



**DOCTORAL SCHOOL OF
ECONOMICS**

RÉSUMÉ
for

Balázs Árpád Szűcs

Intra-Day Forecasting of Stock Volumes

Ph. D. dissertation

Supervisor:

Tamás Makara, Ph.D.
associate professor

Budapest, 2015

Department of Finance

RÉSUMÉ
for

Balázs Árpád Szűcs

Intra-Day Forecasting of Stock Volumes

Ph. D. dissertation

Supervisor:

Tamás Makara, Ph.D.
associate professor

© Balázs Árpád Szűcs

Contents

| | |
|--|-----------|
| 1 Preliminaries | 2 |
| 2 Structure of the thesis, data and methodology | 5 |
| 2.1 Structure of the thesis | 5 |
| 2.2 The data | 5 |
| 2.3 Applied methodology | 6 |
| 3 Main results | 10 |
| 3.1 Identification of the best model in the literature | 10 |
| 3.2 Own models | 12 |
| 3.3 Overview | 18 |
| 4 Main references | 20 |
| 5 Own publications in the subject | 21 |

1 Preliminaries

Research covering stock exchanges usually focuses on the price, therefore much less attention is paid to the turnover. Consequently our knowledge on turnover is much more narrow. Theories modeling the price often ignore the turnover completely. The dominance of modeling price over turnover is probably rooted in the natural desire of investors trying to make money on the stock exchange. A better understanding, and ideally a better forecast of price directly contributes to this goal. It must be noted though that a better understanding, and preferably a better forecast of turnover also contributes to the wealth accumulation of investors.

Let us consider the fact that liquidity is always bounded, which means that above a certain threshold trades cannot be executed without additional costs. This holds to each and every market, only the level of this threshold varies. The reason for this is that an order with a size significantly above the average cannot fully be executed at the price that was observed before the submission. A large order will move the market against the submitter, making them buy at a higher, or sell at a lower price, compared to a much smaller order size. This phenomenon is called price effect, and it has the potential of causing substantial losses to the submitter. However, this price effect can effectively be reduced, or even evaded in the possession of a decent forecast of the turnover, in which case one is able to split their large order into smaller chunks that do not move the price. All practitioners pay close attention to this detail.

On the day with the highest turnover throughout September 2015, stocks with a total worth of 118 billion US dollars were exchanged on the NYSE. This number describes a single day of a sole exchange, whereas in comparison the GDP of Hungary in the year of 2014 amounted for 137 billion US dollars. Assessing the amount of money investors could make on this single day if they cared to split their orders, even if we assume a modest 1% of evaded price effect, we may conclude that a good forecast of turnover can be converted to substantial wealth accumulation for individual investors.

Turnover can be however considered of high importance not only for individual investors, but also for the entire market as a whole. Let us consider market efficiency to see this. The more information is incorporated in the price, and the faster it happens, the more efficient the market is. The highest possible level of market efficiency is therefore desirable. Nevertheless, it must be noted that the only way information can possibly be incorporated into the price is through trading. The higher the intensity (turnover) of the trading, the better the price discovery may be. In other words,

information can only be incorporated into the price through the turnover.

The relevance of turnover forecasts can be recognized in this process as well. As mentioned above, market players split their orders in order to reduce price effect caused by bounded liquidity, and thus they are slowing the trading down. A better forecast of turnovers can diminish the uncertainty about the order size that can be submitted without triggering price effect, and therefore market players can submit larger chunks, which results in shorter execution times of large orders. This makes trading faster in overall, which may contribute to increasing market efficiency, and hence the better general functionality of markets.

Finally, there is a third angle to the relevance of turnover forecasts, besides that of the individual investor and the market as a whole. This angle is related to initial public offerings (IPOs). IPOs clearly contribute to economic growth, the reason for which is that a newly listed company gains a new source of financing, which can be converted to growth of the company. A growing company is likely to create new jobs thus having a positive effect on the entire economy. One can observe that since the early 2000s the number of IPOs has fallen significantly in the United States, which means that fewer companies entered the stock exchange. The reason for this can be sought in the appearance of automated trading, that rendered some stocks extremely liquid, thus making all the rest become relatively less liquid. This results in higher expected returns of these less liquid companies, due to an increased illiquidity premium, which they are not always able to meet. This is exactly the case for newly listed companies, that are necessarily illiquid in the initial phase. Because of the higher expected returns caused by the illiquidity premium, many of them are doomed to fail, and consequently do not enter the stock exchange at all.

This liquidity premium appears in the returns due to the liquidity risk, namely that it is difficult to buy or sell large numbers of shares at good prices. This liquidity risk is a result of the price effect mentioned earlier. As previously explained, price effect can be effectively reduced by order splitting, that is executed based on turnover forecasts, which means that good turnover forecasts can also serve as a tool for risk management. Once liquidity risk is reduced, liquidity premium also diminishes, and this allows more companies to get newly listed.

All in all, a better forecast of stock volumes has positive effects on the wealth of individual investors, on the market itself as a whole, and also on the entire economy.

The topic of the dissertation is forecasting intra-day stock volumes, or equivalently,

turnovers which is simply the percentage form volume. This reasonably new field of research is still evolving. Only few publicly available studies have addressed the forecasting of exchange volumes so far. The first such article was published in 2007, but the data used in it was of daily frequency. The typical intra-day stylized facts of volumes however make it only an indirect antecedent. The first article about the intra-day forecasting of exchange volumes was published in 2008, followed by one in 2011.

The aim of the dissertation is twofold. On the one hand, it aims to review the literature of scientific results achieved in forecasting intra-day stock exchange volumes, including both theoretical and methodological aspects. On the other hand, after running the best methods on own data, it aims to develop new models that perform better than those found in the literature.

2 Structure of the thesis, data and methodology

2.1 Structure of the thesis

The thesis consists of five parts.

Part I. lays the foundations of the research. Chapter 2. presents the most important concepts, while reflecting on the relevance of the studied field. Chapter 3. contains the review of the volume forecasting literature.

Part II. serves as an introduction to the empirical research. Chapter 4. describes the data base at my disposal, while Chapter 5. contains the research questions and hypotheses to be examined later.

Part III. comprises the reproduction of models found in the literature using my own data base, in order to identify the best model available. Chapter 6. sets the common elements of all latter estimations. Chapter 7. and 8. include the estimation of the two relevant models of the literature. Finally, Chapter 9. sets the benchmark to be outperformed later on.

Part IV. contains attempts to set up well performing new models. Chapter 10. includes model propositions and their evaluations based on standard error measures. Chapter 11. evaluates the best ones of the new models according to further error measures, including those suggested in the literature. Chapter 12. reviews the results of the modeling attempts in part IV.

Part V. concludes the dissertation. Chapter 13. is a summary of the first three parts, Chapter 14. provides a short list of the main findings and results, while Chapter 15. suggests some further possible questions to examine in related research later.

2.2 The data

The data base contains stocks included in the Dow Jones Industrial Average (DJIA or Dow 30) index that covers significant companies listed on exchanges in the United States. The index has been computed since 1896. The actual shares included in it somewhat varied since the introduction of the index, which is why the database contains not 30, but 36 tickers. Most of them, namely 33 are listed on the NYSE, the remaining 3 are listed on NASDAQ. The date of the first data point is 02/01/1998, except for stocks that were introduced to the exchange later, in which case the date of the IPO is the first data point. The date of the last data point is 13/07/2012 uniformly for all tickers.

The sample remaining after the data cleaning process ranges from 10/10/2001 to 13/07/2012, a period that is 130 months, nearly 11 years long. The number of tickers remaining in the sample is 33. The original frequency of observations was 1 minute, but I aggregated the data into 15-minute bins, in order to comply with the literature. This resulted in 26 observations every day for each ticker. The stocks remaining in the sample were liquid enough, meaning that every stock had trades in every 15-minute interval, and thus a volume record larger than zero. The data base finally used for analysis thus contains 2.29 million observations.

2.3 Applied methodology

Among the methodologies applied in the dissertation, only those are highlighted here, which are considered new compared to the literature. These are primarily the ones related to the decomposition of the intra-day U shape of turnover.

2.3.1 Polynomial fitting

The U-method being practically an average naturally produces a noisy U shape. It would be worth examining a decomposition that is exempt from such noise, or in other words, which is smoothed.

A smooth U shape could be associated with a polynomial of degree 2, but let us not fix the degree in advance. Let p denote a polynomial of degree n :

$$p_t = \sum_{i=0}^n \beta_i x_t^i \quad (1)$$

where $t \in \{1, 2, \dots, T\}$ and $x_t = t/T$, while T denotes the number of bins in a day. That is, we are looking for $T = 26$ points each day that lay on a polynomial, and fit the $J = 20$ observed days to the highest possible extent. The following problem must be solved to find these p_t points:

$$\sum_{j=1}^J \sum_{t=1}^T (p_t - y_{j,t})^2 \rightarrow \min_{\beta_i} \quad (2)$$

where t is the index for bins, $j \in \{1, 2, \dots, J\}$ is the index for days, and y denotes the observed turnover data.

2.3.2 Exponentially weighted polynomial fitting

Theoretically speaking it seems fair to assume that more recent observations explain tomorrow's turnover better than older ones. On this note it is worth examining a modification of the model above, where larger weights are assigned to more recent observations. In order to make the weights substantially differ, I applied exponential weighting.

The polynomial fitting above changes only slightly, as an additional weight is added to problem (2). The objective function then becomes:

$$\sum_{j=1}^J \sum_{t=1}^T S_j \cdot (p_t - y_{j,t})^2 \rightarrow \min_{\beta_i} \quad (3)$$

where

$$S_j = e^{-\frac{j}{J} \cdot \ln(1/J)} \quad (4)$$

This means that the weights assigned to the errors increase exponentially as the days go by, and reach the maximal value on the day closest to the present.

2.3.3 Spline fitting

Although it is commonly referred to as U shape, the intra-day seasonality of turnover may show some variety compared to the regular letter U. Figure 1. illustrates this through polynomials of degree 14 fitted to different 20-day intervals of the turnover of Kraft Foods. Inc. Apart from a regular U shape, some days resemble more to the letters V, W or J. Furthermore, turnover increases after the first observation sometimes, and it may also fluctuate substantially throughout the day. The need for such flexibility might have contributed to the conclusion of selecting higher degrees for the fitted polynomials.

By contrast, while flexibility is an advantage, an excessively good fit (even higher degree for the polynomial) also made the forecast less accurate. The reason for this might be that in such cases the polynomial reflects fluctuation that is rather just noise, whereas polynomial fitting was primarily introduced to reduce the noise component of the forecasted U shape.

So, on the one hand, flexibility is desirable (see Figure 1.), but on the other hand, noise should be omitted. Both might be achieved through the use of spline functions, that are commonly employed in yield curve fitting for similar motivations. The definition of spline functions is the following.




Figure 1: Polynomials of degree 14 fitted to different 20-day intervals of Kraft Foods Inc. turnover
Source: Own editing

Definition: Let K be a real number, and $t_1 < t_2 < \dots < t_K$ called knots. A spline function of degree N over a given set of knots is a $[t_1, t_K] \rightarrow \mathbb{R}$ continuous function, that has values of a polynomial of degree N between each pair of adjacent knots, and can be continuously differentiated $(N - 1)$ times.

The definition above is valid for $N \geq 2$. It could be extended to cover both $N = 1$ and $N = 0$, but it would not have much added value for this particular application.

With the use of spline functions it is possible to fit polynomials of lower degrees, which guarantees that noise will rather be omitted. But as the parameters of the polynomial vary from interval to interval, the fitted function may show much higher flexibility compared to a regular polynomial of the same degree.

The spline function is to be found in the following form:

$$p_t = \sum_{m=1}^M \beta_m f_m(x_t) \quad (5)$$

This resembles the polynomial fitting in (1), where $t \in \{1, 2, \dots, T\}$, which means that we are looking for $T = 26$ points. Furthermore, $x_t = t/T$. The only difference is that the f function is different, and it still needs to be specified.

The f functions (base functions) should be selected carefully in order to avoid mul-

ticollinearity. An appropriate solution for this is the use of B-spline base functions, the values of which are determined recursively based on the given knots and the degree.

The parameter estimation is performed similarly to the equally weighted polynomial fitting in (2).

The degree of the spline (N), as well as the K number and the exact placement of the knots is to be decided by the modeler. It can be proven that the necessary number of base functions, and thus the M number of parameters to be estimated can be determined as follows.

$$M = N + K - 1 \tag{6}$$

This can be applied to regular polynomials too. For instance, fitting a simple polynomial of degree 3 ($N = 3$) corresponds to 1 interval that has 2 endpoints ($K = 2$), therefore $M = 3 + 2 - 1 = 4$, which is clearly the number of parameters to be estimated, if we consider the constant as well.

Given that spline functions were introduced to be different from regular polynomials, it is reasonable to have at least two intervals, which means that $K \geq 3$.

3 Main results

3.1 Identification of the best model in the literature

The first research question was examined in part III. My aim was to compare the two relevant models of the literature (Bialkowski et al. (2008) and Brownlees et al. (2011)) using identical data and methods in order to find out which one performs better in forecasting intra-day stock turnovers. The winner should be considered as benchmark for my own models.

To this end, some decisions had to be made before starting the investigation. First, I had to determine how the error is measured, and what makes one model better than the other. At this point, I considered the MSE and MAPE error measures calculated for the forecasts, where a smaller value obviously signals better performance. I consider a model better, if it produces lower error measures for a higher number of shares than the other model, and if the average of the error measures for single stocks is smaller. This latter aspect helps to identify scenarios where there is no considerable difference between the forecasts of two models.

The data used for the analysis is the data described previously, that is 130 months for 33 tickers with 15-minute aggregation, which results in 26 daily observations for each stock. I use 20 days for parameter estimation, and the following day for forecasting and evaluation, which implies that parameters are refreshed daily. This results in 2648 forecasted days for each ticker. The information base is updated every 15 minutes.

After these preparations all was set to estimate the models of Bialkowski et al. (2008)¹ and Brownlees et al. (2011)². Estimation of the BDF model was straightforward based on the article, and according to the error measures described above, it clearly outperformed the U-method that is commonly used in practice.

Brownlees et al. (2011) however leaves the reader with some uncertainties regarding the estimation of their model. First, they do not specify the initial values of two variables during the recursion. Second, in the equation of the intra-day periodic component they mention that the number of terms is reduced from 25, but do not specify exactly how. Third, although this is merely a technical issue, it remains unknown how the starting value of θ is specified during the optimization, which turned out to be a key issue in

¹In brief: BDF model

²In brief: BCG model

the success of the estimation. As a result, I had to make my own assumption in these cases, which might only seem minor details, but they make the *perfect* reproduction of the article impossible. Finally, the authors did not provide the model specification that they actually estimated, but this latter was just a minor inconvenience, because they explained what modifications were needed.

Unfortunately the estimation itself was not without difficulties either. The objective function provided for the GMM estimation does not appear to be smooth enough to allow for finding an *acceptable solution* within *acceptable time*. By not acceptable time I mean that with the basic settings it took 60 days to produce the estimation³, which is much longer than the time needed for any other model I dealt with (the longest was 1 day). An acceptable solution would be a forecast the magnitude of which is comparable with the actuals observed later, which was the case for any other model I experimented with, but not the BCG model, not even after the lengthy estimation described above.

Finally, after some modifications in the course of the estimation I managed to produce acceptable forecasts from the BCG model. However, the need for these modifications makes it clear that the specification of the article, at least on my data set, is highly unstable, and therefore the success of the forecasts are rather incidental.

It must be emphasized, that Brownlees et al. (2011) used ETF data, not stock data, but it is doubtful in my opinion that this difference alone could result in such instability.

I compared the acceptable results obtained after the modifications with the U-method and also with the BDF model. Based on my data and the error measures presented previously, I found the forecast of the BCG model better than the forecast of the U-method, but worse than the forecast of the BDF model (especially the SETAR version).

Due to the instability of the BCG model described earlier, I found it unnecessary to compare it to the BDF model along further error measures. Thus an answer is found to the first research question.

Hypothesis H1: Benchmark. I accept the first hypothesis, namely that the BDF model is better than the BCG model, which means that the BDF model is the best one in the literature when it comes to forecasting intra-day stock turnovers.

³This is expressed in machine time, which shows the theoretical waiting time using a single computer with the average performance of the 6 computers I had at my disposal. The actual waiting time was shorter due to the use of several computers.

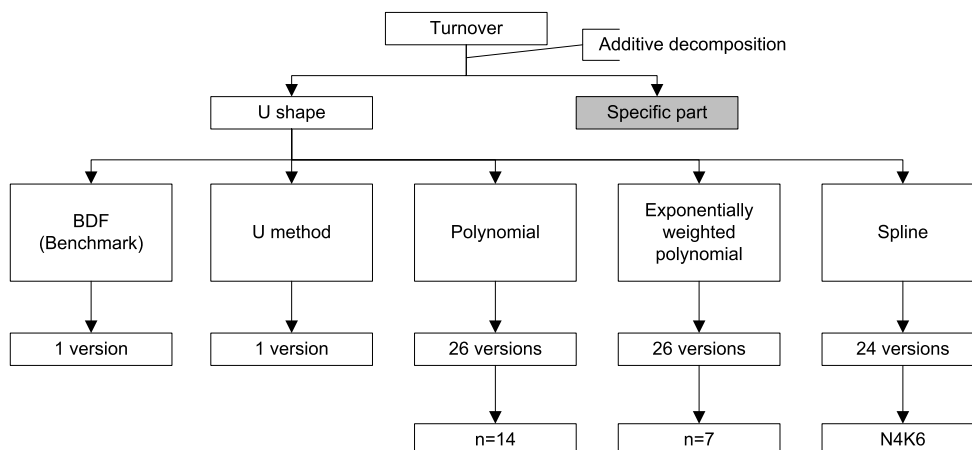


Figure 2: Different methods of decomposing the U shape
Source: Own editing

Consequently, the BDF model is hereinafter considered as the benchmark

3.2 Own models

The second research question and its extensions were examined in part IV. This part aims to find a new model that outperforms the benchmark based on the error measures discussed earlier.

The first step was to check a few models without U decomposition, but in accordance with my preliminary expectations these models did not perform well. The next step was to discover ways of decomposing the U shape that are different from what can be found in the literature. My first approach was similar to that of Bialkowski et al. (2008) in the sense that I also assumed an additive structure (U shape + specific part). I kept the specific part unchanged (and identical to the specific part of the BDF model) throughout part IV. for the sake of better comparability. It was thus possible to isolate the effects of different U decompositions. Figure 2. provides an overview of the examined model variations.

I started with extending the simple U-method with a specific part, and surprisingly this specification was already better than the benchmark. After this, I noticed that the U shapes of both the BDF model and the U method are rather *noisy*. I therefore tried to smooth the U shape, and found several ways to do it.

The first smoothing method was fitting a polynomial of degree n to the turnover series of each stock. Figure 3. provides an illustration of the U shapes of three different




Figure 3: The U-method, the U shape of the BDF model and the polynomial of degree 4 fitted to the first 20 days of Alcoa, Inc.
Source: Own editing

models (the arbitrarily selected degree for the polynomial is 4). The U shape of the BDF model reasonably differs from the other two, whereas the polynomial visually appears to be a smoothed version of the U method.

Given that my data base consists of 26 observations per ticker each day, I considered the n degrees between 1 and 26. I found that $n = 14$ appears to be performing the best, but any choice where $n \geq 7$ holds produced very similar results. The forecasts using polynomials of degree 14 to decompose the U shape performed better than the benchmark, so I tried to refine this method in the hope of even better performance.

The following model to decompose the U shape was fitting exponentially weighted polynomials. Compared to the equally weighted version above, the difference is that this approach assigns increasing weights to days that are closer to the present. I considered the n degrees between 1 and 26 again, and this time found that $n = 7$ performed the best, similarly adding that any choice with $n \geq 7$ produces similar forecasts. However, the use of lower ($n < 7$) degrees are not advised in neither the equally nor the exponentially weighted case. The exponentially weighted variation also performs better than the benchmark, but does not outperform the equally weighted version.

Finally, I tested the method of spline fitting that has some advantages over regular polynomials, therefore it might even improve the forecast. When working with splines,

the degree as well as the number of knots must be selected. In this case, I examined 24 variations and decided to use the degree of 4 and 6 knots (N4K6). This model also performs better than the benchmark, but the regular polynomial is still slightly better. The polynomial also being the simpler of the two, it remains the preferred U decomposition method.

In sum, I suggested new ways of decomposing the U shape, and each variation performed better than the benchmark. The best one of all was the polynomial fitting with a degree of 14, which is kept as benchmark for further experiments. In conclusion, the first extension of the second research question can now be answered.

Hypothesis H2.1: The U shape. According to my investigations, the modeling of the intra-day U shape of turnover does contribute successfully to the forecasting of intra-day turnovers. The model versions excluding U decomposition performed poorly, while the use of the new suggested models for U decomposition alone resulted in a models that outperform the benchmark (which was the BDF model). Each of the U decomposition methods I suggested (U-method, equally weighted polynomials, exponentially weighted polynomials, splines) result in better forecasts compared to the benchmark. The simple polynomial fitting was found to be the best one among all.

It follows from the logic of additive decomposition that after modeling the U shape it is also worth modeling the specific part. I therefore continued by keeping the best U decomposition method (namely the equally weighted polynomial fitting), and looking for different possibilities of extending it with specific part models. Figure 4. provides an overview of the model variations I considered.

There are two directions of my experiments with the specific part: one where there is some kind of price movement indicator included, and the other where there is no such element. So in the first one I experimented with including different price movement indicators besides the turnover data. The motivation behind this idea is based on the literature that reports some covariance between turnovers and prices. The indicators included were: log-return, volatility, gap, range, range percent, true range, true range percent.

I tested the usage of lagged price indicators in several ways, besides the simple lagged values I also tried conditional models based on correlation and Granger causality, but none of them produced favourable results.

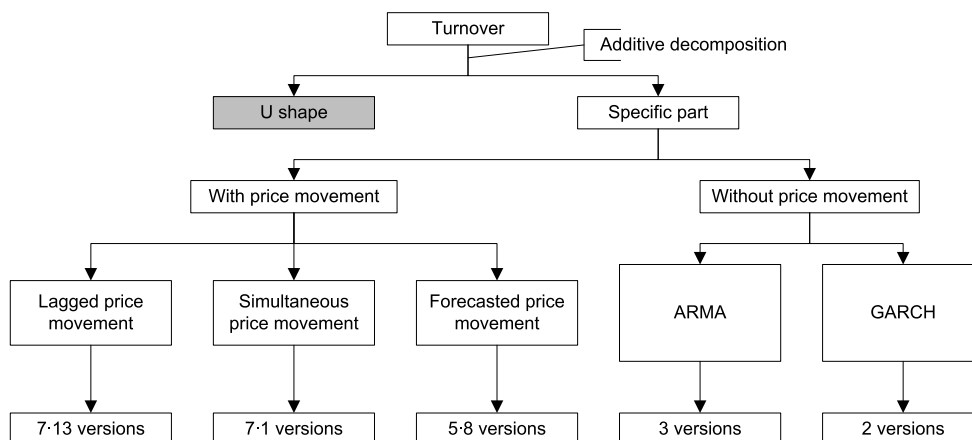


Figure 4: Different methods of forecasting the specific part
 U shape decomposed by the Poli(14) model
 Source: Own editing

After this, I tested a purely theoretical possibility, where the price indicator I use is simultaneous with the forecasted turnover. This solution cannot be implemented in practice, because these price indicators are still unknown when the forecast is carried out. However, this could still be helpful in deciding whether it is worth trying to forecast the price indicators or not. If the simultaneous price indicators do not deliver results, nor can it be expected from forecasted values. According to my findings, the simultaneous price indicators (excluding log-return and gap) contribute significantly to increasing the performance of turnover forecasts, which means that the following step should be the usage of forecasted price indicators (except for log-return and gap, of course).

When forecasting price indicators, I used relatively simple model specifications compared to the development of the field (i.e. forecasting volatility). These specifications however did not yield favourable results. Two remarks must be made here. First, the literature on volatility forecasting is so widespread (it is enough to think of stochastic volatility models or GARCH variants) that covering it is beyond the scope of this dissertation. Consequently, the above failure to successfully use price indicator forecasts does not mean this direction should be rejected. Second, in my opinion however, the usage of simultaneous price indicators can only be successful because it contains the shocks of information that cannot be forecasted by definition. This is the same information that moves the turnover itself, and that could not be foreseen from the lagged turnover values. If this is true, then no matter how developed price indicator forecasting method

we use, it still will be unable to forecast these shocks of information, and therefore will not be able contribute remarkably to forecasting turnover

Hypothesis H2.2: Price indicators. Based on my investigations I could not confirm that price indicators could successfully be used to improve intra-day turnover forecasts.

In the second part of my experiments with specific part models I tested some further specifications that do not include price indicators. I considered ARMA and GARCH variants, and concluded that the simple ARMA(1,1) is the best specific part model among the ones a tested.

Next, I examined some other models that do not fit into the previous additive logic. I started with models that purely use U decomposition, but no specific part. Then I moved on to an error correction model, and finally a multiplicative model that uses the best U decomposition and the best specific part model found so far. Among these, only the latter produced promising forecasts.

In the final chapter of part IV. I introduced some further aspects for evaluating models, and used them to compare the benchmark to my own model suggestions. For an overview of these aspects see Table 1.

In the course of testing models so far, I always produced one-step-ahead forecasts, and updated the information base in every step. The evaluation was done using the MSE and MAPE error measures applied to the deviation of forecasts and actuals. I monitored the number of shares one model showed lower error measures in, and also the average of error measures across shares. The latter helped to spot situations, where there was no significant difference between the errors of the models. The above was the first filter for models, and the following two specifications qualified for further evaluations:

1. Additive structure with polynomial U shape and ARMA specific part
2. Multiplicative structure with polynomial U shape and ARMA specific part

The second aspect considered was full day, that is 26-step (multiple-step-ahead) forecasts, evaluated similarly to the above.

Following Bialkowski et al. (2008), and given that turnover forecasts have an emphasized role in VWAP trading, I introduced two further aspects for evaluation. In both

| Chapter | What is forecasted | What is monitored | Error measure | Comparison |
|---------|--|---------------------|---------------|---------------|
| 10. | One-step ahead turnover values | Simple deviation | MSE | N° of tickers |
| | | | | Average |
| | | | MAPE | N° of tickers |
| | | | | Average |
| 11. | Multiple-step-ahead turnover values | Simple deviation | MSE | N° of tickers |
| | | | | Average |
| | | | MAPE | N° of tickers |
| | | | | Average |
| | Multiple-step-ahead turnover proportions (Static strategy) | Deviation from VWAP | MAPE | N° of tickers |
| | | | | Average |
| | One-step-ahead turnover proportions (Dynamic strategy) | Deviation from VWAP | MAPE | N° of tickers |
| | | | | Average |

Table 1: Different aspects of evaluating turnover forecasts
Source: Own editing

cases, intra-day turnover proportions are forecasted, and supposing that one would perform VWAP trading based on these forecasts, the deviation of actual VWAP and the average price reached on the trade are compared. This deviation is evaluated using the MAPE error measure. The MSE measure cannot be used this time, since the price component of the VWAP would result in significant biases.

The first strategy of focus is the static strategy, in which turnover proportions are forecasted until the end of the day in one step at the beginning of the day. This makes it a multiple-step-ahead forecast. This cannot be a realistic choice, and therefore its importance is only moderate. The second one is the dynamic strategy, where turnover proportions are forecasted only one step at a time (i.e. for 15 minutes), and the information base is updated as time progresses. Bialkowski et al. (2008) argues that the evaluation of such a dynamic strategy based on the VWAP is the ultimate measure for turnover forecasts.

The last coloumn of Table 1. shows that the two of the best own models were compared to the benchmark along 12 different aspects in overall. These aspects include the ones suggested by the benchmark article. We can conclude that both of my models outperform both versions of the benchmark along each of these 12 aspects. Among the

two own models, the multiplicative one could be considered superior, since it performed better in 11 of these 12 aspects compared to the additive one.⁴

Hypothesis H2: Better model. As a result of searching for models, I managed to find a specification that beats the benchmark selected from the literature according to all of the 12 aspects considered. In contrast to the benchmark, it follows a multiplicative logic. The U shape is decomposed in a new manner that cannot be found in the literature, namely by fitting a polynomial. Apart from the multiplicative structure this new decomposition method is the main innovation of my model. This decomposition results in forecasts that outperform the benchmark even if used in an additive context. The model for the specific part also differs from what can be found in the literature, but this difference is less significant compared to the previously mentioned ones.

The suggested multiplicative model beats the benchmark in each ticker under inspection, according to the aspects considered to be the most important by the benchmark article itself. Depending on the specific part of the two benchmark model variations, my multiplicative model presents a 13.6% to 61.9% improvement in average. This improvement was achieved by using 33 times less data compared to the benchmark model, as my model only needs the turnovers of the actual ticker as input, and not the entire market as a whole (which means 33 tickers in this case). The data base I used is over 9 times larger than that of the benchmark article.

3.3 Overview

Let us briefly overview of the main results of the dissertation.

1. Based on identical data (i.e. time period and assets), and using identical evaluation methods I compared the intra-day turnover forecasting models found in the literature in order to find out which one can be considered the best. According to my investigations the model of Bialkowski et al. (2008) is the best intra-day stock turnover (volume) forecasting model in the literature.
2. I suggested a new method (polynomial fitting) for decomposing the intra-day U shape of turnover. This method cannot be found in the literature, and it

⁴The sole exception being the average of MSE* values for the one-step-ahead turnover value forecasts, in which case the additive variant performed somewhat better.

outperforms the methods found in the literature when it comes to contribution to intra-day turnover forecasts.

3. According to my investigations, price indicators cannot contribute significantly to turnover forecasts, which contradicts my expectations based on the descriptive documentation of turnover found in the literature.
4. I suggested a new model that is substantially different from what can be found in the literature. This model is based on a multiplicative structure ($U_shape \cdot specific_part$), and uses the new U decomposition method described above. This model clearly outperforms the benchmark according to all of the 12 aspects considered, which also include those suggested by the benchmark article. Furthermore, my model requires significantly less amounts of input data, given that it only uses data of the actual ticker, unlike the benchmark model, which uses data of the entire market.

The data base at my disposal consists of 130 months and 33 tickers. This is significantly larger than the data bases previously used in the literature, which fosters the robustness of the results.

4 Main references

There are 89 literature references in the dissertation, among which I consider those that directly cover the forecasting of volume to be the most important ones. These are the following.

Bialkowski, J., Darolles, S., and Le Fol, G. (2006). How to reduce the risk of executing VWAP orders? - New approach to modeling intraday volume. Working paper.

Bialkowski, J., Darolles, S., and Le Fol, G. (2008). Improving VWAP strategies: A dynamic volume approach. *Journal of Banking & Finance*, 32:1709–1722. DOI: 10.1016/j.jbankfin.2007.09.023.

Brownlees, C. T., Cipollini, F., and Gallo, G. M. (2011). Intra-daily volume modeling and prediction for algorithmic trading. *Journal of Financial Econometrics*, 9:489–518. DOI: 10.1093/jjfinec/nbq024.

Kaastra, I. and Boyd, M. S. (1995). Forecasting futures trading volume using neural networks. *The Journal of Futures Markets*, Vol. 15, No. 8,:953–970. DOI: 10.1002/fut.3990150806.

Lux, T. and Kaizoji, T. (2004). Forecasting volatility and volume in the Tokyo stock market: The advantage of long memory models. Economics working paper, Christian-Albrechts-Universität Kiel, Department of Economics.

Lux, T. and Kaizoji, T. (2007). Forecasting volatility and volume in the Tokyo Stock Market: Long memory, fractality and regime switching. *Journal of Economic Dynamics & Control*, 31:1808–1843. DOI: 10.1016/j.jedc.2007.01.010.

5 Own publications in the subject

Szûcs, B. A. (2010). Liquidity in practice: Interview series with market players. In *Annual Financial Market Liquidity Conference, Budapest*.

Szûcs, B. A. (2012). A t zsdei forgalom napon bel li el rejelz se. In *Budapesti Corvinus Egyetem K zgazdas gi Doktori Iskola VIII.  ves Konferenci ja*.

Szûcs, B. A. (2013). Verseng  modellek a t zsdei forgalom napon bel li el rejelz s re. In *Budapesti Corvinus Egyetem K zgazdas gi Doktori Iskola IX.  ves Konferenci ja*.

Szûcs, B. A. and Ill s, F. (2015). *Mastering R for quantitative finance*, chapter 3. Forecasting volume, pages 59–75. PACKT Publishing, Birmingham, UK.

Szûcs, B. A. and V radi, K. (2014). Measuring and managing liquidity risk in the Hungarian practice. *Society and Economy*, Volume 36, Number 4:543–563. DOI: 10.1556/socec.36.2014.4.6.