Danila Deliya

Direction-of-change forecasts and trading strategy profitability at intra-day horizons

Working Paper # BSP/2007/087

This paper is based on the Master Thesis prepared at NES in 2007 in the framework of the research project "Issues in predictability in financial and macroeconomic markets" under the supervision of prof. S.A. Anatolyev (NES).

I am grateful to academic advisor and all participants of the project for their helpful comments and discussion.

Moscow
2007

Deliya Danila. Direction-of-change forecasts and trading strategy profitability at intra-day horizons / Working Paper #BSP/2007/87 – Moscow, New Economic School, 2007. – 25 p. (Engl.)

Recent theoretical studies have shown that it is possible to make relevant direction-of-change forecasts using volatility forecasts and mean return forecasts. We show that using ultra high frequency information in liquid markets for direction-of-change forecasts allows building profitable trading strategies operating on short horizons during the day. We use data from Russian stock market (Gazprom) and foreign exchange market (Euro / Japanese Yen). After choosing the optimal length for intra-day intervals, we build models for realized volatility and mean return forecasts, accounting for intra-day and intra-week seasonality. The volatility and mean return forecasts are used to build direction-of-change forecast, which in addition to first and second moments of returns distribution uses bipower variation and realized moments of higher order. The performance of trading strategies built on those direction-of-change forecasts is then studied using out-of-sample forecasts. Results indicate that profitability of the strategies highly depends on the size of transaction costs on the market. The profitability is negative for the Euro / Japanese Yen market, but it is positive for Gazprom market (18% annualized).

Key words: direction forecasts, financial markets, high frequency analysis, trading strategy

Делия Д.С. Предсказание направления изменения и доходность торговой стратегии на внутридневных горизонтах / Препринт #BSP/2007/87 - М.: Российская Экономическая Школа, 2007. – 25 с. (Англ.)

Недавние теоретические исследования показали возможность адекватного предсказания направления изменения цены актива на основе предсказаний волатильности и средней доходности. Мы показываем, что использование высокочастотной информации о ликвидных рынках для предсказания направления изменения позволяет создавать профицитные торговые стратегии, работающие на коротких горизонтах внутри дня. Мы используем данные с российского фондового рынка (обыкновенные акции Газпрома) и с международного валютного рынка FOREX (Евро / Японская йена). Выбрав оптимальный размер внутридневного интервала, мы строим модели для предсказания реализованной волатильности и средней доходности, корректируя на внутридневную и внутринедельную сезонности. Предсказания волатильности и средней доходности используются для построения предсказания направления изменения. Также, в дополнение к первому и второму моментам распределения доходностей, используются двустепенная вариация и реализованные моменты высших порядков. Затем проводится исследование результативности трейдинговой стратегии с использованием прогнозирования за пределы выборки. Результаты показали, что доходность стратегий сильно зависит от величины транзакционных издержек на рынке. Доходность отрицательна для валютного рынка Евро / Японская йена, и положительна для ранка акций Газпром (18% годовых).

Ключевые слова: предсказание направления изменения, финансовые рынки, высокочастотный анализ, торговая стратегия

ISBN

Contents

1 Introduction

A number of trading strategies developed for financial markets are based on direction-of-change forecasts. Despite the fact that future returns are unpredictable, a large amount of recent articles show that the direction-of-change forecasts can be made with great success (see, among others, Breen, Glosten and Jagannathan (1989), Gençay (1998), White (2000), Pesaran and Timmerman (2004), Christoffersen and Diebold (2006)).

Specifically, Christoffersen and Diebold (2006) show that in case of predictable volatility on equity returns, the induced sign dependence may be used to produce direction-of-change forecasts useful for market timing. Christoffersen, Diebold, Mariano, Tay and Tse (2006) advance this topic further, checking various direction-of-change forecasting models on 20 indexes in stock markets all over the world.

The amount of trade operations as well as their volume are growing extremely fast in developed and in emerging markets (see Harris (2003)). For example, the volume of trades on the Moscow Interbank Currency Exchange (MICEx) has increased more than 6 times during 5 years from 2000 to 2005[1], the average daily trading volume on foreign exchange markets increased from $1.0 trillion a day in 1992 to $1.6 trillion a day by 2003[2]. Consequently, the number of trading operations per day has increased too. Larger number of trading operations per unit of time reveals more information about the underlying price change process and allows to make market timing decisions in shorter periods of time.

One of approaches to direction-of-change forecasts is based on the volatility forecast. This subject is studied a lot in the recent literature, and the detailed description of mainstream approaches to volatility forecasting can be found in Andersen, Bollerslev, Christoffersen and Diebold (2006). In this work we use the benefits of realized volatility approach — specifically that it makes it possible to treat volatility as observable (as noted in Andersen, Bollerslev, Diebold and Labys (2000)).

For the purpose of this work we use ultra high frequency data on shares of 'Gazprom' company — one of the most liquid stocks in the Russian stock market, and on the bid-ask quotes of the Euro / Japanese Yen foreign exchange market. MICEx, the stock exchange where 'Gazprom' is traded[3], is a continuous order-driven electronic exchange, which opens its trading

[1] http://www.micex.ru/profile/years/

[2] http://www.moneyforex.com/about-forex.php

[3] Shares of 'Gazprom' are traded on two exchanges in Russia — Russian Trading System (RTS) and MICEx. We choose MICEx for our studies because the liquidity of 'Gazprom' is much higher on it than on RTS.

sessions with a call market auction[4]. On MICEx trading operations happen between buyers and sellers without the intermediation of dealers. Euro / Japanese Yen foreign exchange market is a quote-driven dealer market, where dealers quote the prices and supply all the liquidity.

First of all, we determine the optimal sampling frequency for calculating the realized volatility. We use the "volatility signature plot" described in Andersen, Bollerslev, Diebold and Labys (2000) to find an optimal sampling frequency for both datasets. Then we identify the smallest decision interval for direction-of-change forecasting that will be large enough to capture the benefits of the realized volatility approach. Later we use calculated realized volatility to estimate models for volatility forecasts. As a next step we proceed with various models for mean return forecasts. Both of those forecasts are later used for direction-of-change forecast - that is, estimating the probability that price will go up during the next decision interval. This probability forecast is later used to build a trading strategy.

To check the performance of trading strategies, we introduce a 'virtual' researcher that starts applying direction-of-change forecasting techniques on September 1, 2006. Every time the current decision interval is finished, she adds arrived information on transactions (for 'Gazprom' market) or on quotes (for Euro / Japanese Yen market) to her data pool, reestimates the models and makes a direction-of-change forecast for the next decision interval. After that she acts according to the forecast, taking either a short or a long position. We show that after accounting for transaction costs, cumulative return is negative for the Euro / Japanese Yen market, but it is positive for Gazprom market (18% annualized).

In the remainder of this paper we proceed as follows. We begin in Section 2 by formally describing the models for mean return and direction-of-change forecasts. Next, in Section 3 we summarize the empirical results, describe the data and software used during this work, and present direction-of-change forecasts and trading strategy evaluation. In Section 4 we conclude with suggestions for future research.

2 Theoretical Background

Let P_t denote price of an asset at time t, with unit of time t being equal to the size of decision interval. Then we calculate the return at time t as

$$R_t = 100(\log P_t - \log P_{t-1}).$$

We start with describing approaches to direction-of-change forecasts which use forecasts of

[4]A comprehensive description of different market types is provided in Harris (2003).

volatility and mean return. After that we present a direction-of-change technique based on logit model which does not use those forecasts.

The first of approaches to direction-of-change forecasts is the 'naive' mean forecast. To forecast mean return in this approach, we introduce seasonal dummy variables corresponding to each decision interval inside a week. This gives us 20 dummy variables[5] for Gazprom, and 60 dummy variables[6] for Euro / Japanese Yen market. Then we regress historical returns R_t on those dummies, and use estimated coefficients to make mean return forecast $\hat{\mu}_{t+1|t}$ for the next decision interval. In this way we account for the fact that activity of traders shows significant intra-day and intra-week seasonality.

After that we proceed with volatility forecast. We calculate realized volatility over the historical period in the following way:

$$\sigma_{RV,t} = \sqrt{\sum_{i=1}^{m} r_{t-1+i\cdot\Delta t}^2}, \tag{1}$$

where $\sigma_{RV,t}$ — square root of realized volatility for the decision interval ending at time t, $r_{t-1+i\cdot\Delta t} = 100(\log P_{t-1+i\cdot\Delta t} - \log P_{t-1+(i-1)\cdot\Delta t})$ — return for the interval of time with length Δt which ends at moment of time $t-1+i\cdot\Delta t$, and Δt is the length of sampling interval for realized volatility calculations[7], m — number of sampling intervals inside one decision interval.

Subsequently we regress historical $\log \sigma_{RV,t}$ on seasonal dummies, calculate residuals and estimate coefficients for an ARMA model[8] of residuals. Those coefficients are used to make a one-step-ahead forecast, and after restoring seasonal component we get the volatility forecast for the next decision interval. Now we have all the components required to make direction-of-change forecast.

Some of the transaction costs are taken into account in the direction-of-change forecast. We include in our model only proportional transaction costs — that is costs, taken as part of the whole trade cost. An example of those costs is brokerage commission. For the sake of simplicity, we assume the bid-ask spread to price ratio is constant[9]. We do not include fixed transaction costs, like cost of opening a trading account or fixed monthly fee, because they are either negligible or absent.

[5] 4 decision intervals inside a day × 5 trading days in a week = 20, see Section 3.2 for discussion regarding the decision interval size selection

[6] 12 decision intervals inside a day × 5 trading days in a week = 60

[7] See Section 3.2 for discussion regarding the sampling interval selection.

[8] See Section 3.4 regarding the model selection for residuals.

[9] See Section 3.5 regarding the sample bid-ask spread and transaction costs study.

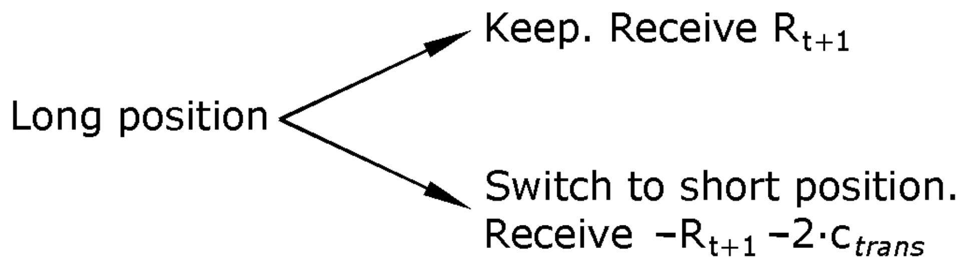

Figure 1: Choice faced when holding a long position in asset.

We account for transaction costs in the following way (see figure 1): suppose the 'virtual' researcher at time t holds a long position in an asset, and has two opportunities: either to keep it, or sell it and take a short position. Then to make a decision to switch position, the researcher has to forecast the probability

$$Pr(-R_{t+1} - 2 \cdot c_{trans} > R_{t+1}|\Omega_t) = Pr(R_{t+1} < -c_{trans}|\Omega_t),$$

where c_{trans} is the proportional transaction cost, Ω_t — set of information available at time t. In the same way, when holding a short position, the researcher forecasts the probability $Pr(R_{t+1} > c_{trans}|\Omega_t)$.

In this paper we follow the approach of Christoffersen and Diebold (2006) to direction-of-change predictability. Specifically, for 'naive' mean forecasts we use

$$\widehat{Pr}(R_{t+1} > c_{trans}|\Omega_t) = 1 - \widehat{F}\left(\frac{c_{trans} - \hat{\mu}_{t+1|t}}{\hat{\sigma}_{t+1|t}}\right), \tag{2}$$

where $\widehat{F}(\cdot)$ is an empirical cumulative density function of $\frac{R_k - \hat{\mu}_{k|k-1}}{\sigma_{RV,k}}$, $k = 1, ..., t$.

The second approach to direction-of-change forecast is 'standardized' mean forecast. We calculate standardized returns for the whole sample of available information:

$$R_{stand,t} = \frac{R_t}{\sigma_{RV,t}}, \tag{3}$$

Then we regress historical standardized returns $R_{stand,t}$ on seasonal dummies, estimate coefficients and use them to forecast $\hat{R}_{stand,t+1|t}$. Next, $\hat{\mu}_{t+1|t}$ is calculated in the following way:

$$\hat{\mu}_{t+1|t} = \hat{R}_{stand,t+1|t} \cdot \hat{\sigma}_{t+1|t}, \tag{4}$$

where $\hat{\sigma}_{t+1|t}$ is the one-step ahead volatility forecast for the decision interval ending at $t + 1$. The volatility and direction-of-change forecasts are made in the same way as for 'naive' mean forecast (see formula (2)).

The third approach to direction-of-change forecast is the 'Gram-Charlier expansion'. Ideas of this approach track back to the works of Gram and Charlier written in the end of XIX'th century. It involves usage of linear sum of parent function (in our case – cumulative density function of normal distribution) and its successive derivatives to build an approximate representation of a given probability distribution (see Samuelson (1943) for discussion of Gram-Charlier series). Application of this technique to direction-of-change forecasts is described in great details in Christoffersen, Diebold, Mariano, Tay and Tse (2006).

The mean return forecast here includes regression of R_t on $\log(\sigma_{RV,t})$, $[\log(\sigma_{RV,t})]^2$ and seasonal dummy variables (this differs from the approach of Christoffersen, Diebold, Mariano, Tay and Tse (2006) as they use only a constant as a deterministic regressor). Estimated coefficients are used to make mean return forecast $\hat{\mu}_{t+1|t}$. This form of dependence was chosen because the quadratic term in this regression is significant for the starting estimation sample in the Euro / Japanese Yen data[10].

For the direction-of-change forecast the following form is used:

$$\begin{aligned}\widehat{Pr}(R_{t+1} > c_{trans}|\Omega_t) &= 1 - F\left(\frac{c_{trans} - \hat{\mu}_{t+1|t}}{\hat{\sigma}_{t+1|t}}\right) \\ &\approx 1 - \Phi\left(\frac{c_{trans} - \hat{\mu}_{t+1|t}}{\hat{\sigma}_{t+1|t}}\right)(\hat{\gamma}_0 + \hat{\gamma}_1/\hat{\sigma}_{t+1|t}) \end{aligned} \tag{5}$$

Parameters γ_0 and γ_1 are estimated by regressing $1 - I(R_k > c_{trans})$ on $\Phi\left(\frac{c_{trans} - \hat{\mu}_{k|k-1}}{\sigma_{RV,k}}\right)$ and $\Phi\left(\frac{c_{trans} - \hat{\mu}_{k|k-1}}{\sigma_{RV,k}}\right)/\sigma_{RV,k}$ for $k = 1, ..., t$. Here $I(\cdot)$ is an indicator function, $\Phi(\cdot)$ — cumulative density function of standard normal distribution.

The fourth approach to direction-of-change forecasting, which we will refer to as the 'dynamic logit model', follows Anatolyev and Gospodinov (2007). We parameterize $Pr(R_{t+1} > c_{trans}|\Omega_t)$ using the dynamic logit model

$$Pr(R_{t+1} > c_{trans}|\Omega_t) = \frac{\exp(\theta_{t+1})}{1 + \exp(\theta_{t+1})} \tag{6}$$

where

$$\theta_{t+1} = d'_{t+1}\omega + \beta I[R_t > c_{trans}] + x'_t\delta \tag{7}$$

The model (7) contains lagged value of indicator function as a regressor. In addition, d_t contains seasonal dummies and x_t includes lagged bipower variation (BPV, see Barndorff-Nielsen and Shephard (2004)), square root of realized volatility (σ_{RV}, see Andersen, Bollerslev,

[10]The same reason for this form of dependence is provided in the original paper of Christoffersen, Diebold, Mariano, Tay and Tse (2006).

Diebold and Labys (2000)), third and fourth moments (RS and RK, accordingly), defined in the following way:

$$BPV_t = \frac{\pi}{2}\frac{m}{m-1}\sum_{i=1}^{m-1}|r_{t-1+i\cdot\Delta t}||r_{t-1+(i+1)\cdot\Delta t}|$$

$$RS_t = \sum_{i=1}^{m} r^3_{t-1+i\cdot\Delta t}$$

$$RK_t = \sum_{i=1}^{m} r^4_{t-1+i\cdot\Delta t}$$

BPV is included in addition to realized volatility because bipower variation is unaffected by presence of jumps, while realized volatility is an estimator of integrated variance plus a jump component (Barndorff-Nielsen and Shephard (2004)).

Equations used to forecast $Pr(R_{T+1} < -c_{trans})$ are defined analogously.

3 Empirical Results

3.1 Data and software

The database for Gasprom, the world's largest gas company[11], was downloaded from the website www.finam.ru[12]. The chosen sample covers the period from 23 January 2006 to 19 December 2006, composed of 230 trading days. Only the 'normal' trading regime is studied in this paper. It lasts from 10:30 in the morning to 18:45 in the evening (Moscow time — GMT+3), without stopping for lunch break. Nevertheless, the source data contains information on trades that occurred outside of the boundaries of the 'regular' trading regime. No corrections were made in this paper to account for the overnight volatility, which may improve the empirical results.

MICEx, the stock exchange where Gazprom is traded, is an continuous order-driven market. Every market participant is able to post two types of orders: a limit order or a market order. A limit order contains information regarding the amount of shares that the market participant wants to sell or buy, and the price. A market order specifies the amount of shares the participant wants to sell or buy immediately (sometimes, maximum or minimum price which is not to be exceeded may be included in the order). When a new market order arrives to the stock exchange, it is matched with the best possible limit orders presented in the system. The matching is made

[11] http://www.gazprom.com/eng/articles/article8511.shtml

[12] The precise address is http://www.finam.ru/analysis/export/default.asp for russian-language interface, http://www.fin-rus.com/analysis/export/default.asp for english-language interface

first according to the price set in limit orders, and then (if two limit orders set the same price) priority in execution receives the limit order that arrived first. If one limit order is not enough to fill the amount of shares requested in a market order, the matching process is continued. Every time a matching of orders occurs, the record about this transaction appears in a transaction database.

Each record in the database contains information about the time when this transaction occurred (with a precision of 1 second), price and volume information. After keeping only trades from the 'normal' regime, we aggregate sales according to Jasiak (1999), Anatolyev and Shakin (2007). However, a slight modification to their algorithm is proposed. In case of several transactions recorded at the same time, they leave only the first set with non-decreasing or non-increasing prices. All consequent information recorded at this second is discarded. In our empirical task only price and time information are required — as soon as the number of simultaneous transactions is very high, removing deals will result in loss of information, especially at high-volatility periods. Thus the transactions that occurred on the same second are aggregated into one transaction no matter which direction the price is moving during this second and whether or not it changes direction. This approach is used to smooth the price process for realized volatility estimation. After aggregation, the average number of transactions per day is 8800.

The database for Euro / Japanese Yen exchange rate was downloaded from the GAIN Capital web site[13]. The sample is from 3 January 2006 to 29 December 2006, from 00:00 Monday to 00:00 Saturday (GMT time), all day long. The FOREX market is a quote-driven dealer market. Dealers quote the prices at which they want to sell or buy currency and supply the liquidity for market participants. The database contains bid and ask quotes information as well as the time (with precision up to a second) when this quote was set by a dealer. The average number of quotes per day is 13000.

Main part of calculations was performed in Matlab using standard functions and Econometrics Toolbox by James P. LeSage (1999). To prepare database for analysis, Microsoft Excel with Visual Basic for Applications were used. Some supplementary tasks were done in software Ox (see Doornik (2002)) and Matrixer by Alexander Tsyplakov[14].

[13]http://ratedata.gaincapital.com/

[14]http://www.matrixer.narod.ru/

3.2 Selection of sampling frequency

First of all we need to choose the sampling frequency for realized volatility estimation and select the appropriate division of available data on parts for which the realized volatility is estimated. In this task we rely on the 'volatility signature plot' approach, introduced in Andersen, Bollerslev, Diebold and Labys (2000). Their insight was to build the graph representing a relation between average realized volatility calculated using sampling frequencies from 1 minute to 120 minutes inside the day, and sampling frequency, used for estimation.

To compute sampled returns for Gazprom we use the grid technique, proposed in Fleming, Kirby and Ostdiek (2003). First, for a given sampling frequency Δt seconds, we construct a grid of intervals with length Δt seconds that spans the trading day. Then we identify the nearest transactions that occurred before or after each of the grid point. If the point of the grid coincides with a transaction, we take the price of this transaction as the price at this point. In other cases, we use linear interpolation to estimate the log price at each grid point.

In the Euro / Japanese Yen market we have the bid-ask spread, and, consequently, mid-price information (calculated as an average of bid and ask prices). Thus at every moment of time we know which price was available on the market, and as the price for a grid point we use the last observed price on the market before or at this grid point.

After that we take first differences of log prices to obtain the continuously compounded Δt-seconds returns. Those returns are used to calculate realized volatility $\sigma_{RV,t}$ as sum of squared Δt-seconds returns (formula (1)).

Applying the 'volatility signature plot' methodology for Gazprom and Euro / Japanese Yen data we get the following picture — figure 2. In choosing the sampling frequency we face a trade-off: if we take sampling interval too small, the impact from market microstructure noise will be high. If we take it too large, the relevant information from the price process will be lost. For Gazprom, on small sampling intervals on this graph (less than 150 seconds), we capture the bias from market microstructure noise, contained in the intra-day data. For intervals smaller that 30 seconds noise leads to overestimation of volatility, for intervals from 30 to 150 seconds — tends to underestimate it. This behavior of realized volatility is consistent with the same calculations, performed for the stock of Alcoa, Inc. by Phillips and Yu (2005, see figure 2), and for the stock of IBM by Bandi and Russell (2005, see figure 1). For Euro / Japanese Yen the graph steadily declines with the growth of sampling frequency.

According to the methodology we should choose the point where the average volatility stabilizes. On our graph it stabilises at roughly 200–300 seconds for Gazprom, at 200 seconds for Euro / Japanese Yen.

Having approximately chosen the sampling frequency we proceed with selecting the decision interval size. In Andersen, Bollerslev, Diebold and Labys (2000) and Fleming, Kirby and Ostdiek (2003) the number of grid points used for realized volatility calculation varies from around 30 (when using returns sampled at 15 minute to estimate daily volatility) to around 100 grid points per interval of time. With our 200–300 seconds interval between grid points (for Gazprom) this corresponds to 100 to 500 minutes per decision interval. However, the length of trading day is just 495 minutes (8 hours and 15 minutes). So we decided to use decision intervals with size of 124 minutes (4 decision intervals inside a day, the last one being 1 minute shorter than others) and sampling interval of 240 seconds. For Euro / Japanese Yen we use decision intervals with size of 2 hours (12 decision intervals inside a day) and sampling interval of 200 seconds.

Distributions of returns R_t and standardized returns $\frac{R_t}{\sigma_{RV,t}}$ are presented on figures 3 and 4. Standardized returns are much closer to the usual Gaussian distribution. Jarque-Bera test cannot reject the null hypothesis of normality for both Gazprom and Euro / Japanese Yen markets.

3.3 Seasonal data adjustment

Anatolyev and Shakin (2007) discuss the importance of seasonal adjustments in Russian stock market. Thus, mean values are estimated for all 5 days of week, for every part inside the day. During this process, we accounted for the large number of holidays in Russia. As a result, some working weeks last less than 5 days, and some of them start not on Monday. This was corrected manually - a special table marking which day corresponds to which day of week was organized, to properly work with such things as 'first working day of week'.

Seasonal mean for logarithms of square root of realized volatility ($\log(\sigma_{RV,t})$) were calculated and subtracted from the realized volatility. The seasonal component is presented on figure 5. As we see, for Gazprom the price process is most volatile during first part each day, and the volatility decreases in the middle of trading day. Jarque-Bera test rejects the normality hypothesis for the residuals with p-value of 0.

To construct realized volatility estimates for $t+1$, we estimate seasonal component in $\log(\sigma_{RV,k})$ for $k = 1..t$. Then we remove seasonal component from $\log(\sigma_{RV,k})$ and use residuals to make a one-step-ahead forecast. After this we restore seasonal component in our forecast.

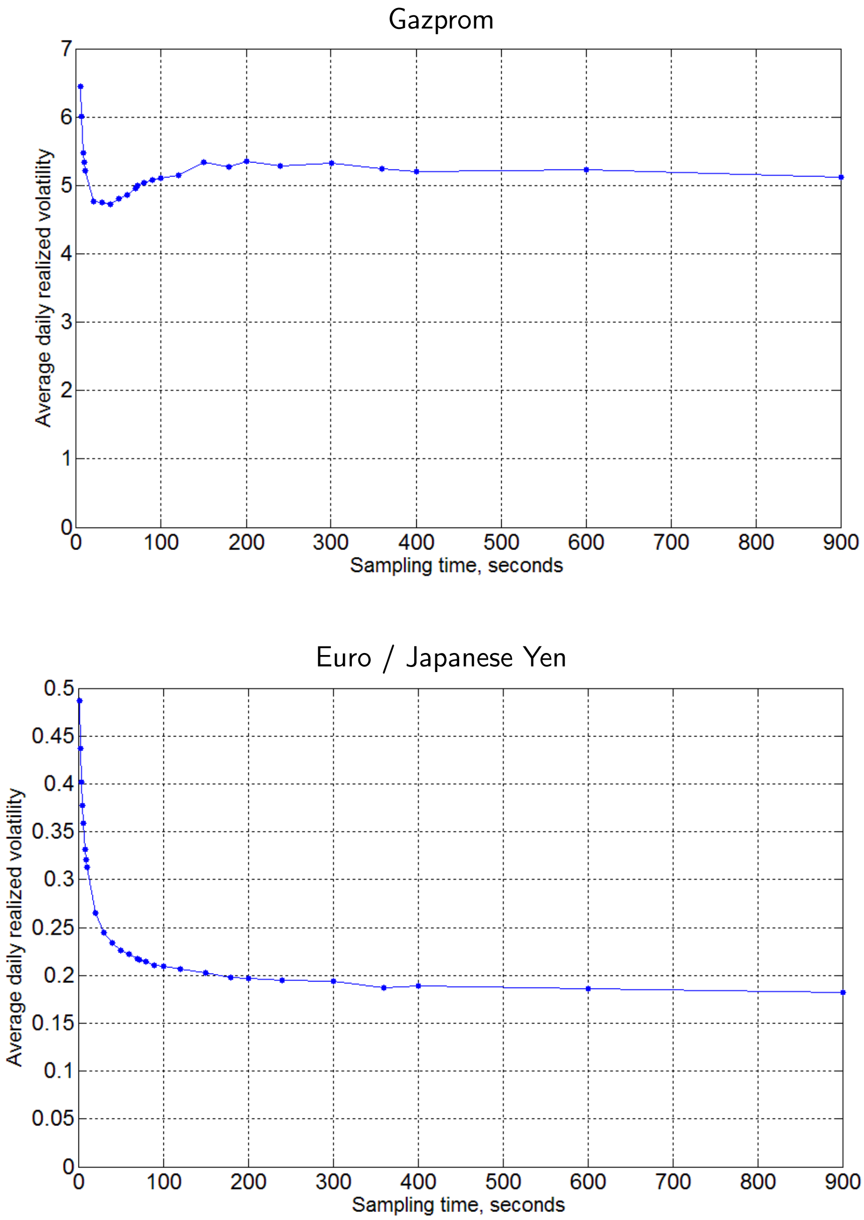

Figure 2: Average realized volatility for daily intervals with different sampling time for Gazprom and Euro / Japanese Yen markets.

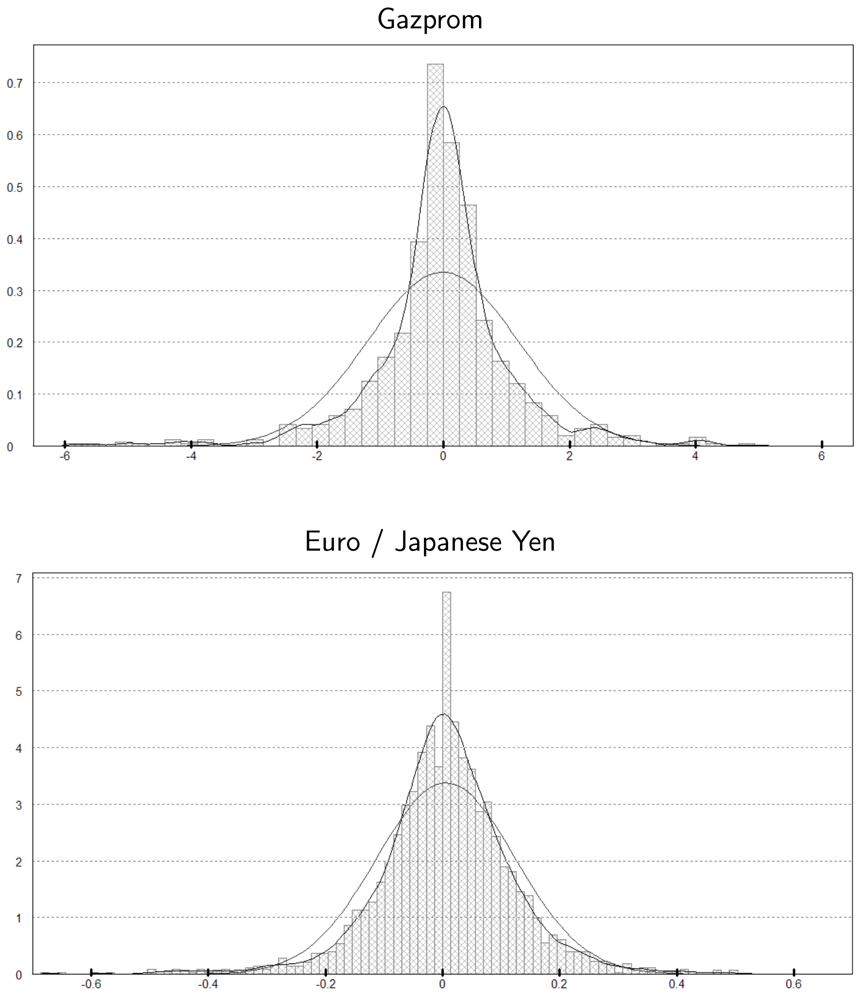

Figure 3: Histogram and kernel estimate for the returns sampled at 124 minutes (Gazprom) and 120 minutes (Euro / Japanese Yen). Gaussian probability density function with the same mean and variance is plotted for comparison.

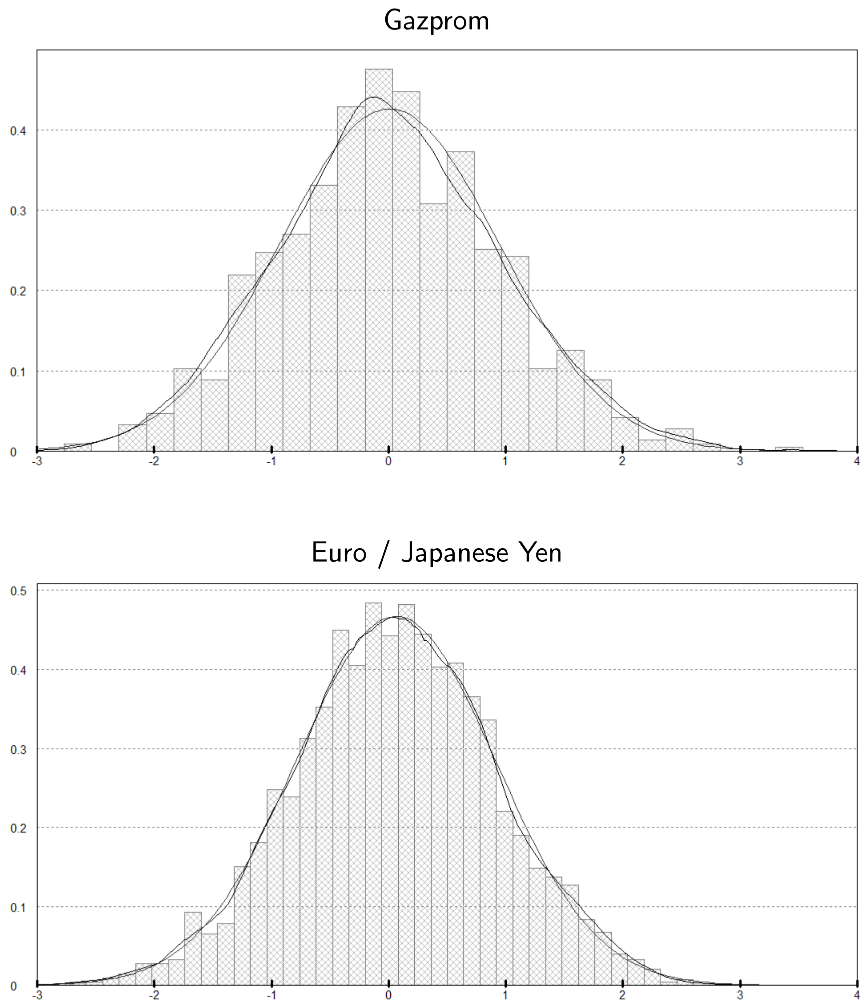

Figure 4: Histogram and kernel estimate for the standardized returns $\frac{R_t}{\sigma_{RV,t}}$ sampled at 124 minutes (Gazprom) and 120 minutes (Euro / Japanese Yen). Gaussian probability density function with the same mean and variance is plotted for comparison.

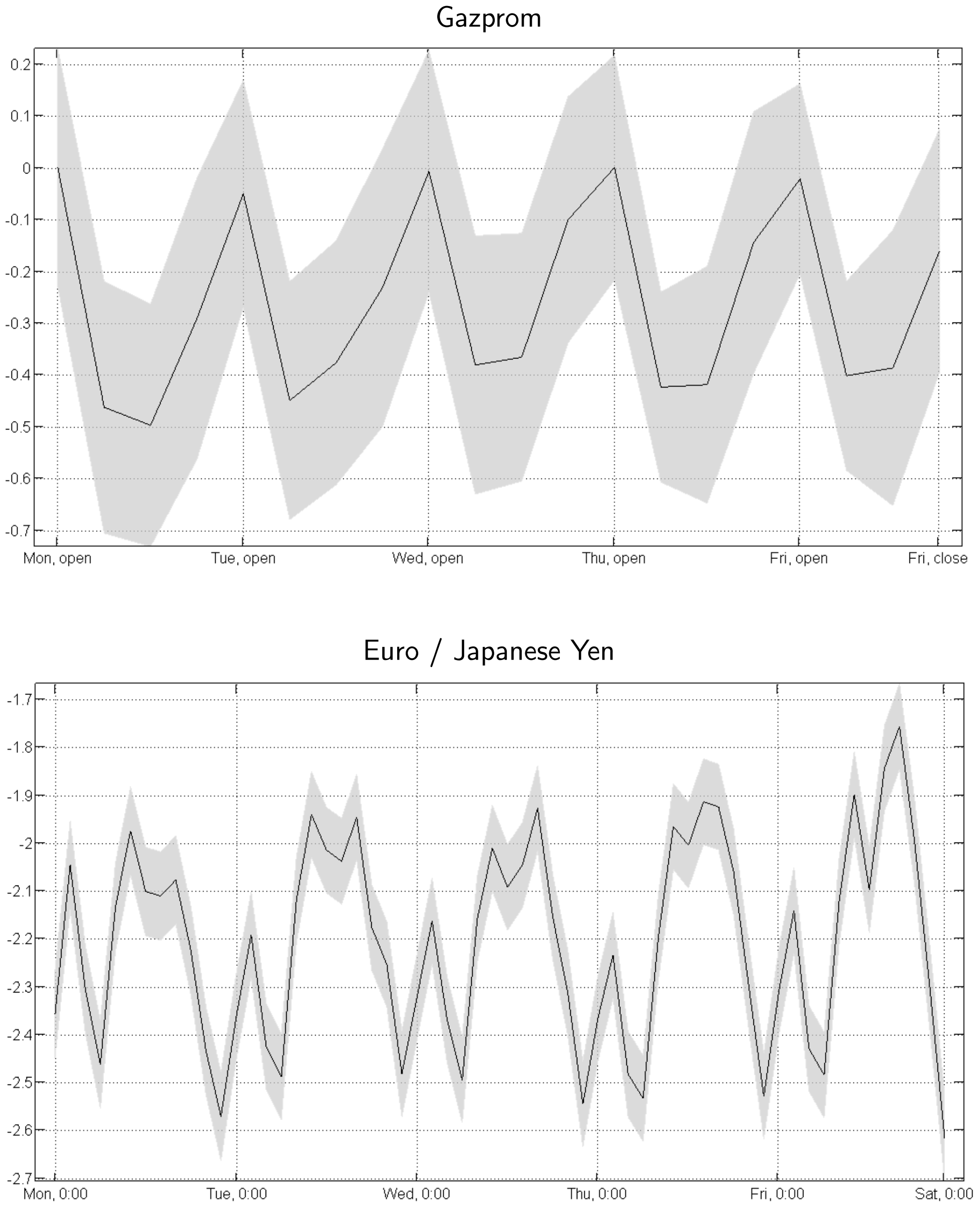

Figure 5: Seasonal component in the log realized volatility with 95% confidence band.

3.4 Realized volatility forecasting

The whole sample of realized volatilities consists of 920 observations for Gazprom[15] and 3086 observations for Euro / Japanese Yen[16]. On figure 6 the autocerrelation function for realized volatility with removed seasonality is presented. Series are highly autocorrelated. In addition, Andersen, Bollerslev, Diebold and Labys (2000) mention that realized volatility series show long-memory characteristics - consequently, they recommend to allow for fractional integration in the specification and use ARFIMA models.

For realized volatility forecasts, we proceed in the following way. Every time the current decision interval is finished, our 'virtual' researcher estimates seasonal component in $\log \sigma_{RV,t}$ and removes it. After that she chooses an ARMA(p,q) model for residuals by minimization of BIC on the grid $(p, q) \in [0 : 3, 0 : 3]$. Most of the time (except for a few days) this minimization selects ARMA(2,2) model for both markets[17].

After the parameters of model are estimated for the current period, the researcher makes a one-step-ahead prediction for realized volatility, restoring seasonal component.

ARMA models were estimated using MATLAB GARCH toolbox, ARFIMA models were estimated using Ox version 4.00 (see Doornik (2002)) and the Arfima package version 1.00 (Doornik and Ooms (2003)).

3.5 Spreads and other transaction costs

As noted in Section 3.1, the Gazprom database contain only transaction prices, without noting whether it is a buyer or seller initiated trade. One way to restore a proxy of bid-ask spread[18] is to take the minimum and maximum prices observed on the market in a a few minutes window centered on the point of interest. There is a tradeoff in this window size selection: if it is too small, then it may contain only (for example) buyer-initiated transactions, thus underestimating the spread. On the other hand, if it is too large, the spread will be overestimated because it will take price movements into account. We choose to present spreads distributions for three window sizes: 1, 2 and 3 minutes. For Euro / Japanese Yen market the spreads are directly

[15]There are 230 trading days from 1 January 2006 to 19 December 2006. 230 days $*$ 4 decision intervals a day = 920 observations.

[16]257 trading days from 1 January 2006 to 29 December 2006. 257 days $*$ 12 decision intervals a day = 3084 observations.

[17]We considered ARFIMA models as an alternative for ARMA models. However, the difference in forecasts of the models become insignificant when seasonality is added.

[18]Author thanks Alexander Gerko for this idea.

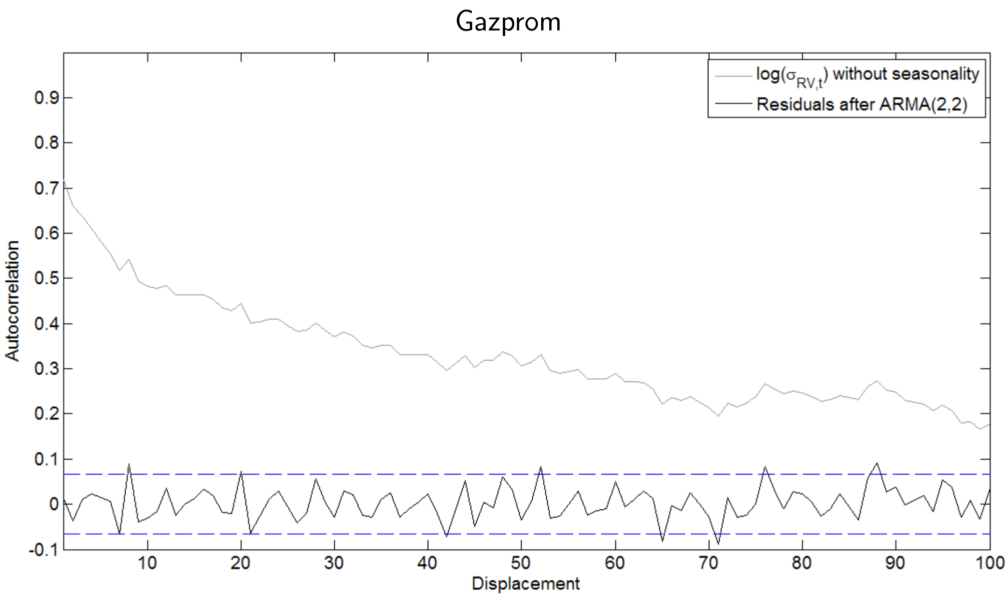

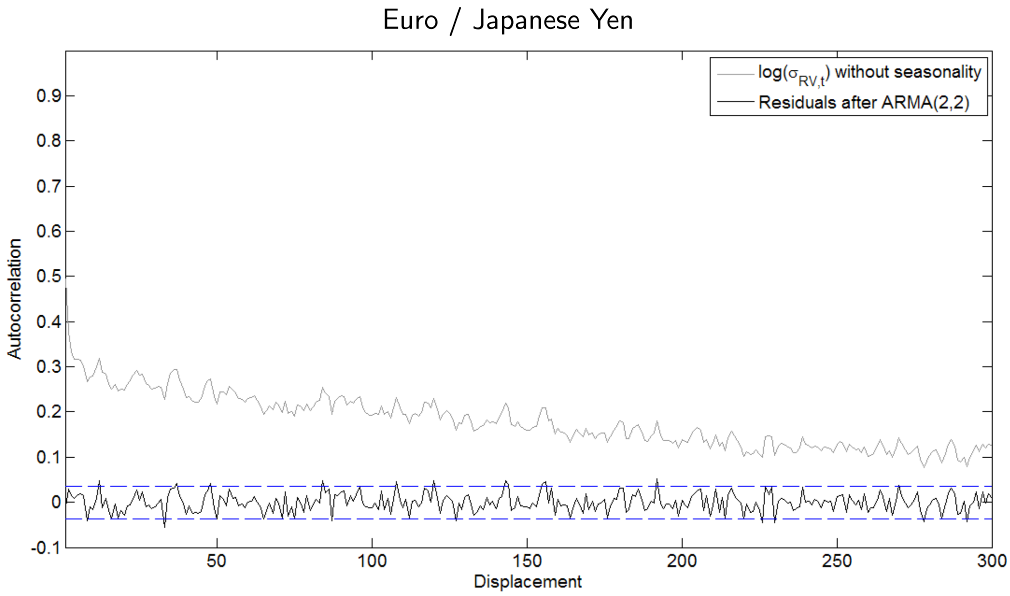

Figure 6: Autocorrelation function for log realized volatility with removed seasonality and residuals of ARMA(2,2) model. Two horizontal lines - 95% confidence interval bands.

available from our database. The spreads distribution for the subsample where the profitability will be estimated is presented in Table 1. The mean spread gradually increases as the window size grow.

We calculate the spread to price ratio (Table 2) to estimate proportional costs a trader incurs because of bid-ask spread. We choose the average transaction cost from bid-ask spread as 0.002 for Gazprom, 0.0003 for Euro / Japanese Yen.

Table 1. Spreads distribution.

	Min	Max	Mean	Median	Std. Dev.	Skewness	Kurtosis
				Gazprom[a]			
1 minute window	0.01	2.95	0.37	0.26	0.36	2.75	14.2
2 minute window	0.02	3.51	0.53	0.37	0.45	2.51	12.2
3 minute window	0.06	3.51	0.61	0.47	0.48	2.34	11.0
				Euro / Japanese Yen[b]			
	0.02	0.09	0.04	0.04	0.0022	8.67	266

[a]The spreads cover September 1, 2006 through December 19, 2006.

[b]The spreads cover September 1, 2006 through December 29, 2006.

Table 2. Spread to price ratio distribution.

	Min	Max	Mean	Median	Std. Dev.	Skewness	Kurtosis
				Gazprom[a]			
1 minute window	0.0000	0.010	0.0013	0.0009	0.0012	2.76	14.0
2 minute window	0.0001	0.012	0.0018	0.0013	0.0016	2.54	12.1
3 minute window	0.0002	0.012	0.0021	0.0016	0.0016	2.37	10.9
				Euro / Japanese Yen[b]			
	0.0001	0.0006	0.0003	0.0003	0.00002	7.67	228

[a]The sample covers September 1, 2006 through December 19, 2006.

[b]The sample covers September 1, 2006 through December 29, 2006.

The transaction cost incurred by trader consists of several components. First of them is brokerage commission. It is calculated as a percentage of the total transaction, and varies from 0.00% to 0.1% depending on the type of account a trader has with a broker and the volume of daily transactions[19]. Various MICEx commissions aggregate to 0.01%[20]. We choose 0.05% as

[19]See, for example, http://rustock.onlinebroker.ru/micex/tarif.asp or http://www.troika.ru/rus/Capital_Markets/Inter_tr/tariffs/index.wbp.

[20]http://www.micex.ru/stock/fees/

total transaction costs in excess of bid-ask spread. This corresponds to the 'Universal' account on the VTB24 OnlineBroker internet-trading service[21] with daily turnover of more than 300 000 rubles (approximately 8 700 euros as of September 1, 2006). As total transaction costs in direction-of-change forecasts we will use 0.25% (brokerage commissions plus average bid-ask spread in terms of price).

Operations on the FOREX market are free of brokerage commissions[22]. So as transaction costs for it we will use 0.03%.

3.6 Direction-of-change predictions and trading strategy valuation

To make one step ahead direction-of-change forecasts we use techniques described above in this paper. When the next decision interval is finished, 'virtual' researcher proceeds as follows. First of all, she makes a volatility forecast as described in Section 3.4. Then she makes direction-of-change forecast according to descriptions in Section 2. Then, if at time t she is holding a long position in an asset, and $\widehat{Pr}(R_{t+1} < -c_{trans}|\Omega_t) > 0.5$, she sells her assets and takes a short position. If she is holding a short position in an asset at time t, and $\widehat{Pr}(R_{t+1} > c_{trans}|\Omega_t) > 0.5$, she closes short position and takes a long one. Otherwise no changes to asset position is made. After that, researcher waits until the next decision interval is finished, and repeats the process.

To calculate the cumulative payoff we assume that the researcher reinvests all her initial wealth and all the money she received during previous trades. We estimate profitability of trading strategy under under zero brokerage commission and actual spreads for Euro / Japanese Yen market. For Gazprom market we check profitability under 0.05% brokerage commissions and spreads restored from transaction prices (we use windows with length of 1, 2 and 3 minutes). The results of those trading strategies are presented in Table 3. We include profits from a simple buy and hold strategy for comparison. The best results in both asset classes is shown by 'Gram-Charlier expansion' approach to direction-of-change forecast.

[21] http://www.onlinebroker.ru/

[22] http://itrading.onlinebroker.ru/olb/tarif.asp

Table 3. Trading strategies profitability for September – December 2006,

Gazprom

	Approach to direction-of-change forecast				
	'Naive'	'Standard.'	'Gram-Charlier'	'Dynamic logit'	Buy and hold
Number of operations	109	105	125	130	1
Profit (1 minute spr.)	2.0%	0.2%	4.7%	-30.4%	-3.8%
2 minute spread	-5.6%	-7.6%	-4.8%	-36.5%	-3.8%
3 minute spread	-9.2%	-10.9%	-8.6%	-38.8%	-3.8%

Euro / Japanese Yen

	Approach to direction-of-change forecast				
	'Naive'	'Standard.'	'Gram-Charlier'	'Dynamic logit'	Buy and hold
Number of operations	476	496	454	570	1
Profit	-4.0%	-7.3%	-3.9%	-16.9%	4.5%

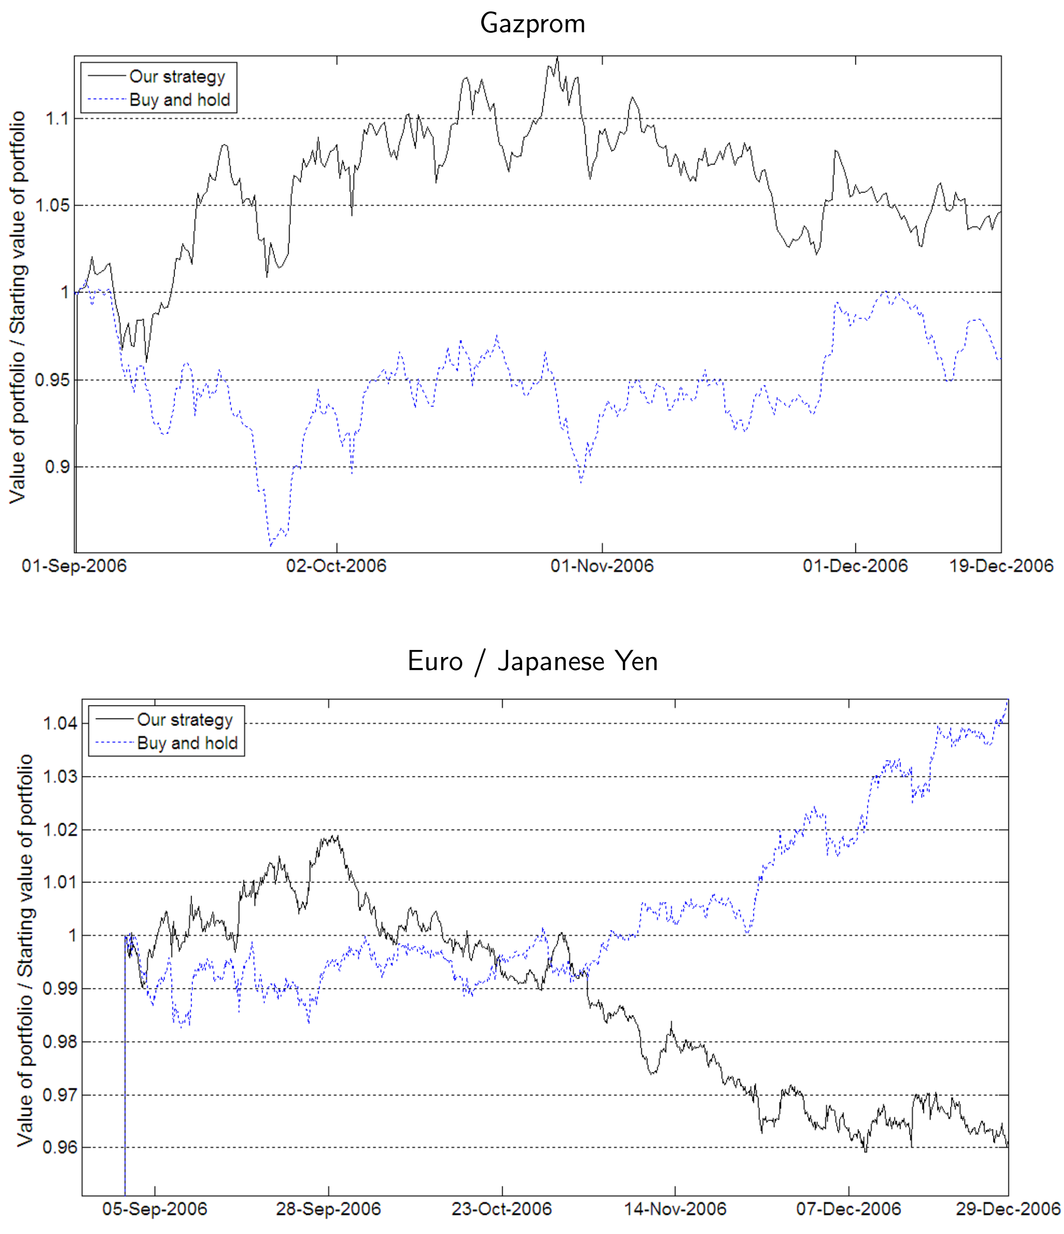

Figure 7: Value of portfolio (as share of initial value) under the trading strategy which uses 'Gram-Charlier expansion' approach to direction-of-change forecast. Value of portfolio under a simple buy and hold strategy is presented for comparison.

4 Conclusion and directions for future research

As we have shown, using ultra high frequency information in liquid markets allows to build profitable trading strategies operating on short horizons during the day. However, the profitability of those strategies highly depend on the transaction costs. The positive result for Gazprom stock market (4.7% for 4 months) corresponds to some market inefficiency which we were able to capture. Nevertheless, even this positive result disappears if we take less optimistic spread estimates. The negative result for Euro / Japanese Yen market, despite the absence of transaction costs, occurred probably because of market efficiency.

In future work, we look forward to building a trading strategy accounting for limitations on short positions and leverage. In addition, the impact of overnight volatility, which can be somehow measured by comparing opening and closing prises, may contain additional relevant information.

5 Acknowledgements

Author thanks Stanislav Anatolyev, his project leader, for advices and guidance during creation of this work. Author thanks Alexey Belkin, Viktor Kitov and Viktoria Stepanova for useful comments regarding ideas of this work. Special thanks go to Sergey Nikonov, the person with an online trading terminal on his notebook for insights regarding the stock market operations. Author thanks Alexander Gerko for comments regarding spread estimation on Russian markets. Author emphasizes that all errors remain his alone.

References

Anatolyev S. and N. Gospodinov (2007) Modeling Financial Return Dynamics by Decomposition, Centre for Economic and Financial Research at New Economic School Working Paper No 95, available at http://www.cefir.ru/papers/WP95Anatolyev.pdf.

Anatolyev S. and D. Shakin (2007) Trade Intensity in the Russian Stock Market: Dynamics, Distribution and Determinants, *Applied Financial Economics*, **17**, 87–104.

Andersen T.G., T. Bollerslev, P.F. Christoffersen and F.X. Diebold (2006) Volatility and Correlation Forecasting, in G. Elliott, C. Granger and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 777–878. Amsterdam: North-Holland.

Andersen T.G., T. Bollerslev, F.X. Diebold and P. Labys (2000) Great Realizations, *Risk Magazine*, **18**, 105–8.

Bandi F.M. and J.R. Russell (2005) Volatility, in J.R. Birge and V. Linetski (eds.), *Handbook of Financial Engineering*, Elsevier North-Holland, forthcoming.

Barndorff-Nielsen O.E. and N. Shephard (2004) Power and Bipower Variation with Stochastic Volatility and Jumps (with discussion), *Journal of Financial Econometrics*, **2**, 1–48.

Breen W., L. R. Glosten and R. Jagannathan (1989) Economic Significance of Predictable Variations in Stock Index Returns, *Journal of Finance*, **44**, 1177–89.

Christoffersen P.F. and F.X. Diebold (2006) Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics, *Management Science*, **52**, 1273–87.

Christoffersen P.F., F.X. Diebold, R.S. Mariano, A.S. Tay and Y.K. Tse (2006) Direction-of-Change Forecasts Based on Conditional Variance, Skewness and Kurtosis Dynamics: International Evidence, PIER Working Paper No. 06-016. Available at SSRN: http://ssrn.com/abstract=908317.

Doornik J.A. (2002) Object-Oriented Matrix Programming Using Ox, 3rd ed. London: Timberlake Consultants Press and Oxford: www.doornik.com.

Doornik J.A. and M. Ooms (2003) Computational Aspects of Maximum Likelihood Estimation of Autoregressive Fractionally Integrated Moving Average Models, *Computational Statistics and Data Analysis*, **41**, 333–48. Also see www.doornik.com.

Fleming J., C. Kirby and B. Ostdiek (2003) The Economic Value of Volatility Timing Using "Realized" Volatility, *Journal of Financial Economics*, **67**, 473–509.

Gençay R. (1998) Optimization of Technical Trading Strategies and the Profitability in Security Markets, *Econometric Letters*, **59**, 249–54.

Harris L., (2003) Trading and Exchanges, Oxford University Press.

Jasiak J. (1999) Persistence in Intertrade Durations, *Finance*, **19**, 166–95.

LeSage J.P. (1999) Applied Econometrics using MATLAB, unpublished manuscript, University of Toledo. Available at: www.spatial-econometrics.com/html/mbook.pdf.

Pesaran M.H. and A.G. Timmerman (2004) How Costly Is It to Ignore Breaks When Forecasting The Direction of a Time Series?, *International Journal of Forecasting*, **20**, 411–25.

Phillips P.B.C. and J. Yu (2005) Comment on "Realized Variance and Market Microstructure Noise" by Peter R. Hansen and Asger Lunde, Singapore Management University, School of Economics and Social Sciences, Working Papers 13-2005.

Samuelson P.A. (1943) Fitting General Gram-Charlier Series, *The Annals of Mathematical Statistics*, **14**, 179–87.

White H. (2000) A Reality Check for Data Snooping, *Econometrica*, **68**, 1097–126.