time stamps are rounded at the one-minute level (see Table 1). The period considered extends from January 1990 to May 2008. The time series of the realized daily stock–bond covariances is plotted in Figure 8.

Insert Figure 8 about here

First, to appreciate the remarkable difference between the daily realized covariances measured by using tick-by-tick data and the standard cross-products of daily returns (the usual, inaccurate proxy for daily covariances in standard multivariate volatility models), both measures are plotted together on the same scale. These differences are of central importance when assessing the fitness of a model for volatilities or correlations (see, for example, Patton, 2011). Using inaccurate proxies for daily covariance may lead to the choice of a less accurate model for forecasting second-order dynamics.

5.2 Assessing performance with volatility-timing strategies

We also evaluate the economic benefit of the proposed covariance estimators, employing the methodology suggested by West et al. (1993) and Fleming et al. (2001, 2003), which compares the utility level obtained by different variance-covariance estimators in the context of portfolio allocation.⁵ In particular, we consider a risk-averse investor who uses a conditional mean-variance optimization rule to allocate funds across stocks (S&P 500) and bonds (30-year US Treasury bond), re-balancing his portfolio daily. To avoid short selling restrictions and minimize transaction costs, the allocation decisions are implemented by trading on futures contracts.

Defining R_{t+1}^f and R_{t+1} as the risk-free rate and the bivariate vector of daily return from t , $t + 1$, respectively, we denote $\mu_t = E_t[R_{t+1}]$ the vector of conditional means and $\Sigma_t = E_t[(R_{t+1} - \mu_{t+1})(R_{t+1} - \mu_{t+1})']$ the conditional covariance matrix of R_{t+1} . The investor then solves the following quadratic program:

$$\min_{w_t} \quad w_t' \Sigma_t w_t$$

subject to

$$w_t' \mu_t + (1 - w_t' \mathbf{1}_2) R_t^f = \mu_p$$

where w_t is the vector of portfolio weights, $\mathbf{1}_2$ is the 2×1 unit vector, and μ_p is a target expected return of the portfolio. The solution to this classical mean-variance optimization problem is:

$$w_t = \frac{(\mu_p - R_t^f) \Sigma_t^{-1} (\mu_t - R_t^f \mathbf{1}_2)}{(\mu_t - R_t^f \mathbf{1}_2) \Sigma_t^{-1} (\mu_t - R_t^f \mathbf{1}_2)}$$

hence, the portfolio weights w_t identify the daily re-balanced portfolio that minimizes conditional variance for any choice of expected return μ_p . The optimal portfolio weights vary through time as both μ_t and Σ_t change. However, since there is no evidence of predictability in the expected returns at the daily level, we assume that the investor sets μ_p equal to R_{t+1}^f

⁵As an alternative, one may think to assess the goodness of the different realized covariance estimators in other empirical applications like, for example, risk management (see, among others, Fleming and Kirby, 2003.). This is left for future research.

plus a constant spread (10% for the S&P 500 and 5% for the US treasury Bond). In a mean-variance framework, this is therefore equivalent to following a volatility-timing strategy where the weights vary only with Σ_t .

The variance–covariance matrix Σ_t is estimated daily using the high frequency information at time $t - 1$. To quantify the economic value of volatility timing using alternative realized covariance measures, we compute the daily variance–covariance matrix with the following covariance estimators: First–Last, Needlework, standard one–minute, standard five–minute, Scholes and Williams, and Cohen 12 lead–lags. To focus our attention on the impact on portfolio allocation of the different covariance estimates, we use the same variance measure for all the different variance–covariance matrices: the Two-Scale realized volatility estimator of Zhang et al. (2005).

Following Engle and Colacito (2006), Bandi et al. (2008), and Audrino and Trojani (2011), we quantify the economic differences between the alternative covariance estimates using the variance component of an investor’s long-run mean–variance utility:

$$AU = \frac{\lambda}{2} \frac{1}{N} \sum_{t=1}^N (R_t^p - \bar{R}^p)^2 \quad (17)$$

where

$$R_t^p = R_t^f + w'_{t-1}(R_t - R_t^f \mathbf{1}_2)$$

is the return of the portfolio constructed at time $t - 1$, $\bar{R}^p$ is the sample mean of portfolio returns, and λ is a coefficient of risk–aversion. The difference between two AU measures obtained with different covariance estimators can be interpreted as the fee that the investor would be willing to pay to switch from one covariance measure to the other.

5.3 Empirical results

We compute the AU quantity of equation (17) for the different high frequency covariance estimators over the full sample of more than 18 years from January 1990 to May 2008. We choose the US three-month rate as proxy for the risk-free rate R_t^f . Table 5 reports the annualized fees (expressed in basis points) that an investor following a volatility timing strategy would be willing to pay to employ the First–Last covariance measure in place of the other realized

covariance estimators. Results are reported for three conventional values of the risk-aversion parameter $\lambda = 2, 7$, and 10 and three targets of the portfolio expected return $\mu_p = 6\%, 10\%$ and 15% .

Insert Table 5 about here

The figures reported in Table 5 shows that an investor following a volatility timing strategy would always be willing to pay a positive amount of money (the annual fees) in order to switch to the First–Last realized covariance measure (although the economic difference between First–Last and Needlework is negligible). The economic gains in employing the proposed tick-by-tick covariance estimators are often remarkably large, especially for greater values of μ_p and λ .

6 Conclusions

We adapted the approach of computing realized covariance with tick-by-tick data to the case where price time stamps are rounded to a frequency lower than the typical arrival frequency of the asset encountered in the real world. The proposed methods can be seen as general interpolation schemes allowing to compute unbiased realized covariance at any desired time interval. In fact, although the time stamp rounding induces a natural calendar time grid, the proposed estimators can be computed at any lower time frequency by employing the sub-sampling and averaging scheme which ensures consistency and robustness in presence of market microstructure noise.

Monte Carlo simulations performed with a wide range of noise-to-signal ratios, realistic dependence structures of the microstructure noise and true lead-lag cross covariance, show that the proposed First–Last and Needlework tick-by-tick covariance estimators turn out to be systematically the best performing among the realized covariance estimators that can be feasibly computed in presence of rounded time stamps. Moreover, thanks to higher robustness to large microstructure noise (with the sub-sampled versions), and significant lead-lag cross-covariance they even outperform (in terms of RMSE) the unfeasible Hayashi and Yoshida tick-by-tick estimator which, under these conditions, shows a substantial bias.

An empirical application considering a risk-averse investor who uses a conditional mean-variance optimization rule to allocate funds across stocks (S&P 500) and bonds (30 year US

Treasury bond) shows that such an investor would always be willing to pay a positive and often large amount of money to switch from a standard covariance measure to the proposed tick-by-tick ones, confirming the simulation results.

References

- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P.: 2001, The distribution of realized exchange rate volatility, *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P.: 2003, Modeling and forecasting realized volatility, *Econometrica* **71**, 579–625.
- Audrino, F. and Trojani, F.: 2011, A general multivariate threshold GARCH model with dynamic conditional correlations, *Journal of Business and Economic Statistics* **29**(1), 138–149.
- Bandi, F., Russell, J. and Zhu, Y.: 2008, Using high-frequency data in dynamic portfolio choice, *Econometric Reviews* **27**(1), 163–198.
- Barndorff-Nielsen, O. E. and Shephard, N.: 2001, Non-gaussian ornstein-uhlenbeck-based models and some of their uses in financial economics, *Journal of the Royal Statistical Society, Series B* **63**, 167–241.
- Barndorff-Nielsen, O. E. and Shephard, N.: 2002a, Econometric analysis of realized volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society, Series B* **64**, 253–280.
- Barndorff-Nielsen, O. E. and Shephard, N.: 2002b, Estimating quadratic variation using realized variance, *Journal of Applied Econometrics* **17**(5), 457–477.
- Barndorff-Nielsen, O. E. and Shephard, N.: 2005, How accurate is the asymptotic approximation to the distribution of realized volatility?, in D. W. F. Andrews and J. H. Stock (eds), *Identification and Inference for Econometric Models. A Festschrift in Honour of T.J. Rothenberg*, Cambridge University Press, pp. 306–331.
- Bollerslev, T. and Zhang, B. Y. B.: 2003, Measuring and modeling systematic risk in factor pricing models using high-frequency data, *Journal of Empirical Finance* **10**(5), 533–558.

- Cohen, K., Hawawini, G. A., Maier, S. F., R., R. S. and D., D. W.: 1983, Friction in the trading process and the estimation of systematic risk, *Journal of Financial Economics* **12**, 263–278.
- Comte, F. and Renault, E.: 2001, Long memory in continuous time stochastic volatility models, *Mathematical Finance* **8**, 291–323.
- Corsi, F.: 2005, *Measuring and modelling realized volatility: from tick-by-tick to long memory*, PhD thesis, University of Lugano.
- Dacorogna, M. and Lundin, M.: 1999, Correlation of high-frequency financial time series, in P. Lequeux (ed.), *Financial Markets Tick by Tick*, Wiley & Sons.
- Engle, R. and Colacito, R.: 2006, Testing and valuing dynamic correlations for asset allocation, *Journal of Business and Economic Statistics* **24**(2), 238–253.
- Epps, T.: 1979, Comovements in stock prices in the very short run, *Journal of the American Statistical Association* **74**(366), 291–298.
- Fleming, J. and Kirby, C.: 2003, A closer look at the relation between GARCH and stochastic autoregressive volatility, *Journal of Financial Econometrics* **1**(3), 365–419.
- Fleming, J., Kirby, C. and Ostdiek, B.: 2001, The economic value of volatility timing, *Journal of Finance* **56**(1), 329–352.
- Fleming, J., Kirby, C. and Ostdiek, B.: 2003, The economic value of volatility timing using realized volatility, *Journal of Financial Economics* **67**(3), 473–509.
- Frost, P. A. and Savarino, J. E.: 1986, An empirical bayes approach to portfolio selection, *Journal of Financial and Quantitative Analysis* **21**, 293–305.
- Griffin, J. and Oomen, R.: 2011, Covariance measurement in the presence of non-synchronous trading and market microstructure noise, *Journal of Econometrics* **160**(1), 58–68.
- Hautsch, N., Kyj, L. and Oomen, R.: 2010, A blocking and regularization approach to high-dimensional realized covariance estimation, *Journal of Applied Econometrics* .

- Hayashi, T. and Yoshida, N.: 2005, On covariance estimation of non-synchronously observed diffusion processes, *Bernoulli* **11**(2), 359–379.
- Heston, S.: 1993, A closed-form solution for options with stochastic volatility with applications to bonds and currency options, *Review of Financial Studies* **6**, 327–343.
- Kalnina, I.: 2010, Subsampling high frequency data, *Journal of Econometrics* **161**(2), 262 – 283.
- Ledoit, O. and Wolf, M.: 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, *Journal of Empirical Finance* **10**(5), 603–621.
- Ledoit, O. and Wolf, M.: 2004a, Honey, i shrunk the sample covariance matrix, *Journal of Portfolio Management* **31**(1).
- Ledoit, O. and Wolf, M.: 2004b, A well-conditioned estimator for large-dimensional covariance matrices, *Journal of Multivariate Analysis* **88**(2), 365–411.
- Lo, A. and Andrew, W.: 1990, An econometric analysis of nonsynchronous trading, *Journal of Econometrics* **45**(1-2), 181–211.
- Martens, M.: 2004, Estimating unbiased and precise realized covariances. Social Science Research Network Electronic Library.
- Oomen, R.: 2006, Properties of realized variance under alternative sampling schemes, *Journal of Business and Economic Statistics* **24**(2), 219–237.
- Palandri, A.: 2006, Consistent realized covariance for asynchronous observations contaminated by market microstructure noise. Unpublished Manuscript.
- Patton, A.: 2011, Volatility forecast comparison using imperfect volatility proxies, *Journal of Econometrics* **160**(1), 246–256.
- Renó, R.: 2003, A closer look at the epps effect, *International Journal of Theoretical and Applied Finance* **6**(1), 87–102.

- Scholes, M. and Williams, J.: 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* **5**, 181–212.
- Voev, V. and Lunde, A.: 2007, Integrated covariance estimation using high-frequency data in the presence of noise, *Journal of Financial Econometrics* **5**, 68–104.
- West, K., Edison, H. and Cho, D.: 1993, A utility-based comparison of some models of exchange rate volatility, *Journal of International Economics* **35**, 23–45.
- Zhang, L.: 2011, Estimating covariation: Epps effect, microstructure noise, *Journal of Econometrics* **160**(1), 33–47.
- Zhang, L., Aït-Sahalia, Y. and Mykland, P. A.: 2005, A tale of two time scales: determining integrated volatility with noisy high-frequency data, *Journal of the American Statistical Association* **100**, 1394–1411.

S&P 500 data sample

Date	Time	Name	Price
20030724	1322	SP_03U	994.3
20030724	1322	SP_03U,	994.5
20030724	1322	SP_03U	994.7
20030724	1322	SP_03U	994.8
20030724	1322	SP_03U	994.9
20030724	1322	SP_03U	995.0
20030724	1322	SP_03U	995.2
20030724	1322	SP_03U	995.1
20030724	1322	SP_03U	995.0
20030724	1323	SP_03U	994.8
20030724	1323	SP_03U	994.7
20030724	1323	SP_03U	994.6
20030724	1323	SP_03U	994.5
20030724	1324	SP_03U	994.3
20030724	1324	SP_03U	994.2
20030724	1324	SP_03U	994.0

Table 1: Small data sample example for the S&P 500 index future on 24 July 2003 (20030724). The data are from the Price-data.com database. Time is written in the format hours minutes (hhmm).

Different noise-to-signal ratios simulation results: standard setting

noise-to-signal ratio	0	0.5	1	1.5	2	2.5	3
Lo-MacKinlay	0.465	0.884	2.832	6.176	10,92	16.927	24.350
1 min no correction	0.312	0.326	0.347	0.396	0.483	0.599	0.758
5 min no correction	0.327	0.336	0.347	0.367	0.418	0.470	0.527
Scholes and Williams Cov	0.244	0.246	0.273	0.312	0.378	0.472	0.595
10 sec Cohen 12 leads-lags	0.256	0.263	0.293	0.345	0.450	0.592	0.779
Zhang Two-Scales Cov	0.392	0.399	0.400	0.396	0.408	0.407	0.410
First-Last	0.190	0.193	0.210	0.234	0.282	0.3490	0.433
First-Last Optimal Sub. (optimal K)	0.190 (1)	0.193 (1)	0.210 (1)	0.234 (1)	0.298 (2)	0.311 (2)	0.351 (3)
Needlework	0.187	0.196	0.225	0.275	0.343	0.454	0.571
Needlework Optimal Sub. (optimal K)	0.187 (1)	0.196 (1)	0.225 (1)	0.271 (3)	0.312 (4)	0.340 (5)	0.374(6)
Hayashi and Yoshida on exact time stamp	0.126	0.141	0.210	0.348	0.554	0.837	1,175

Table 2: Root mean squared errors (RMSE) of the estimation errors on the annualized covariance for a simulation set-up having i.i.d. microstructure noise and noise-to-signal ratios of $0, 0.5, 1, \dots, 3$. The First-Last and Needlework estimators are computed with no sub-sampling and with optimal sub-sampling K (reported in parenthesis). Being unfeasible on rounded time stamps data, as a comparison the Hayashi and Yoshida estimator is computed (unlike all the other estimators) on the exact time stamps.

Calibrated S&P US bond simulation			
	bias	std	RMSE
Lo and MacKinlay	-0.0009	0.8090	0.8090
1-min no correction	-0.1764	0.2732	0.3252
5-min no correction	-0.0374	0.3392	0.3413
Scholes and Williams	-0.0073	0.2592	0.2593
Cohen et al. 12 leads-lags	-0.0057	0.2719	0.2719
Zhang Two-Scales Cov	-0.0353	0.4035	0.4050
First–Last	-0.0032	0.1973	0.1973
Needlework	-0.0042	0.2060	0.2060
Hayashi and Yoshida	0.0002	0.1602	0.1602

Table 3: Mean, standard deviation and root mean squared errors (RMSE) of the estimation errors on the annualized covariance for a simulation set-up that reproduces the statistical properties of the S&P 500 and US bond future data. As a comparison, the Hayashi and Yoshida estimator is computed (unlike all the other estimators) on the exact time stamps.

Calibrated S&P US bond simulation with lead-lag cross-covariances

	bias	std	RMSE
1-min no correction	-0.4698	0.1710	0.5000
5-min no correction	-0.0802	0.4405	0.4478
Scholes and Williams	0.0111	0.2945	0.2947
Cohen et al. 12 leads-lags	0.0110	0.3158	0.3159
Zhang Two-Scales Cov	0.0806	0.3533	0.3624
First-Last	0.0018	0.1938	0.1938
Needlework	-0.0081	0.2046	0.2047
Lo and MacKinlay	-1.3238	0.7719	1.5324
Hayashi and Yoshida	-0.2077	0.1513	0.2570

Table 4: Mean, standard deviation and root mean squared errors (RMSE) of the estimation errors on the annualized covariance for a simulation set up that reproduces the statistical properties of the S&P 500 and US bond future data, including lead-lag cross-dependence. As a comparison, the Hayashi and Yoshida estimator is computed (unlike all the other estimators) on the exact time stamps.

Volatility timing comparison: Annualized fees (in basis points)				
Target	Realized Covariance Measure	$\lambda = 2$	$\lambda = 7$	$\lambda = 10$
6%	Needlework	0.109	0.381	0.545
	One–Minute	3.382	11.838	16.912
	Five–Minute	5.770	20.195	28.850
	Scholes and Williams	3.955	13.842	19.774
	Cohen et al. 12 leads-lags	9.904	34.662	49.518
10%	Needlework	0.045	0.159	0.227
	One–Minute	10.535	36.873	52.676
	Five–Minute	19.584	68.545	97.922
	Scholes and Williams	13.967	48.885	69.836
	Cohen et al. 12 leads-lags	39.283	137.489	196.413
15%	Needlework	0.253	0.885	1.265
	One–Minute	25.452	89.081	127.259
	Five–Minute	49.344	172.703	246.718
	Scholes and Williams	37.345	130.707	186.724
	Cohen et al. 12 leads-lags	113.944	398.804	569.720

Table 5: Annualized fees (expressed in basis points) that an investor following a volatility timing strategy would be willing to pay to employ the First–Last covariance measure in place of the alternative realized covariance estimators. The portfolio weights are obtained minimizing the conditional variance of a portfolio containing the S&P 500 and 30–year US Treasury bond over the full sample from January 1990 to May 2008 (4,480 daily observations).

Graphical illustration of the Epps effect




Figure 1: Graphical representation of the Epps effect arising from standard covariance estimators when returns are interpolated on a regular grid (vertical lines) using previous tick interpolation, that is considering the last price tick of the interval (represented with a small square) as the price prevailing at the end of the grid. The shaded area is the portion of covariance lost because of the wrong imputation of a portion of the second asset return to the next interval.

Graphical illustration of the First–Last covariance estimator




Figure 2: Graphical representation of the first component $\sum_{s=1}^{M_{i,t}^{(F)}} \sum_{q=1}^{M_{j,t}^{(L)}} r_{i,s}^{(F)} r_{j,q}^{(L)} I(\delta_{q,s} > 0)$ of the First–Last covariance estimator, where the first asset i is interpolated with next-tick interpolation (small circles represent the first ticks) and asset j with previous tick interpolation (with last ticks being the small squares).

Graphical illustration of the needlework covariance estimator




Figure 3: Graphical illustration of the needlework covariance estimator with the two cross-bin returns $r_{i,\tau+1}^{(LF)}$ and $r_{j,\tau+1}^{(LF)}$ (small inner arch) multiplying the last-tick interpolated returns of the other asset in order to correct for the lost portion of covariance induced by previous tick interpolation.




Figure 4: Tick-by-tick autocorrelations of S&P 500 (top) and US bond (bottom) empirical (left) and simulated (right) returns. In order to mimic the dynamics of the S&P 500, an MA(2) microstructure noise has been simulated with $\theta_1 = 0.85$, $\theta_2 = 0.25$ and noise to signal of 0.45. For the same reason, for the US bonds the microstructure noise is an AR(1) with $\phi_1 = -0.65$ and noise to signal of 0.6.

Covariance estimation errors for empirically calibrated simulations




Figure 5: Comparison of the probability density function of the covariance estimation errors for a simulation set-up that reproduces the statistical properties of the S&P 500 futures and US bond future data. Only the best (in terms of MSE) five estimators are plotted.

Simulated true lead-lag cross-covariance




Figure 6: Lead-lag cross-covariances of simulated data obtained from a Lo and MacKinlay (1990) non-synchronous trading model with market microstructure noise and common factor following a slow mean reverting Ornstein-Uhlenbeck process.

Covariance estimation errors with true lead-lag cross-covariance




Figure 7: Comparison of the probability density function of the covariance estimation errors for a simulation set-up that reproduces the statistical properties of the S&P 500 futures and US bond future data and the presence of true lead-lag cross-covariance between the two series. Only the best (in terms of MSE) five estimators are plotted.

S&P US bond covariance from 1990 to 2008




Figure 8: Time series of daily realized covariances constructed using tick-by-tick data (solid line) superimposed on the daily cross-product returns (dotted line) of S&P 500 futures 30-year US treasury bonds from 1990 to 2008.