

School of Finance



University of St.Gallen

PRECIOUS METALS UNDER THE MICROSCOPE: A HIGH-FREQUENCY ANALYSIS

**MASSIMILIANO CAPORIN
ANGELO RANALDO
GABRIEL G. VELO**

WORKING PAPERS ON FINANCE NO. 2014/09

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

JANUARY 2014



Precious metals under the microscope: A high-frequency analysis*

Massimiliano Caporin[†] Angelo Ranaldo[‡] Gabriel G. Velo[§]

January 21, 2014

Abstract: Taking advantage of a trades-and-quotes high-frequency database, we document the main stylized facts and dynamic properties of spot precious metals, i.e. gold, silver, palladium, and platinum. We analyze the behaviors of spot prices, returns, volume, and selected liquidity measures. We find clear evidence of periodic patterns matching the trading hours of the most active markets round-the-clock. The time series of spot returns have thus properties similar to those of traditional financial assets with fat tails, asymmetry, periodic behaviors in the conditional variances, and volatility clustering. The gold (platinum) is the most (least) liquid and less (most) volatile asset. Commonality in liquidities of precious metals is very strong.

Keywords: precious metals, high-frequency data, liquidity, commonality in liquidity, intradaily periodicity.

JEL Classifications: C58, C22, C52, G10.

*An early version of this paper circulated under the title: *Stylized facts and dynamic modeling of high-frequency data on precious metals*. We thank the Editor and two referees of this Journal for their insightful comments.

[†]Department of Economics and Management *Marco Fanno*, University of Padova; corresponding author: Via del Santo 22, Padova, Italy, Tel: +39 049/8274258. Fax: +39 049/8274211. Email: massimiliano.caporin@unipd.it.

[‡]University of St. Gallen, Swiss Institute of Banking and Finance, Rosenbergstrasse 52, St.Gallen, Switzerland, angelo.ranaldo@unisg.ch.

[§]Department of Economics and Management *Marco Fanno*, University of Padova, Via del Santo 22, Padova, Italy, gabriel.velo@unipd.it.

1 Introduction

In the last twenty years there has been a wider access to high-frequency data allowing researchers to document several stylized facts and to test economic and financial theories from new perspectives. While the high-frequency literature has widely been applied to stocks and bonds, see Dacorogna et al. (2001), Parlour and Seppi (2008), and Abergel et al. (2012), among others, none of the previous studies examines the spot precious metals. We provide the first high-frequency analysis on spot precious metals by describing the main stylized facts and dynamic properties of gold, silver, palladium, and platinum.

Only a few authors have discussed the properties of high-frequency data on precious metals. Cai et al. (2001) focus on 5-minute gold future returns. They study the periodicity in absolute returns and link its movements to macroeconomic announcements. Fleming et al. (2003) make use of realized volatility on several assets and determine the economic value of volatility timing. Their study includes 5-minute returns of gold futures contracts. Baillie et al. (2007) analyze the dynamic behavior of several daily and 5-minute commodity futures prices, including the gold future. Bannouh et al. (2009) work on co-range computation (an estimator of the integrated covariance) with a trade dataset that includes the gold future. Finally, Khalifa et al. (2011) focus on volatility measurement and forecasting of several commodity future prices, including gold and silver. All of the aforementioned studies have a common element: They work on New York futures data. Moreover, to the best of our knowledge, gold has piqued the greatest amount of interest, whereas silver has only attracted limited attention. We support this by the larger amount of interest in gold, compared with other precious metals, and the perception of gold as a safe-haven investment.

Our work belongs to the strand of literature focusing on high-frequency data on precious metals. We differ from previous studies in at least five respects. First, we have access to high-frequency data about four different precious metals: gold, silver, palladium, and platinum. None of the previous papers has studied a comprehensive set of precious metals. Second, our data are related to spot prices and not futures contracts. Therefore, the time-series behaviors might be different as a consequence of the different uses associated with spot and future contracts. To our knowledge, there is no literature on palladium and platinum based on high-frequency data. Third, we have access to a brokerage house database that operates on a 24-hours basis. As a result, the movement we observe in high-frequency time series can be linked to the activity in different precious metals markets, including New York but also the European markets (London and Zurich) and the Asian markets. This is of relevant interest because previous works only focused on a specific market. Fourth, we differ from previous works because of our access to a trade-and-quotes dataset. From the trade side, we have trading prices and volumes as well as the trade direction (i.e. the exact identification if a trade was buyer- or seller-initiated). On the quote side, we have access to the limit order book up to ten levels (i.e. the best then ask and bid quotes). We observe the limit order price and volume, which makes feasible the measurement of more dimensions of market liquidity. To our knowledge, this is the first study on market liquidity and commonality in liquidities of precious metals. Finally, the time frequency of the database is extremely high, with book updates reaching a 100-millisecond frequency. Because of the novelty of our database, the pure statistical analysis of high-frequency spot data on precious metals is per se interesting and relevant. In this work, we focus on the stylized facts and dynamic properties of the precious metals data. We start from the most traditional time series, prices, and returns, and then

focus on the volatility, preliminary measured from squared returns and then filtered with periodic components and GARCH-type models. Notably, we introduce an EGARCH specification augmented with exogenous variables and HAR terms, e.g. Corsi (2009), to capture both periodicity and long-range dependence in conditional variances. We move later to volume time series and to two specific liquidity measures, the order flow, and the percentage quoted spread. These two quantities provide a first look at the information content one might extract from a trade-and-quotes database on precious metals. The analyses reported here are, by construction, preliminary to further economic applications at both the univariate and multivariate level. Later studies will take advantage of the findings in the current paper.

There are several areas in which our results might be useful: first, our study should help understand better the price formation process of precious metals which is fundamentally different from that of stocks and bonds. For instance, precious metals pay no dividends or coupon rates. In a fiat currency regime, the value of precious metals should essentially be driven by some genuine demand and supply mechanisms that are different from other assets. For instance, managing a large amount of gold reserves and pursuing other goals than profit-maximization, central banks are certainly particular players and should substantially affect the gold prices (and indirectly those of all precious metals).¹ A priori, it is not entirely clear what drives price fluctuations of spot precious metals and this study contributes to understand better the price discovery process of these commodities. Second, this paper should provide valuable insights to traders of precious metals. We propose models that might be used to explain and forecast prices, volumes, and liquidity. The model specifications we propose, given their dynamic nature, can be easily used to compute forecasts of the different quantities relevant for trading decisions. Third, this paper can be of interest for risk managers. Gold and other precious metals can be used as hedging instruments or stress indicators because of their safe-haven properties. This study sheds light on the volatility and liquidity behaviors of these refuge assets. Fourth, this study can be useful to perform dedicated event studies. For instance, gold is typically seen as a hedge for inflation risk. The dynamic behavior of prices, returns, volume and liquidity reported in this paper might be used to investigate the impact of central bank announcements on Quantitative Easing programs on investors inflation fear. Finally, we document common patterns in the time-series and cross-sectional variations of market liquidity of precious metals. Although very preliminary, our results indicate that precious metals exhibit strong commonality in liquidity. Since commonality in liquidity implies liquidity risk premiums, the question whether spot precious metals bear liquidity risk premiums can be addressed in future research.

The time-series analyzed here are characterized by the presence of a periodic behavior in their levels and/or in the second-order moment. Such a finding is well expected as we consider intradaily data. However, the periodic movement is associated with the trading activity of the main markets active in precious metals trading, which can be identified, for instance, with peaks in the volatility of the returns time series. This result has relevant affinities to the periodic behavior observed by e.g. Dacorogna et al. (1993) on the foreign exchange market. For each of the analyzed series, we suggest the use of a dynamic model for capturing the serial dependence in the mean and/or in the variance. Our methodological contribution is to extend the filtering techniques previously applied to

¹A recent example is represented by the substantial price impact due to the rumours on the intention of the Central Bank of Cyprus to sell an important part of its gold reserves to deal with its banking crisis, see Financial Times (2013).

intraday volatility to market liquidity measures of precious metals. We provide evidence supporting the idea that the proposed specifications seem appropriate in the whitening of the analyzed time series and allow the development of research based on time-series forecasting. However, the evaluation of forecasting performances of alternative models for precious metals returns, volatility, volume and liquidity is left to future researches.²

This paper proceeds as follows: Section 2 describes the dataset and the variables of interest; Section 3 focuses on the stylized facts of prices, returns, volumes, and liquidity time series; Section 4 deals with dynamic models; and Section 5 concludes.

2 Database description, data handling, and the analyzed quantities

The database used in this study has been provided by ICAP through its platform called Electronic Brokerage Services (EBS). EBS is the leading interdealer trading platform for currencies, and data provided by EBS have been already used in academic research (see, for instance, Berger et al., 2008 and Mancini et al., 2013). However, this is the first study on precious metals using EBS data. The database we access is equivalent to that adopted in Mancini et al. (2013) and it includes more than fifty currencies and four precious metals. For all assets made available, the trades and quotes refer to spot prices. In this work we focus on the precious metals: gold (identified by the ticker XAU), silver (XAG), palladium (XPD), and platinum (XPT). We consider spot prices against the U.S. dollar and refer to one ounce (for instance, gold quotes are U.S. dollar per one gold ounce). The data we analyze span the period beginning with December 27, 2008 to November 30, 2010. The dataset includes trading and order information. Trading data include transaction price, volume and the exact time stamp of the trade. The presence of the maker and taker sides allows us construct the order flow. Notably, as observed by Mancini et al. (2013), trade direction data are known and do not need to be inferred by means of rules such as done by Lee and Ready (1991). Order data include the binding bid and ask quotes, together with the related order volume. Every time there is a change in the order book, the database provides us a snap of the order book up to the tenth level (i.e. the ten bid quotes with the highest limit prices and the ten ask quotes with the lowest limit prices). However, here we present analyses based on the trade data and on the best bid and ask quotes. Moreover, we discuss the dynamic features of the data, comparing our findings to those of previous studies.

In the EBS database, data are recorded at a very high frequency (let us call them nanofrequency data.³) From December 27, 2008 to the end of August 2009, the observation frequency is 250 milliseconds; from the first of September 2009 to the end of the sample, the observation frequency increases to 100 milliseconds.⁴ As a result, new information is recorded by the system every 100/250 milliseconds if at least one of the following events takes place: an order is executed and/or a new quote is entered, deleted or revised. The last case produces a new flash of the book because each flash includes the number of quotes pending in the book for each side/price together with the available volume. Table

²The computation of forecasts based on the models presented in the following sections requires limited efforts. For the sake of focus, however, we refrain from extending our study to a forecasting exercise.

³We thank Michael McAleer for suggesting us the use of this name for our dataset.

⁴The 100-millisecond frequency begins August 28, 2009; however, to test the 100-millisecond frequency EBS organized already a few days of July 2009 at this frequency. Moreover, both frequencies (100 and 250 milliseconds) are present on July 21, 2009, and August 28, 2009.

Table 1: Number of quotes and trades

	XAU	XAG	XPd	XPT
Quotes bid	184.254.422	51.970.220	8.649.896	22.487.653
Quotes offer	169.512.628	48.992.734	7.908.232	23.365.209
Quotes total	353.767.050	100.962.954	16.558.128	45.852.862
Trades total	184.929	27.638	21.428	27.357
Outside trade Sunday	385	8	54	38
Outside trade Friday	32	7	17	40
Outside trade Saturday	2	0	1	0
Total outside trade	419	15	72	78
Traded volume (in 1000 oz)	233.610	1.173.425	165.580	170.140

(1) reports the total number of trades (buyer/seller initiated) and quotes (by market side) for the precious metals included in the database. We note that gold is the metal with the highest activity, both in terms of trades and quotes. Silver has less than one-third of the quotes of gold and is followed by platinum and then palladium. The last has less than a twentieth of gold quotes.⁵ Excluding gold, with about 185 thousand trades, the other precious metals have a similar number of trades. It is worth noting that an order ratio of more than 50 is pretty in line with other financial markets and it is not an indication of high-frequency trading. For instance, on July 2, 2012, the NASDAQ adopted an excess order fee for members with an individual order ratio (formally called Order Entry Ratio) of more than 100.

Data are recorded on a GMT time scale, on a 24-hour basis, 7 days a week. As a result, there are quotes and trades executed from buyers/sellers located in different geographical areas. To the best of our knowledge, this study provides for the first time an analysis on precious metal trades jointly covering the most active world markets. Compared with the previous studies of Barkoulas et al. (1997), Baillie et al. (2007), and Khalifa et al. (2011), our analysis includes trading activity originating from Asian markets as well as the complete activity of European-based traders. Previous studies have focused on U.S.-based data. Therefore, the daily time coverage of our database is a distinctive feature of our contribution.

In the following analyses, we limit the trading period to five full days, ranging from 10 p.m. Sunday to 10 p.m. Friday. We delete all trades and quotes outside this range. As shown in Table (1), the weekend time excluded from the database is really of minor relevance. As expected, in absolute terms, the deleted information is higher for gold compared with the remaining metals. The procedure we adopt is similar to that employed for the analysis of currency data (see Andersen et al., 2003 and Berger et al., 2008, among others).

The database reports the prices of precious metals expressed in U.S. Dollars per ounce, while the volume is expressed in multiples of the minimum tradable amount (MTA), which

⁵The number of quotes reported in Table (1) refers to the entire book which is available in our dataset. Therefore, the rough total number of best bids (best asks) is obtained by dividing the total number of bid (ask) by 10. If we take only the best quotes, the order entry ratio (number of best quotes) divided by number of trades varies between about 50 for Palladium to around 200 for Silver.

is fixed as follows: 1,000 ounces for Gold; 25,000 ounces for silver; and 500 ounces for palladium and platinum. As a consequence, the time series of physical volume of executed trades might be sensibly different across precious metals, both for the different trading activity, and for the different sizes of the MTA. In fact, as shown in Table (1), the traded volume is the highest for silver (more than 1.1 billion of traded ounces) and lower for gold, platinum, and palladium. Moreover, we observe that the trade size is the highest for silver (about 42.5 MTA per executed trade) and the lowest for gold (less than 1.3 MTA per trade). Platinum and palladium generally have a trade size larger than gold, 6.2 and 7.7 MTA per trade, respectively.

2.1 The variables of interest

From the whole database, we extract the relevant nanofrequency data, which are analyzed and aggregated in different ways, depending on the quantities of interest. We first recover prices and volume from the recorded trades. Within a given time interval, the price is defined as the price of the last trade recorded in the interval. If no trades occur, we replicate the last price from the previous interval. Returns are then computed as changes of log transaction prices. Volume is equal to the sum of the amount exchanged in the trades recorded within a given time interval, regardless of the trade side (buyer/seller initiated). If no trades occurred during a given time interval, then the volume is equal to zero. In the following, we denote the daily time index by t , the intradaily period by $i = 1, 2, \dots, N$ with N denoting the number of intradaily periods in day t . Intradaily periods have length equal to $1/N$ days, or $1440/N$ minutes. Furthermore, we indicate the end of period i of day t as (i, t) , and we use the following notation. The intradaily price sequence is denoted by $p_{i,t}$, while the intradaily volume time series is given as $v_{i,t}$; the intradaily log-returns are defined as $r_{i,t} = \log(p_{i,t}) - \log(p_{i-1,t})$, $i = 2, 3, \dots, N$, and $r_{1,t} = \log(p_{1,t}) - \log(p_{N,t-1})$.⁶

Finally, we filter quote data from outliers, most likely associated to errors in matching the asset identifier and the price.⁷ We choose a simple approach, excluding all quotes with a value 20% higher (for buyer initiated) or lower (for seller initiated) than the average trade price of the day.

The literature in financial economics and econometrics includes several liquidity measures (see the survey by Gabrielsen et al., 2011). In this study, we focus on two specific quantities that are exploiting part of the informative content of the database, namely the order flow, $OF_{i,t}$, and the percentage quoted spread, $QS_{i,t}$, defined as follows.

- Order flow: Let h denote the execution time of an order, and denote by x_h the trade direction of the order recorded at time h ; the trade direction is equal to -1 for seller-initiated trades and $+1$ for buyer-initiated trades; the order flow for interval i of day t is equal to $OF_{i,t} = \sum_{(i-1,t) < j \leq (i,t)} x_h$; by definition, the order flow assumes only integer values and can be either positive, zero or negative.⁸

⁶When computing the first return of the week on Sunday evening, we compute the returns with respect to the last price observed on Friday evening. In addition, we have only removed a few bank holidays (New Year's Day [2009 and 2010], Good Friday [2009 and 2010], and Christmas Day [2009]). Two days are not included in the database because of missing data: They are May 17, 2010, and May 18, 2010 (a Monday and a Tuesday).

⁷See, for example, Brownlees and Gallo (2006) for high-frequency data-cleaning techniques.

⁸Note that different orders (at different prices) might be recorded at the same time; these must all

- Percentage quoted spread: Let $A_{i,t}$ and $B_{i,t}$ denote the best ask and bid Prices available in the book at the end of period i at day t ; we define the midquote as $M_{i,t} = 0.5 (A_{i,t} + B_{i,t})$ and the percentage quoted spread as $QS_{i,t} = \frac{A_{i,t} - B_{i,t}}{M_{i,t}}$.

The percentage quoted spread is a standardized quantity allowing a direct comparison across precious metals because it is not dependent on the price of the MTA.

We study time series of the previously defined quantities at different frequencies. We analyze returns and volumes at the 5- and 60-minutes frequencies, where the first case is only considered for gold and silver, due to the higher number of trades recorded for these two metals. We inspect liquidity measures at the 15- and 60-minutes frequencies. When the frequency is set at 60 minutes, the dataset contains 11,880 observations. The number increases to 47,520 when considering the 15-minutes frequency and to 142,560 at the 5-minute frequency. We do not consider higher frequencies because of the large number of zeros that would be observed in the time series of interest (see the following section for further details).

3 Stylized facts of precious metals price, return, volume, and liquidity

In this section, we focus on those features that characterize the time series of returns, volume, and liquidity of our four precious metals. In particular, we evaluate the serial correlation properties of the first- and second-order moments, the distribution, and the existence of periodic patterns.

3.1 Prices and returns


The prices of precious metals follow an upward-sloping behavior across our sample period (see Figure (1)). This might depend on the so-called safe-haven effect, that is, the outflow from risky financial assets toward the precious metals, which are perceived as safer investments during periods of high uncertainty. We first analyze the series to determine their integration properties. ADF tests confirm that the log-prices follow a Random Walk model and suggest focusing on returns time series.⁹

Returns time series have patterns similar to those of equity returns and are characterized by volatility clustering, as well as the presence of extreme movements (see Figure (1)). The descriptive analysis, reported in Table (2), suggests that gold returns are less volatile than other precious metals, at both the 5- and 60-minutes frequencies. This finding is also matched with an inverse relation in average returns, the lowest being that of gold. We support that by the larger interest for gold compared with silver, platinum, and palladium, which might lead to higher efficiency for this precious metal. In turn, this leads to smaller risk and lower returns. The kurtosis is similar across metals at the hourly frequency, whereas at the 5-minute level, gold has a much smaller kurtosis than silver. Nevertheless, as we discuss in the following paragraph, this might be influenced by the large amount of zeros in the silver time series. Finally, we observe that the skewness is negative for silver, palladium, and platinum at the 60-minute frequency, a behavior

be considered in the construction of the order flow.

⁹ADF tests are available upon request.

Figure 1: Prices and returns plots.



Plots of 5-minutes price and returns data across the four precious metals. Prices on the left panels, returns on the right panels

Table 2: Descriptive analysis: Returns and Volume across metals

	Return						Volume					
	XAU	XAU	XAG	XAG	XPD	XPT	XAU	XAU	XAG	XAG	XPD	XPT
Frequency	5	60	5	60	60	60	5	60	5	60	60	60
Mean	0.0003	0.0038	0.0006	0.0079	0.0115	0.0049	1.6387	19.664	8.2311	98.773	13.937	14.321
Median	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	12	0.0000	25	0.0000	5
Std	0.0763	0.2582	0.1240	0.4499	0.5205	0.3369	3.5198	25.574	30.515	173.15	30.221	24.502
Kurtosis	30.09	15.643	62.863	15.361	17.266	10.577	51.952	18.375	100.90	22.274	41.431	30.443
Skewness	0.1031	0.3347	0.5407	-0.059	-0.569	-0.305	5.2819	3.0949	7.3061	3.5571	4.9074	3.9287
5% quant.	-0.104	-0.385	-0.039	-0.701	-0.795	-0.540	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
50% quant.	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	12	0.0000	25	0.0000	5
95% quant.	0.1070	0.3841	0.0545	0.7305	0.8564	0.5458	8	68	50	450	65	60
n. of 0	83592	1772	126534	5513	7063	5572	79842	1580	125250	5268	6354	5087

similar to that observed in equities. However, gold returns have positive skewness. Considering the range analyzed in this paper, 2009 and 2010, the trend of price time series, and the considerable interest in gold, this result is unsurprising. It suggests that trades on gold led to a consistent increase in the price of gold over time. When focusing on the returns ACF, we have statistically significant positive correlations for the first lag at the 5-minute level, whereas at the 60-minutes frequency, the first ACF becomes negative but is still statistically significant. The presence of negative correlation is expected on high-frequency data and might have different explanations: bid-ask bounce (Bollerslev and Domowitz, 1993), order imbalance (Flood, 1994), and/or trade behavior (Engle and Russell, 2009, among others). Such serial correlation has to be taken into account when developing models for the returns time series.

Regarding the trade activity, the proportion of zeros in a given interval of time can be seen as a measure of market liquidity (or illiquidity). For example, Bekaert et al. (2007) examine the impact of liquidity, proxied by the proportion of zero daily firm returns averaged over the month in emerging equity markets, on expected returns. Table (2) highlights that the number of zero returns is sensibly high. The percentage of zeros on hourly data is close to 45% for silver and platinum, peaks at 60% for palladium, and is minimum, about 15%, for gold. We report 5-minute descriptive statistics for gold and silver, the two most traded metals, showing that the number of zeros increases to 60% for gold and 90% for silver. The existence of such a large number of zeros makes the analysis on those time series challenging. In fact, on the one side, the zeros could make the identification of dynamic properties more difficult, but on the other side, the occurrence of zeros during the day, and their potential concentration during specific time ranges, is informative. As reported in Figure (2), zeros are more present during specific hours of the day, in particular at the 60-minute frequency.¹⁰ Following previous works, this means that those intervals are the most illiquid times during the 24-hour trading day. As we expected, they occur after the closing time of the main markets (the graphs time index starts at 10 PM GMT and the vertical lines refer to 8 AM GMT). Zeros are maximum between the closing time of North American markets (right side of the plots, around 8 PM GMT)

¹⁰Figure (2) refers to the gold. Similar patterns are present for the other metals, with the largest frequency of zeros for palladium and platinum. Graphs are available upon request.

and the opening of Asian markets (around midnight GMT, left side of the plots). On the contrary, during the opening hours of Asian Markets, about 3 AM to 9 AM GMT, European Markets, 9 AM to 5 PM GMT, and North American markets, about 3 PM to 8 PM GMT, the frequency of zeros is much lower, and starts increasing when the trading activity originated in Europe decreases. These findings suggest the possible presence of a periodic pattern in the return means. However, additional analysis does not support this idea.¹¹

Precious metals returns provide relevant information on the evolution of the conditional variances. The analysis of squared returns shows evidence of heteroskedasticity, with a clear periodic behavior (see the upper panel of Figure (3) for the gold time series). The pattern is more regular in the gold and silver cases compared to palladium and platinum. This is an expected result consistent with the large number of zeros present in the last two time series. Similar patterns can be identified at the 5-minute frequency for gold and silver. Notably, the oscillations of the ACF have a period of one day (24 hours). Confirmation of the periodic evolution of the intradaily variance is given in Figure (4), where we report, for the gold series, the hourly average squared return, computed as $\bar{r}_i^2 = \frac{1}{T} \sum_{t=1}^T r_{i,t}^2$. The graph shows evidence, within the day, of three relevant periods, which we can associate with the trading activities in different geographical areas (those are associated with three shaded areas corresponding to the trading activity in Tokyo, London-Zurich, and New York). The first increase in volatility corresponds to the Asian markets' trading activity, peaking around the opening of the Tokyo market around 3 AM GMT. The average squared returns then decrease until the opening of European markets. The increase peaks around 10 AM GMT and then it decreases until 1 PM GMT. The average returns sharply increase when American markets open and peak when American markets are active and European markets are closing, around 4 PM GMT. The same behavior is observed in the other three metals. This finding is, to our best knowledge, completely novel in the literature on precious metals, and resemble the results observed for round-the-clock data on exchange rates, see e.g. Dacorogna et al. (1993).

The periodic behavior of the intradaily volatility might be estimated and filtered out from the returns time series following different approaches.¹² Among the possible methods, we mention those of Andersen and Bollerslev (1997a) and Boudt et al. (2011). The first approach assumes that returns follow a multiplicative model


$$r_{i,t} = s_{i,t} \sigma_{i,t} \eta_{i,t}, \quad (1)$$

where $s_{i,t}$ is a periodic deterministic component, $\sigma_{i,t}$ is a GARCH-type variance, and $\eta_{i,t}$ is a standardized innovation with unit variance. The periodic term $s_{i,t}$ is estimated by means of a parametric regression model. The work of Boudt et al. (2011), despite using the same multiplicative model for the return, differs from the parametric approach of Andersen and Bollerslev (1997a) in two elements. First, $\sigma_{i,t}$ is an average volatility factor kept constant in a local window around $r_{i,t}$, and second, the deterministic component $s_{i,t}$ is estimated by means of a nonparametric approach. Notably, the approach of Boudt et al.

¹¹Graphical analyses supporting this claim are available upon request. We also estimated models allowing for periodic evolution of mean returns, but results were against this evidence with parameters statistically not significant.

¹²The same procedures to estimate periodic deterministic components have been implemented in different areas. For example, Caporin and Preš (2012) model and forecast wind speed intensity and Caporin and Preš (2011) forecast temperature indices density.

Figure 2: Frequency of zero returns.

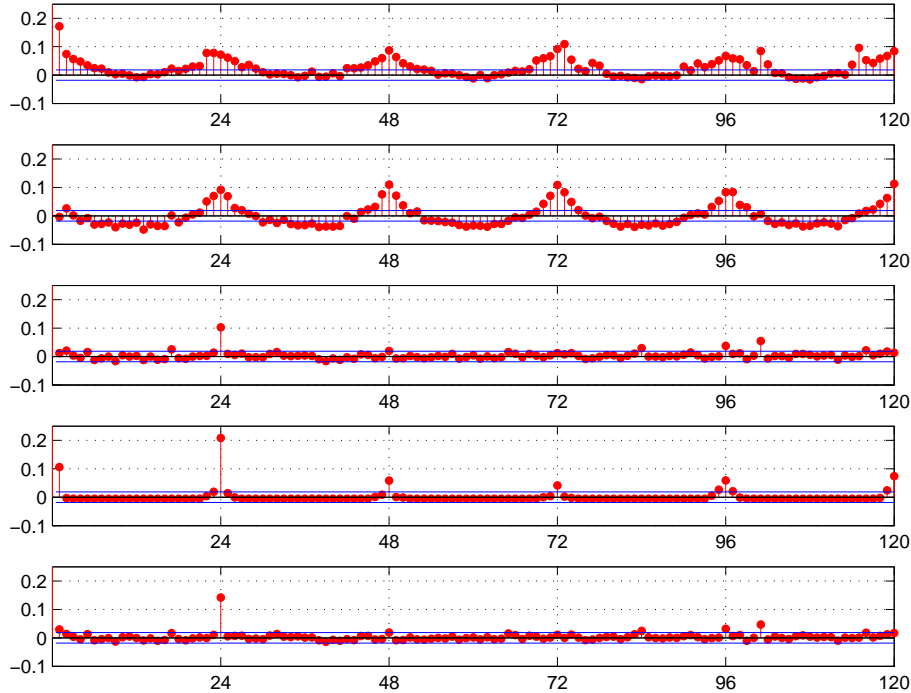


Occurrence of zero returns during the trading day for 5-minutes (left) and hourly (right) gold time series. Bars report percentages across the 5-minute/hourly intervals included in the dataset. The time index starts at 10 PM GMT. The vertical line refers to 8 AM GMT.

(2011) is robust to the presence of jumps in the high frequency returns. The estimation of the periodic component $s_{i,t}$ in Andersen and Bollerslev (1997a) considers a regression on harmonics of the log-transformed return series (in deviations from their unconditional mean). The two proposed approaches might provide similar results, and there is not a general criterion to be followed for the choice of the most appropriate method. Nevertheless, some metrics can be used. We might prefer flexibility, and thus choose Boudt et al. (2011). On the contrary, if we appreciate simplicity of implementation, Andersen and Bollerslev (1997a) is a good choice. Some insight for the comparison of the methods is given by the graphical contrast of the average squared returns \bar{r}_i^2 with the patterns estimated by two competing methods. The preferred approach is that with estimated patterns close to the empirical and naive estimate based on \bar{r}_i^2 . Finally, the search for periodic patterns in standardized returns can shed some light on the performances of both Boudt et al. (2011) and Andersen and Bollerslev (1997a) methods. Finally, we note that the presence of a large number of zeros might affect the performances of the Andersen and Bollerslev (1997a) approach. As this is the case for our returns data, we give our preference to the non-parametric approach of Boudt et al. (2011). Nevertheless, we consider a variance of the Andersen and Bollerslev (1997a) approach when dealing with volumes, characterized by a periodic pattern in their mean (as opposed to the returns case where the periodic evolution is observed on the variances).

Boudt et al. (2011) base the estimation of the periodic component in the intraday volatility on the standardized high-frequency return ($\bar{r}_{i,t}$), where the standardization factor is given by the square root of the normalized realized bipower variation of Barndorff-Nielsen and Sheppard (2004). By construction, $\bar{r}_{i,t}$ is distributed with mean zero and variance equal to the squared periodicity factor. Boudt et al. (2011) propose to estimate the periodic component using a nonparametric estimator of the scale of the standardized returns $\bar{r}_{i,t}$. They suggest three possible estimators: (1) the nonparametric periodicity estimator SD presented by Taylor and Xu (1997), which is based on the standard deviation of returns belonging to a local window, (2) the $ShortH$ estimator for the periodicity factor, based on the Shortest Half scale estimator (see Boudt et al., 2011, for details), and (3) the Weighted Standard Deviation estimator, WSD , of Boudt et al. (2011). Note that

Figure 3: ACF of the squared standardized returns for gold.




ACF of squared returns (upper panel), squared returns normalized with the realized bipower variation, RBV, (second panel) and of RBV-normalized squared returns once the periodic component has been filtered out with alternative methods: *SD*, *ShortH* and *WSD* (third, fourth and bottom panels, respectively).

the second and third estimators are also robust to the presence of price jumps. Given the nonparametric estimators of the periodic component, and similar to the approach of Andersen and Bollerslev (1997a), it is possible to recover the standardized return series.

Figure (4) displays the estimated periodic component for the hourly return series of gold with the three nonparametric methods presented by Boudt et al. (2011). The plot presents a behavior similar to the average squared returns. However, the ACF of the squared returns standardized series ($r_{i,t}/s_{i,t}$) (see Figure (3)), shows evidence of some residual periodic behavior. This finding is not influenced by the estimator adopted to capture the periodic behavior of squared returns. As a consequence, the nonparametric methods of Boudt et al. (2011) are not able to completely capture the periodic evolution characterizing the volatility of precious metals returns.¹³ This result suggests the possible presence of a stochastic behavior in the periodic component, which might be captured within an appropriate time-series framework. In the following section, we propose a parametric model that captures both the periodic evolution and the serial correlation of squared returns.

Figure 4: Estimated periodic component for gold squared returns.



Periodic component estimated by sample averages (black thick line) and by the alternative filtering methods of Boudt et al. (2011): *SD*, circles; *ShortH*, squares, and *WSD*, triangles. Shaded areas refer to the opening hours of Tokyo, London-Zurich, and New York exchanges, respectively (the last two partially overlap).

3.2 Volume

The volume time series are characterized by a percentage of zeros comparable to the returns time series. Note that volume and returns do not necessarily have the same occurrence of zeros as trades could be executed at the same price over consecutive intradaily periods. Table (2) presents descriptive analysis of the volume data, measured in numbers of MTA. We observe that the average volume is the highest for silver, at both the 5- and 60-minute frequencies. Moreover, when we recast the MTAs in ounces, we highlight that gold has the lowest average volume at the 60-minute frequency. This is further confirmed by the volume quantiles (see Table (2)). The volume time series show evidence of a strong periodic pattern: The correlograms are characterized by a cyclical behavior, and the intradaily volume averages, $\bar{v}_i = \frac{1}{T} \sum_{t=1}^T v_{i,t}$, are higher during the opening hours of the most active precious metals markets. The volume peaks first during the opening hours of Asian markets, then during the morning activity of European markets, and, finally, when both Europe and North America are active. The local troughs are associated with the break between Asia and Europe trading activity, and with the lunch break in Europe coupled with a pause waiting for first signals coming from North America before the afternoon trading.¹⁴ The increases in the volume level has a timing comparable to the increase in squared returns observed in Figure (4). We thus find patterns very close to those of the volatility, which supports the relevant impact of the main international

¹³Despite being sub-optimal, we also estimated the periodic component with the method proposed by Andersen and Bollerslev (1997a). Results are qualitatively similar to those obtained with the Boudt et al. (2011) approach, and thus suggest the possible presence of a stochastic periodic behavior.

¹⁴An example of the periodic pattern of trading volume is observable in Figure 6, where shaded areas refer to the trading activity in the main precious metals markets.

markets on the trading activity of precious metals. The similarities between volume and volatility are also associated with the literature exploring the relationships across those two quantities, see Tauchen and Pitts (1983), Admati and Pfleiderer (1988), and Andersen (1996), among many others.

Similarly to volatility, the periodic behavior of volume time series might be deterministic, stochastic, or a mixture of both deterministic and stochastic elements. Among the different approaches that are available to filter the periodic component from volume time series, we mention the use of seasonal adjustment methods, which can be based on moving averages or regression approaches and might be either multiplicative or additive. We prefer the latter as it could be less sensitive to the relevant fraction of zeros present in the volume time series. As a first analysis of volume, we propose the use of regression methods based on harmonics. We assume that the volume mean is given as follows (the time index evolves at an intradaily frequency):

$$v_t = \alpha + \sum_{i=1}^p \delta_i t^i + \sum_{j=1}^q \left(\gamma_j \cos \left(\frac{2\pi jt}{24} \right) + \phi_j \sin \left(\frac{2\pi jt}{24} \right) \right) + \zeta_t. \quad (2)$$

The periodic mean component is composed by a constant, a polynomial trend, and a combination of harmonics that capture the intradaily periodic behavior. The use of harmonics makes the estimation of periodic components similar to that adopted by Andersen and Bollerslev (1997a) for the volatility. We estimate the parameters with ordinary least squares (OLS) using robust standard errors due to the possible presence of serial correlation and heteroskedasticity in the innovations. We note that the regression provides an expected hourly volume replicating the periodic behavior, but residuals are still characterized by a strong periodic evolution. Moreover, the slow decay of both the volume ACF and the volume residuals ACF might suggest the presence of long memory.¹⁵ In the next section, we consider time-series models that capture both the periodic intradaily dynamics and the behavior of the volume.

3.3 Liquidity: Order flow and percentage quoted spread

Moving to the liquidity measures, we first point out a relevant difference between order flow (OF) and percentage quoted spread (PQS): OF has a number of zeros comparable to those observed for returns and volume (see Table (3)), while the PQS time series does not have zeros. This is a consequence of the different NF data used to evaluate the two time series. In fact, OF comes from trade data, whereas PQS depends on book-level data. OF time series of gold shows evidence of much larger variability, compared with the other precious metals. This is a consequence of gold attracting the largest number of trades. Notably, the OF is, on average, negative for gold and positive for silver, platinum, and palladium.

The OF time series are negatively skewed (with the exception of platinum) and highly leptokurtic (due to the overconcentration around the mean; see the quantiles reported in Table (3)). The PQS time series have similar unconditional behavior at the 15- and

¹⁵The ACF of the residuals and the estimated periodic component for the gold series are available in the appendix.

Table 3: Descriptive analysis: OF and QS across metals

	Order Flow						Quoted Spread					
	XAU	XAU	XAG	XAG	XPD	XPT	XAU	XAU	XAG	XAG	XPD	XPT
Frequency	15	60	15	60	60	60	15	60	15	60	60	60
Mean	-0.174	-0.699	0.0007	0.0030	0.1084	0.1733	0.0013	0.0012	0.0035	0.0034	0.0098	0.0052
Median	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0005	0.0006	0.0024	0.0025	0.0073	0.0038
Std	3.7337	9.0484	1.0721	2.3514	2.2681	2.5741	0.0049	0.0042	0.0045	0.0039	0.0080	0.0055
Kurtosis	22.350	15.720	57.206	18.481	20.964	21.106	1084.5	1553.0	245.13	85.934	29.615	115.47
Skewness	-0.342	-0.547	-1.048	-0.093	-0.369	0.7425	26.257	30.551	10.422	7.0002	3.8557	7.75
5% quant.	-6	-14	-1	-4	-3	-3	0.0003	0.0003	0.0010	0.0010	0.0036	0.0016
50% quant.	0	0	0	0	0	0	0.0005	0.0006	0.0024	0.0025	0.0073	0.0038
95% quant.	5	12	1	4	3	4	0.0040	0.0035	0.0089	0.0083	0.0242	0.0126
n. of 0	18424	2305	35787	5937	6852	5737	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

60-minute frequencies (see Table (3)). This result is a by-product of the methodology adopted to compute these quantities, which are determined as intradaily periods averages. As expected, the spreads are on average smaller for gold and higher for platinum and palladium. The dispersion is minimum for silver and maximum for palladium. The PQS series show evidence of positive asymmetry and of extremely long upper tails (see the quantiles reported in Table (3)): For Gold at the 15-minute frequency, the average spread between best bid and ask quotes is around 13 basis points, whereas the upper 99% quantile of PQS reaches the 150 basis points; large values are observed for palladium, where the average spread is close to 100 basis points but peaks at more than 425 basis points at the 99% quantile. Such large values of the PQS depend on the activity in the EBS platform, which further depends on the timing of the day. In fact, PQS time series have a clear intradaily pattern, with the largest values observed between the closing of American markets and the opening of Asian markets. This pattern is observed in the average hourly PQS of gold, $\bar{Q}S_i = \frac{1}{T} \sum_{t=1}^T QS_{i,t}$, and the ACF of the PQS time series. Oscillations from the opening of Asian markets to the close of American ones are much less pronounced.¹⁶

The OF time series do not show peculiar periodic behaviors in their mean, even if there is a clear evidence of serial correlation. As a further descriptive analysis, we also evaluate the serial correlation and periodic behavior of the squared order flow, which can be considered a proxy of volatility. In fact, an increase in the order flow, regardless of the sign, shows evidence of an increase in trading activity in the market in one specific direction (either an increase in seller- or buyer-initiated trades). The squared OF has intradaily patterns similar to the squared returns, with a clear increase during the opening hours of the most active precious metals markets, peaking at opening of Asian and European markets, and while Europe and North America are jointly active. The evidence of intra-daily periodic evolution in liquidity measures, in particular within a round-the-clock framework, and its association with the trading activity originated from different geographical areas, is a novel contribution to the literature of precious metals as well as of market liquidity. In fact, previous studies were just observing periodic patterns associated with a single market: among others, Brock and Kleidon (1992) focus on intra-daily patterns of spreads and

¹⁶Similar patterns are present for silver, palladium, and platinum. Figures are included in the appendix.

volume for the New York market; Chan et al. (1995) consider the bid-ask spread of stocks and options on the New York and Chicago exchanges; Chung et al. (1999) use again New York bid-ask spreads. We stress that the evidence of periodic behavior is observed on two different liquidity measures, and is associated with similar patterns on volatility and volume.

The estimation of the periodic behavior of OF and PQS might follow the same approaches outlined for the return volatility and the volume. Given the specific features of the two time series, we adapted the method of Boudt et al. (2011) on the OF, while the Andersen and Bollerslev (1997a) approach has been considered for the PQS time series. The analysis of OF and PQS time series standardized with the estimated deterministic periodic patterns, highlights that the two methods largely but not completely capture the periodic evolution, a result qualitatively similar to that observed for the volume. We thus apply in the following models whose aim is to estimate both the deterministic and the stochastic behavior of the series.

4 Dynamic modeling of precious metals time series

The previous section shows that precious metals time series are characterized by periodic behaviors. Those patterns are generally captured a priori, before the implementation of dynamic models of, say, the ARMA and GARCH families, which are used to describe the dynamic evolution of high-frequency time series.¹⁷ However, the use of alternative methodologies for filtering periodic patterns leads to standardized series that are still affected by periodic components. Therefore, the use of two-stage estimation methods might not be appropriate for precious metals time series and calls for time-series models capturing both the nonperiodic dynamic and the periodic behavior. The literature has proposed several models starting with the use of Seasonal ARMA models, with an appropriate selection of the period, up to the models with periodic long memory in the mean (Gray et al., 1988, and Woodward et al., 1998). Moreover, Bollerslev and Ghysels (1996), Guegan (2000), and Bordignon et al. (2007, 2009) propose GARCH-type models with periodic coefficients and periodic long memory. Nevertheless, the periodic behavior and the nonperiodic dynamic can be captured resorting to models in which the ARMA- and GARCH-type equations are given as a combination of both deterministic and stochastic components. In fact, those approaches allow for the presence of both a deterministic and a stochastic periodic behavior of a given time series. We thus specify models that fulfill this purpose. As we might have alternative specifications for a given time series, general approaches for model comparison can be used. Those might point at the diagnostic checking of residuals, the best model being that with residuals closer to a white noise. Alternatively, tests can be used across nested specifications. Finally, models can be compared ones a specific application of the model outcomes is chosen. In the last case, we have both statistical forecasting of given quantities, or the construction of trading strategies, among others. We finally stress that other elements might play a role, such as the model flexibility or its computational complexity. In the present paper, we focus on simple diagnostic checking tools based on the serial correlation properties of estimated residuals. Further evaluations are left to future researches.

¹⁷Two-step approaches are computationally simple but imply a loss of efficiency compared with models where periodic patterns are estimated together with the parameters driving the series dynamic.

Table 4: Estimation results for gold returns

μ	δ_1	ω	θ_1	θ_2	θ_3	θ_6	θ_{12}	θ_{24}	β_1	β_2	β_3
$-1.4e-09$ ($2.2e-05$)	-0.038^b (0.019)	-0.241 (0.383)	0.123^a (0.025)	0.138^b (0.055)	0.078 (0.053)	-0.009 (0.022)	-0.016^c (0.008)	0.154^a (0.042)	0.007 (0.034)	-0.012^c (0.007)	0.233^a (0.058)
β_6	β_{12}	β_{24}	γ_1	ϕ_1	γ_2	ϕ_2	γ_3	ϕ_3	γ_4	ϕ_4	γ_5
0.190 (0.126)	0.166^a (0.051)	0.390^a (0.091)	-0.516^a (0.130)	-0.372^a (0.087)	-0.123 (0.105)	0.253^a (0.049)	-0.177^b (0.090)	0.350^a (0.048)	-0.199^a (0.054)	-0.014 (0.024)	-0.251^a (0.046)
ϕ_5											<i>LLF</i>
0.140^a (0.039)											56787.14
μ	δ_1	ω	θ_1	θ_2	θ_3	θ_6	θ_{12}	θ_{24}	β_1	β_2	β_3
$-1.4e-09$ ($2.2e-05$)	-0.038^b (0.019)	-0.107 (0.067)	0.139^a (0.034)	-	0.097^c (0.056)	-0.107^b (0.045)	0.054 (0.121)	0.051 (0.193)	0.625^a (0.061)	-	-0.243^b (0.099)
β_6	β_{12}	β_{24}	γ_1	ϕ_1	γ_2	ϕ_2	γ_3	ϕ_3	γ_4	ϕ_4	γ_5
0.658^a (0.232)	-0.779^c (0.416)	0.728 (0.526)	-0.071^b (0.034)	-0.138^c (0.079)	-0.048^a (0.014)	0.350^a (0.049)	-0.040 (0.029)	0.292^a (0.033)	-0.257^a (0.051)	0.150^a (0.036)	-0.073 (0.048)
ϕ_5											<i>LLF</i>
0.257^a (0.036)											56745.73

Note: Estimation results for the return series of gold (XAU): EGARCHX (top panel) and EGARCHX-HAR (bottom panel). Sample period: 27th of December 2008 to the 30th of November 2010, hourly series, 11880 observations. *LLF* is the Log-likelihood function. Standard errors in parentheses; "a", "b" and "c" indicate significance at the 1%, 5% and 10% confidence level, respectively.


4.1 Return

For the returns, we specify an EGARCHX(P,O,Q) with periodic explanatory variables. The EGARCH model couples flexibility with computational simplifications; in fact, by resorting to exponential specifications, we avoid the introduction of parameter restrictions that lead to conditional variance positivity. This is particularly useful when the model includes exogenous variables as in our case, and makes this model class preferable to other GARCH-type specifications. To simplify the notation, we assume that the time index evolves at an intradaily step. The proposed model has the following structure:

$$\begin{aligned}
 r_t &= \mu + \sum_{j=i}^p \delta_j r_{t-j} + \sigma_t \varepsilon_t \\
 \ln(\sigma_t^2) &= \omega + \sum_{j=1}^Q \beta_j \ln(\sigma_{t-j}^2) + \sum_{j=1}^O \alpha_j \varepsilon_{t-j} + \sum_{j=1}^P \theta_j |\varepsilon_{t-j}| + \\
 &\quad + \sum_{j=1}^q \left(\gamma_j \cos\left(\frac{2\pi jt}{24}\right) + \phi_j \sin\left(\frac{2\pi jt}{24}\right) \right).
 \end{aligned} \tag{3}$$

The model includes an autoregressive component that captures the limited serial correlation in the mean. The variance dynamic depends on a standard EGARCH structure, with orders governing the autoregressive behavior of the log-conditional variances (order Q), as well as the impact of shocks size and sign (orders P and O , respectively). Moreover,

Figure 5: Estimation results for gold returns.



Estimates based on hourly data. On the left, ACF of standardized squared residuals of two returns models (the upper is an EGARCHX, the lower is an EGARCHX-HAR). On the right, mean hourly squared return and estimated periodic patterns from the two fitted models: EGARCHX with circles and EGARCHX-HAR with squares; the black thick line refers to the mean of squared returns. Shaded areas refer to the opening hours of Tokyo, London-Zurich, and New York exchanges, respectively (the last two partially overlap).

the introduction of q harmonics captures the deterministic evolution of log-conditional variances. We also stress that the model orders Q , P , and O can be increased beyond the common practice of restricting them to one. In fact, the orders set equal to the number of intradaily intervals per day, can detect Seasonal GARCH-type behaviors.

Several authors have also pointed out the presence of long memory in high-frequency returns volatility (see Andersen and Bollerslev, 1997b, 1998, and Bordignon et al., 2007, 2009, among others). Long-range dependence might be captured by resorting to long memory EGARCH specifications, as in Bollerslev and Mikkelsen (1996). However, the introduction of long memory in the conditional variance equation further increases the model's complexity. We thus prefer to specify the variance dynamic following a HAR-type structure (e.g. Corsi (2009)).

We suggest the estimation of the following EGARCHX-HAR(P,O,Q) model. This specification approximates the long-memory behavior reproducing the volatility persistence terms of the HAR model of Corsi (2009). Whereas in the HAR model the auto-regressive dynamic is associated with the target period of different market operators (daily, weekly, and monthly), in our specification the volatility evolves according to terms related to daily and intradaily periods: the day, the half-day, etc. Depending on the frequency of the time series, we include sums of past volatilities or shocks over different horizons; for hourly data we consider periods equal to the day, 24 hours, the half-day, 12 hours, and to 6 and 3 hours. The EGARCHX-HAR(P,O,Q) is characterised by the following equation:

$$\begin{aligned}
 \ln(\sigma_t^2) = & \omega + \sum_{j=1,3,6,12,24} \frac{\beta_j}{j} (\ln(\sigma_{t-j}^2) + \dots + \ln(\sigma_{t-1}^2)) + \\
 & + \sum_{j=1,3,6,12,24} \frac{\alpha_j}{j} (\varepsilon_{t-1} + \dots + \varepsilon_{t-j}) + \sum_{j=1,3,6,12,24} \frac{\theta_j}{j} (|\varepsilon_{t-1}| + \dots + |\varepsilon_{t-j}|) + \\
 & + \sum_{j=1}^q \left(\gamma_j \cos\left(\frac{2\pi jt}{24}\right) + \phi_j \sin\left(\frac{2\pi jt}{24}\right) \right). \quad (4)
 \end{aligned}$$

In the analysis of the precious metals returns, we consider different combinations of the EGARCHX and EGARCHX-HAR model orders, as well as different number of harmonics. We augment the model by the introduction of an AR(1) term, which is needed to capture the limited serial correlation observed on mean returns. Table (4) reports the best specifications for the gold series. They include five harmonics and lags up to order 24 (the day when focusing on hourly time series). Notably, the shock's sign was irrelevant (the EGARCHX order O was then set to zero). In both the EGARCHX and EGARCHX-HAR specifications, the lag 24 and the five harmonics parameters are statistically significant. The left panel of Figure (5) presents ACF of the standardized squared returns $((r_t/\hat{\sigma}_t)^2)$ for the EGARCHX(P,O,Q) and for the EGARCHX HAR(P,O,Q) models. The squared standardized residual series show evidence of serial correlation for both specifications. In particular, the first lag in the EGARCHX residuals is significant, and a daily periodic component remains in the EGARCHX-HAR residuals case. The serial correlation in the ACF of the EGARCHX standardized residuals might signal the existence of mild long-memory behavior which is not appropriately taken into account by the model.¹⁸ The EGARCHX-HAR model captures the potential long-range dependence of the volatility, but it is not able to completely remove the periodic component. Nevertheless, both specifications seem adequate for modeling returns, with the latter being preferred as it provides a higher value of the likelihood function and of the associated information criteria.¹⁹ Similar results are obtained for the other precious metals. The dynamic evolution of precious metals returns is similar to that observed on high frequency data for equities, with a limited and negative serial correlation. Moreover, the conditional variance evolution requires two relevant elements: first, the periodic component, capturing the increase in variance associated with the different active markets; second, the HAR term mimicking the long-range dependence, and capturing the persistent effect of shocks.

4.2 Volume

As pointed out in the previous section, the volume time series shows evidence supporting the presence of a stochastic periodic behavior, coupled with the possible presence of long-range dependence.²⁰ To capture such a feature, we consider a multifactor GARMA model that allows for long-memory behavior which might be associated with specific periodic frequencies.

Following Woodward et. al (1998), the multifactor GARMA model is defined by

$$\Phi(L) \prod_{j=0}^h (1 - 2\cos(\omega_j)L + L^2)^{d_j} (y_t - \mu) = \Theta(L)\epsilon_t, \quad (5)$$

where h is an integer, ϵ_t is a white noise with variance σ_ϵ^2 , μ is the mean of the process, ω_j (with $j = 0, \dots, h$) are the frequencies at which the long-memory behavior occurs, d_j (with

¹⁸In this work we prefer to keep a balance between model complexity and model feasibility. The estimation of a GARCH model with long-memory behavior, such as the FIEGARCH of Bollerslev and Mikkelsen (1996), could have provided a better fit. However, the estimation of the model is computationally complex, in particular when it is coupled with exogenous variables in the variance equation. We thus decided to stick to the simpler EGARCHX-HAR specification.

¹⁹The two specifications are not nested and a direct likelihood-ratio test is thus not available.

²⁰Bollerslev and Jubinski (1999) and Lobato and Velasco (2000), among others, document long memory in stock-market trading volume.

Table 5: Estimation results for gold Volume.

d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}	d_{11}	d_{12}
0.292 ^a	0.468 ^a	0.382 ^a	0.362 ^a	0.368 ^a	0.317 ^a	0.323 ^a	0.320 ^a	0.314 ^a	0.316 ^a	0.313 ^a	0.293 ^a
(0.008)	(0.031)	(0.035)	(0.081)	(0.053)	(0.046)	(0.055)	(0.059)	(0.061)	(0.108)	(0.057)	(0.107)
d_{13}	ϕ_1	ϕ_{24}	ϕ_{48}	ϕ_{72}	ϕ_{96}	ϕ_{120}	θ_1	μ	σ_ϵ^2	LLF	
0.143 ^a	-0.018	-0.257 ^a	-0.150 ^a	-0.097 ^c	-0.070 ^a	-0.002	-0.517 ^a	-26.62 ^a	356.8 ^a	-51768.53	
(0.010)	(0.159)	(0.069)	(0.025)	(0.055)	(0.013)	(0.061)	(0.104)	(6.688)	(35.77)		

Note: Estimation results for a GARMA model fitted on the volume series of gold (XAU). Sample period: 27th of December 2008 to the 30th of November 2010, hourly series, 11880 observations. LLF is the Log-likelihood function. Standard errors in parentheses; "a", "b" and "c" indicate significance at the 1%, 5% and 10% confidence level, respectively.


$j = 0, \dots, h$) are the long-memory parameters associated to each frequency, and $\Phi(L)$ and $\Theta(L)$ are the short-memory autoregressive and moving average polynomials with roots satisfying the usual stationarity and invertibility conditions. Stationarity of the GARMA model is achieved if the memory coefficients assume values below 0.5 for $1 \leq j \leq h-1$ and below 0.25 for $j = 0$ and $j = h$ (see Woodward et al., 1998). The most relevant element of the multifactor GARMA model is the so-called generalized Gegenbauer polynomial, given by

$$P(L) = \prod_{j=0}^h (1 - 2\cos(\omega_j)L + L^2)^{d_j}, \quad (6)$$

which may be considered as a generalized long-memory filter for the long-memory periodic behavior at $h + 1$ frequencies. The ω_j 's are the driving frequencies of a cyclical pattern of length S , where $\omega_j = (2\pi j/S)$, $h + 1 = [S/2] + 1$, and $[\cdot]$ refers to the integer part. Previous studies have shown that the GARMA model is able to replicate the periodic patterns similar to those observed in the volume time series (see Bordignon et al., 2007 and 2009). To estimate the $(h + 1)$ -factor GARMA model in (5), we implement an autoregressive approximation technique. Following Chung (1996), it is in fact possible to recover an $MA(\infty)$ or $AR(\infty)$ expansion of the model, and thus to estimate the model parameters through a quasi-maximum likelihood (QML) approach.

As previously observed, the autocorrelation function oscillates and decays slowly toward zero, suggesting a stochastic cyclical or periodic behavior coupled with long memory. From the previous section, we know that the periodic behavior has a length of 24 hours (on an hourly time series), leading to $S = 24$. In turn, such a value implies the presence of thirteen factors that capture the long-memory and cyclical behavior ($h = 12$ in the GARMA model of equation (5)). Table (5) reports the estimation results for the volume time series of gold. The memory parameters are statistically significant at all frequencies. However, we note that the memory coefficients associated with the zero frequency leads to a nonstationary component (the memory parameter is marginally larger than 0.25). We observe the same result in two other metals. The introduction of autoregressive and moving average components result in an improvement of the fit of the model compared with a pure long-memory specification; the short-memory coefficients are all statistically significant. We observe that the lags we introduce in the AR polynomial mimic the lags associated with the length of the periodic oscillation and its multiples. Such a result

Figure 6: Estimation results for gold volume.



Estimates based on hourly data. On the left, ACF of the residuals of two volume models (the upper is a GARMA without short-memory dynamic, the lower is a GARMA with short-memory dynamic). On the right, mean hourly volume and estimated periodic patterns from the two fitted models: GARMA without short memory in circles and GARMA with short-memory in squares; the black thick line refers to the hourly mean of volume. Shaded areas refer to the opening hours of Tokyo, London-Zurich, and New York exchanges, respectively (the last two partially overlap).

might suggest the presence of a weekly (five days) periodic pattern. Nevertheless, the estimation with $S = 120$ (one-week cyclical behavior) provides inferior fit to the data.²¹ The left panel of Figure (6) reports the ACF of the GARMA model residuals for the two different specifications for the short-memory component of the GARMA model. Notably, when considering the larger specification, few values of the correlograms are statistically significant. Moreover, the right panel of Figure (6) includes a comparison of the average hourly volume and of the average hourly fitted volume. The two quantities are close, further supporting the ability of the GARMA model in replicating the behavior of the volume time series. Our findings suggest the relevance of long-memory in the dynamic evolution the volume, coupled with the presence of a stochastic periodic behavior. The latter element is associated with the trading hours of the most relevant financial markets. These results could be used for an appropriate point forecasting of volume levels with potential applications in precious metal trading.

Finally, unreported ACF of the squared residuals show the presence of a periodic pattern. This feature might be captured by introducing a GARCH or EGARCH equation with periodic exogenous variables, similarly to what have been considered for the returns. Such an additional model component could be of interest for those aiming at forecasting the volume density.

4.3 Order flow and percentage quoted spread

The PQS and OF time series follow behaviors similar to trading volume and returns, respectively. In fact, the PQS has a relevant periodic pattern in the mean, whereas the OF mean values show little evidence of serial correlation. On the contrary, the absolute or squared values of OF are characterized by a strong periodic behavior. Moreover, deter-

²¹The GARMA model might collapse on a Seasonal ARFIMA specification, see the next section for a brief description of the latter model. The two models are equivalent if the memory coefficients of the GARMA satisfies $2d_1 = 2_2 = \dots = d_{12} = 2d_{13}$. In the present case this null hypothesis is rejected within a likelihood ratio framework.

Table 6: Estimation results for gold QS .

d	ϕ_1	ϕ_{24}	ϕ_{48}	ϕ_{72}	ϕ_{96}	ϕ_{120}	θ_1	μ	γ_1	ϕ_1	γ_2
0.076 ^a	-0.072 ^a	0.090 ^a	0.036 ^a	0.049 ^a	0.070 ^a	0.292 ^a	0.227 ^a	-0.134 ^a	-0.120 ^a	0.014	-0.101 ^a
(0.010)	(0.022)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.022)	(0.014)	(0.011)	(0.011)	(0.011)
ϕ_2	γ_3	ϕ_3	γ_4	ϕ_4	γ_5	ϕ_5	σ_ϵ^2				LLF
-0.002	-0.085 ^a	-0.005	-0.071 ^a	-0.009	-0.057 ^a	-0.011	0.128 ^a				-4674.295
(0.011)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.001)				

Note: Estimation results for an ARFIMA model augmented with harmonics fitted on the QS series of gold (XAU). Sample period: 27th of December 2008 to the 30th of November 2010, hourly series, 11880 observations. LLF is the Log-likelihood function. Standard errors in parentheses; "a", "b" and "c" indicate significance at the 1%, 5% and 10% confidence level, respectively.

ministic periodic filters are not effective in removing the periodic behavior of the liquidity time series. Given these findings, the liquidity measures' dynamic features might be captured using the approaches we applied to the time series of trading volume and returns, namely the GARMA and EGARCH specifications. As an alternative methodology, we consider the ARFIMA and Seasonal ARFIMA (SARFIMA) models extended with the inclusion of periodic explanatory mean variables. We estimate these two models on the PQS mean and on the squared OF values. While for the PQS the models represent the most appropriate specifications, for the OF we find that the direct modeling of the levels was not satisfactory, and thus decided to analyze the squared OF, a proxy of the volatility, thus sharing similar features. The SARFIMA is a special case of the multifactor GARMA model. Similar to the most general model, it provides periodic behavior coupled with long memory. The SARFIMA model is given as follows:


$$\Phi(L)(1 - L^S)^d(y_t - \mu - \sum_{j=1}^q \left(\gamma_j \cos\left(\frac{2\pi jt}{24}\right) + \phi_j \sin\left(\frac{2\pi jt}{24}\right) \right)) = \Theta(L)\epsilon_t, \quad (7)$$

where ϵ_t is a white noise with variance σ_ϵ^2 . The autoregressive and moving average polynomials $\Phi(L)$ and $\Theta(L)$ satisfy the usual restrictions for stationarity and invertibility, whereas the memory parameter d gives a stationary model if its value is below 0.5. The seasonal long-memory behavior influences the observed variable y_t in deviation from its unconditional mean μ and from a deterministic periodic behavior captured by the q harmonics. The relation between SARFIMA and GARMA is given by the following decomposition of the seasonal long-memory filter

$$(1 - L^S) = (1 - 2\cos(\omega_0)L + L^2)^{\frac{1}{2}} \left[\prod_{j=1}^{S-1} (1 - 2\cos(\omega_j)L + L^2) \right] (1 - 2\cos(\omega_s)L + L^2)^{\frac{1}{2}}, \quad (8)$$

where $\omega_0 = 1$ and $\omega_s = -1$. The previous decomposition takes into account the roots of the polynomial $(1 - L^S)$, which are associated with frequencies in $0 - \pi$. In particular, the frequencies are $0, \pi$ (if S is even), and a set of frequencies depending on the value of S , each associated with a pair of roots of the polynomial. Notably, such a decomposition corresponds to a product of Gegenbauer polynomials. If we introduce long memory and

Figure 7: Estimation results for gold QS .



Estimates based on hourly data. On the left, ACF of the residuals of two QS ARFIMA models differing in their short-memory dynamic, the upper model includes a limited set of periodic lags compared to the lower panel model. On the right, mean hourly QS and estimated periodic patterns from the two fitted models: restricted ARFIMA in circles and the more flexible ARFIMA with squares; the black thick line refers to the hourly mean of volume. Shaded areas refer to the opening hours of Tokyo, London-Zurich, and New York exchanges, respectively (the last two partially overlap).

consider $(1 - L^S)^d$, the exponent of each Gegenbauer polynomial in (8) is either equal to d or to $d/2$ (this happens for frequencies ω_0 and ω_1). As a consequence, the SARFIMA model is a special case of the multifactor GARMA under a restriction on the memory coefficient, and if the frequencies over which the GARMA model is specified are exactly the same set of frequencies associated with the decomposition of the seasonal filter $(1 - L^S)$. In preliminary estimates, not reported for reasons of space, the estimate of GARMA specifications and the subsequent tests for the equality of memory coefficients, suggested the use of the less parameterized SARFIMA.

In the analysis of the liquidity series,²² we consider different specifications for the models. We set three different values for the seasonal length (S) 1, 24, and 120. With $S = 1$, we specify a pure long-memory model for the hourly series, whereas in the case of $S = 24$ or 120, we specify a daily and weekly seasonal integration pattern. To capture the periodic behavior of the series, we consider up to six daily harmonics, and we include a weekly harmonic. Finally, to model the short-memory component, we introduce different autoregressive and moving average specifications, considering lags up to the week.

For the QS , Table (6) presents the results of the model that best fits the gold series. Although we try with the ARFIMA and SARFIMA models, we concentrate on the first kind of models ($S = 1$). Estimation results and the ACF of the residuals of SARFIMA models are similar. The long-memory parameter is significant for all the metals, but it is lower for gold than the other three metals. In the first case, it is equal to 0.076 whereas it is near 0.28 for silver and palladium and 0.493 for platinum. The introduction of the autoregressive lags improves the fitting of the model. Note that the lags are all statistically significant.²³ The analysis of the ACF of the residuals, in the left panel of Figure (7), favors the introduction of short-memory component at the daily and its multiple lags. Moreover, it displays significant correlations associated with particular lags. This fact is more evident in the gold series. We believe that they are neither associated

²²The estimated series are equal to $QS \times 100$ and $OF^2/100$, respectively.

²³A likelihood ratio test between an ARFIMA(1 24 48 72 96 120,d,1) and ARFIMA(1 24 120,d,1) specification rejects the restricted model in the four metals at 5% level.

Table 7: Estimation results for gold OF^2

d	ϕ_1	ϕ_{24}	ϕ_{48}	ϕ_{72}	ϕ_{96}	ϕ_{120}	θ_1	μ	γ_1	ϕ_1	γ_2
0.046 ^a	0.294 ^a	0.029 ^a	0.027 ^a	0.021 ^b	0.020 ^b	0.048 ^a	-0.191 ^a	-0.826 ^a	0.587 ^a	0.452 ^a	0.309 ^a
(0.014)	(0.062)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.058)	(0.061)	(0.059)	(0.059)	(0.055)
ϕ_2	γ_3	ϕ_3	γ_4	ϕ_4	γ_5	ϕ_5	σ_ϵ^2				LLF
-0.268 ^a	-0.150 ^a	-0.245 ^a	0.093 ^c	0.237 ^a	0.017	-0.013	9.220 ^a				-30052.31
(0.055)	(0.051)	(0.051)	(0.048)	(0.048)	(0.045)	(0.045)	(0.119)				

Note: Estimation results for an ARFIMA model augmented with harmonics fitted on the OF^2 series of gold (XAU). Sample period: 27th of December 2008 to the 30th of November 2010, hourly series, 11880 observations. LLF is the Log-likelihood function. Standard errors in parentheses; "a", "b" and "c" indicate significance at the 1%, 5% and 10% confidence level, respectively.

with the long-memory component nor with the periodic pattern, which have both been correctly removed. We considered different model specifications, without, however, better results. The right panel of Figure (7) presents the mean fitted series, which substantially replicates the periodic component observed in the QS time series, picking when the most active markets are closed (between the close of US markets and the opening of Asian markets).

For the OF^2 , Table (7) displays the estimation result for the gold series. As in the previous case, we consider ARFIMA and Seasonal ARFIMA specifications and we find very similar outcomes. Then we focus on the pure long-memory model ($S = 1$). Estimations of the d range from 0.046 to 0.168 for the different models, and they are always highly significant. Regarding the short-memory specification, the autoregressive lags produce an improvement in the fitting of the model. We consider an ARFIMA(1 24 120,d,1) and an ARFIMA(1 24 48 72 96 120,d,1) model. In the case of the gold time series, the full specification presents significant coefficients, and it is preferred to the restricted one when formally testing.²⁴ Finally, ACF of the residuals displays a good performance of the models, and the mean estimated series reproducing the periodic component present in the OF^2 series, in Figure (8).


For both liquidity measures the proposed models largely capture the periodic evolution and long-range dependence. The former characteristic can be seen as a further confirmation of the relevance of active markets in the evolution of market liquidity. The latter characteristic points at the persistent impact of shocks. These two elements can be useful to forecast liquidity measures with potential applications in trading strategies for precious metals.

4.4 A first look at commonality across precious metals

The previous analyses show evidence of common patterns across precious metals suggesting the presence of commonality in liquidities of these commodities. While commonality in liquidities has been studied for stocks (e.g. Chordia et al. (2000), Hasbrouck and Seppi, 2001, Huberman and Halka, 2001), bonds (e.g. Chordia et al., 2005) and more recently

²⁴A likelihood ratio test between an ARFIMA(1 24 48 72 96 120,d,1) and ARFIMA(1 24 120,d,1) specification rejects the restricted model in the gold and palladium series at 5% level.

Figure 8: Estimation results for gold OF^2 .



Estimates based on hourly data. On the left, ACF of the residuals of two OF^2 ARFIMA models differing in their short-memory dynamic, the upper model includes a limited set of periodic lags compared to the lower panel model. On the right, mean hourly OF^2 and estimated periodic patterns from the two fitted models: restricted ARFIMA in circles and the more flexible ARFIMA with squares; the black thick line refers to the hourly mean of OF^2 . Shaded areas refer to the opening hours of Tokyo, London-Zurich, and New York exchanges, respectively (the last two partially overlap).

for foreign exchange markets (Mancini et al., 2013, and Karnaukh et al., 2013)), to our best knowledge only one paper analyzes commonality in liquidities for commodities using futures data, i.e. Marshall et al. (2012). However, there is no literature on commonality in liquidities of spot precious metals. As a first step, we standardize all series of interest in order to avoid any impact of the measurement scale. Then, following e.g. Hasbrouck and Seppi (2001) and Mancini et al. (2013) we analyze the common behaviors by looking at the first principal component of the time series of squared returns, trading volume, quoted spread and squared order flow. To do this, we use the full sample covariance across precious metals. Table (8) reports the main quantities. Two main results emerge. First, the first principal component (PC1) has always a positive impact on all metals for all variables of interest. Moreover, PC1 explains a large fraction of the total variance (from 34.5% of quoted spread up to the 51.5% of volume). The relevance of PC1 is further confirmed by the R^2 of the regression of each metal on PC1, see again Table (8), which is in most cases higher than 0.5. However, by extending the analysis to the first and second component (PC1+PC2) there is an improvement of the fit: the R-squared for the regression of each metal on the first and second components are sensibly higher, with a minimum at 0.48 and a maximum at 0.99; the fraction of variance explained by PC1 and PC2 ranges from 59% up to 74%. Second, commonality in liquidity of precious metals is strong. Compared to other markets, it is much stronger than that documented in stock markets.²⁵ and comparable to that observed in FX markets. Notice that our measures of commonality in liquidities are based on intraday data. We replicated our analysis using daily data and we found very similar results (available upon request). A novel contribution of our study is to document commonality in liquidities intradaily rather than at daily or lower frequencies, as done in the previous literature. Our results suggest that commonality in liquidities of precious metals depends on seasonalities of trading activities in different geographical areas.²⁶

²⁵For example Korajczyk and Sadka (2008) report adjusted-R2 ranging between 2% and 30%, depending on the liquidity measure

²⁶Plots of the hourly averages for the first principal component of squared return, volume, percentage quoted spread and order flow squared are available upon request.

Table 8: First results on commonality.

	Loadings		R-squared		Loadings		R-squared	
	PC1	PC2	PC1	PC1+PC2	PC1	PC2	PC1	PC1+PC2
	Squared returns				Volume			
Gold	0.55	0.48	0.57	0.78	0.53	0.46	0.57	0.77
Silver	0.58	0.37	0.62	0.75	0.52	0.48	0.56	0.77
Palladium	0.42	-0.60	0.33	0.67	0.46	-0.59	0.44	0.75
Platinum	0.44	-0.52	0.35	0.61	0.49	-0.45	0.50	0.69
Var. Expl.			46.7	70.5			51.7	74.4
	Quoted spread				Squared OF			
Gold	0.43	-0.56	0.26	0.57	0.36	-0.92	0.24	0.99
Silver	0.43	-0.55	0.26	0.56	0.51	0.11	0.47	0.48
Palladium	0.55	0.45	0.42	0.63	0.57	0.21	0.59	0.63
Platinum	0.56	0.41	0.44	0.61	0.54	0.31	0.53	0.61
Var. Expl.			34.5	59.2			45.9	67.8

Note: First and second columns (fifth and sixth): loadings for the row precious metal on the first and second principal components, respectively. Third and fourth (seventh and eighth) columns: R-squared for regressions of each precious metal on first and first and second principal components. Fifth and tenth rows: fraction of variance explained by the first and by the first plus the second principal components.

5 Conclusions

We provide a pioneering description of the stylized facts and dynamic properties of precious metals time series extracted from a unique nanofrequency database that includes trade and quotes data. The most innovative elements of our paper are (1) the time frequency of the database, up to 100 millisecond, (2) the analysis of spot precious metals including palladium and platinum that were not studied so far at least at this time granularity, (3) the analysis of trading activity recorded around the clock, and (4) the analysis of liquidity and commonality in liquidities of precious metals. Overall, our results show that the price, return, and volume time series have features comparable to those of traditional assets. Moreover, market liquidity of precious metals is characterized by intraday seasonalities and very strong commonality.

This work represents a preliminary research on the statistical aspects and data analysis of the precious metals. Building on our results, future research might be extended in several respects, for instance, to understand better the price discovery process and liquidity issues on spot precious metals. Alternatively, multivariate approaches might be considered, focusing on the joint analyses on different precious metals or on models capturing the interdependence between returns, volume, volatility, and liquidity.

References

- [1] Abergel, F., Bouchaud, J., Foucault, T., Lehalle, C., and Rosenbaum, M., 2012, *Market Microstructure: Confronting Many Viewpoints*, Wiley.
- [2] Admati, A.R. and Pfleiderer, P., 1988, A Theory of Intraday Trading Patterns: Volume and Price Variability, *Review of Financial Studies*, 1, 3-40.
- [3] Andersen, T.G., 1996, Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility. *Journal of Finance*, 51, 169-204.
- [4] Andersen, T.G., and Bollerslev, T., 1997a, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115-158.

- [5] Andersen, T.G., and Bollerslev, T., 1997b, Heterogeneous information arrivals and return volatility dynamics: uncovering the long run in high volatility returns. *Journal of Finance*, 52, 975-1005.
- [6] Andersen, T.G., and Bollerslev, T., 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review*, 39, 885-905.
- [7] Andersen, T.G., Bollerslev, T., Diebold, F.X., and Vega, C., 2003, Micro effects of macro announcements: real-time price discovery in foreign exchange, *American Economic Review*, 93, 38-62.
- [8] Baillie, R.T., Han, Y., Myers, R.J., and Song, J., 2007, Long-memory models for daily and high-frequency commodity future returns, *Journal of Future Markets*, 27-7, 643-668.
- [9] Bannouh, K., van Dijk, D., and Martens, M., 2009, Range-Based Covariance Estimation Using High-Frequency Data: The Realized Co-Range, *Journal of Financial Econometrics*, 7-4, 341-372.
- [10] Barkoulas, J., Labys, W.C., Onochie, J., 1997, Fractional Dynamics in International Commodity Prices. *Journal of Future Markets*, 17-2, 161-189.
- [11] Barndorff-Nielsen, O.E., and Shephard, N., 2004, Power and bipower variation with stochastic volatility and jumps, *Journal of Financial Econometrics*, 2, 1-37.
- [12] Bekaert, G., Harvey, C., and Lundblad, C., 2007, Liquidity and Expected Returns: Lessons from Emerging Markets, *Review of Financial Studies*, 20-6, 1783-1831.
- [13] Berger, D.W., Chaboud, A.P., Chernenko, S.V., Howorka, E., and Wright, J.H., 2008, Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data, *Journal of International Economics*, 75, 93-109.
- [14] Bollerslev, T., and Domowitz, I., 1993, Trading Patterns and Prices in the Interbank Foreign Exchange Market, *Journal of Finance*, 48, 1421-43.
- [15] Bollerslev, T., and Ghysel, E., 1996, Periodic Autoregressive Conditional Heteroscedasticity, *Journal of Business and Economic Statistics*, 14, 139-151.
- [16] Bollerslev, T., and Jubinski, D., 1999, Equity Trading Volume and Volatility: Latent Information Arrivals and Common Long-Run Dependencies, *Journal of Business and Economic Statistics*, 17-1, 9-21.
- [17] Bollerslev, T., and Mikkelsen, H.O., 1996, Modeling and pricing long memory in stock market volatility. *Journal of Econometrics*, 73, 151-184.
- [18] Bordignon, S., Caporin, M., and Lisi, F., 2007, Generalised Long Memory GARCH models for intradaily volatility, *Computational Statistics & Data Analysis*, 51-12, 5900-5912.
- [19] Bordignon, S., Caporin, M., and Lisi, F., 2009, Periodic Long Memory GARCH models, *Econometric Reviews*, 28, 60-82.

- [20] Boudt, K, Croux, C., and Laurent, S., 2011, Robust estimation of intraweek periodicity in volatility and jump detection, *Journal of Empirical Finance*, 18, 353-367.
- [21] Brock, W.A., and Kleidon, A.W., 1992, Periodic market closure and trading volume: a model of intraday bids and asks, *Journal of Economics Dynamics and Control*, 16, 451-489.
- [22] Brownlees, C.T., and Gallo, G.M., 2006, Financial econometric analysis at ultra-high-frequency: Data handling concerns, *Computational Statistics & Data Analysis*, 51 - 4, 2232-2245.
- [23] Cai, J., Cheung, Y., and Song, M.C.S., 2001, What Moves the Gold Market?, *Journal of Future Markets*, 21-3, 257-278.
- [24] Caporin, M., and Preś, J., 2011, Forecasting Temperature Indices Density with Time-Varying Long-Memory Models, *Journal of Forecasting*.
- [25] Caporin, M., and Preś, J., 2012, modeling and forecasting wind speed intensity for weather risk management, *Computational Statistics & Data Analysis*, 56 - 11, 3459-3476.
- [26] Chan, K., Chung, Y.P., and Johnson, H., 1995, The intraday behavior of bid-ask spreads for NYSE stocks and CBO options, *Journal of Financial and Quantitative Analysis*, 30 - 3, 329-346.
- [27] Chordia, T., Roll, R., and Subrahmanyam, A., 2000, Commonality in liquidity, *Journal of Finance*, 52, 328.
- [28] Chordia, T., Sarkar, A., and Subrahmanyam, A., 2005, An empirical analysis of stock and bond market liquidity, *Review of Finance*, 18, 851-29.
- [29] Chung, K.H., Van Ness, B.F., and Van Ness, R.A., 1999, Limit orders and the bid-ask spread. *Journal of Financial Economics*, 53, 255-287.
- [30] Chung, C., 1996, Estimating a generalized long memory process. *Journal of Econometrics*, 73, 237-259.
- [31] Corsi, F., 2009, A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7, 174-196.
- [32] Dacorogna, M.M., Muller U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V., 1993, A geographical model for the daily and weekly seasonal volatility in the foreign exchange market, *Journal of International Money and Finance* 12, 413-438.
- [33] Dacorogna, M.M., Gencay, R., Muller, U.A., Olsen, R., Pictet, O.V., 2001, *An introduction to high-frequency finance*, Academic Press, London.
- [34] Engle, R.F., Russell, J.R., 2009, Analysis of high-frequency Data, *Handbook of Financial Econometrics*, Eds. Y. Ait-Sahalia and L.P. Hansen, Elsevier, North Holland.
- [35] *Financial Times*, 2013, Cyprus to dive into its gold reserves, April 10th.

- [36] Fleming, J., Kirby, C., and Ostdiek, B., 2003, The economic value of volatility timing using realized volatility, *Journal of Financial Economics*, 67, 473-509.
- [37] Flood, M.D., 1994, Market structure and inefficiency in the foreign exchange market, *Journal of International Money and Finance*, 13, 131-158.
- [38] Gabrielsen, A., Marzo, M., and Zagaglia, P., 2011, Measuring market liquidity: an introductory survey, MPRA paper 35829, available at <http://mpra.ub.uni-muenchen.de/35829/>
- [39] Gray, H.L., Zhang, N., and Woodward, W., 1988, On generalized fractional processes, *Journal of Time Series Analysis*, 10, 233-257.
- [40] Guegan, D., 2000, A new model: the k-factor GIGARCH process, *Journal of Signal Processing*, 4, 265-271.
- [41] Hasbrouck, J., and Seppi, D.J., 2001, Common factors in prices, order flows, and liquidity, *Journal of Financial Economics*, 59, 383-411.
- [42] Huberman, G., and Halka, D., 2001, Systematic Liquidity, *Journal of Financial Research*, 24-2, 161-78.
- [43] Karnaukh, N., Ranaldo, A., and Soderlind, P., 2013, Understanding FX Liquidity, University of St. Gallen, School of Finance Working Paper.
- [44] Khalifa, A.A.A., Miao, H., and Ramchander, S., 2011, Return distribution and volatility forecasting in metal futures markets: evidence from gold, silver and copper, *Journal of Future Markets*, 31-1, 55-80.
- [45] Korajczyk, R.A., and Sadka, R., 2008, Pricing the commonality across alternative measures of liquidity, *Journal of Financial Economics* 87, 45-72.
- [46] Lee, C., and Ready, M., 1991, Inferring Trade Direction from Intraday Data, *Journal of Finance*, 46, 733-746.
- [47] Lobato, I., and Velasco, C., 2000, Long Memory in Stock-Market Trading Volume, *Journal of Business and Economic Statistics*, 18-4, 410-427.
- [48] Mancini, L., Ranaldo, A., and Wrampelmeyer, J., 2013, Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums, *Journal of Finance*, 68-5, 1805-1841.
- [49] Marshall, B.R., Nguyen, N.H., and Visaltanachoti, N., 2012, Commodity Liquidity Measurement and Transaction Costs, *Review of Financial Studies*, 25, 599-638.
- [50] Parlour, C.A., and Seppi, D.J., Limit order markets: a survey, *Handbook of financial intermediation and banking*, Thakor, A.V., and Boot A.W.A. (Eds.), North-Holland, Elsevier, Amsterdam.
- [51] Taylor, S.J., and Xu, X., 1997, The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance* 4, 317-340.
- [52] Tauchen, G.E., and Pitts, M., 1983, The Price Variability-Volume Relationship on Speculative Markets, *Econometrica* 51, 485-505.

- [53] Woodward, W.A., Cheng, Q.C., and Gray, H., 1998, A k-factor GARMA long-memory model, *Journal of Time Series Analysis* 19, 485-504.