NTNU
Norwegian University of
Science and Technology

Stock Return Prediction Using Artificial Neural Networks and Google Search Volumes

Rakel Håkegård
Cecilie Teisberg

Industrial Economics and Technology Management
Submission date: June 2018
Supervisor: Peter Molnar, IØT

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

Problem Description

Through decades of research, stock returns have proven to be difficult to predict. Complex relationships are evident in financial market data, and investigation of more sophisticated modelling techniques than traditional regression methods is therefore of interest. The aim of this thesis is to investigate the predictability of stock returns and examine whether artificial neural networks and Google search volume data can improve such predictions.

Preface

This Master's thesis examines predictability of stock returns using artificial neural networks and Google search volume data. The thesis concludes our Master of Science degree in Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU) in the spring of 2018.

We would like to thank our supervisor, Associate Professor Peter Molnár at the University of Stavanger Business School for helpful guidance and advice. We also extend gratitude to Katarina Lucivjanska at the Department of Economic and Financial Mathematics at the Pavol Jozef Safarik University in Kosice, Slovakia without whom we would not have been able to gather the data needed for this study in such short time.

Trondheim, 4 June 2018

Abstract

We investigate the predictability of abnormal stock returns using artificial neural networks, and examine whether Google search volume data can enhance such predictions. Our results show that the neural network models significantly outperform linear and semi-parametric methods for abnormal stock return prediction. We implement a trading strategy, buying the 50% of stocks with highest predicted abnormal return, and selling the 50% of stocks with lowest predicted abnormal return. The neural network trading strategy with a quarterly trading horizon has an average annual return that is 5 percentage points higher than the equally weighted portfolio, after accounting for trading costs. The performance is also improved in terms of portfolio volatility, and thus the risk-adjusted return is exceptional. We find that both the horizon of the input data and the prediction horizon impact the accuracy of predictions, with enhanced performance for longer horizons. Further, we examine the impact of Google search volume data on predictions, by comparing the neural network trading strategy with a benchmark excluding Google search volume data. We find that Google search volume has significant predictive power of abnormal stock returns for a quarterly trading horizon. Our results suggest that a decrease in Google searches over a period of one quarter is associated with significantly decreased abnormal return over the next quarter. However, for an increase in searches there is no such pronounced effect.

Contents

Chapter 1

Introduction

Prediction of stock returns has long been a topic of interest for academics and investors alike. If one could reliably predict stock returns, making money on the stock market would be a trivial task. Yet, beating the market has proven to be quite difficult. Although some stock fund managers have been able to achieve consistent excess returns over the market, the question of whether this has been a result of skill or sheer luck is highly disputed. The efficient market hypothesis, a cornerstone of modern financial theory, asserts that stock prices always reflect all relevant and currently available information, rendering any efforts to predict future stock prices based on the information available useless. However, there is general agreement in the academic community that the efficient market hypothesis in its most extreme form is not applicable in the real world.

Criticism of the efficient market hypothesis mainly involve the argument of imperfect and irrational human behaviour. Furthermore, even if completely rational behaviour was exhibited, humans do not have unlimited capacity to process all information that exist. Consequentially, stock market investors must allocate their attention to a limited number of stocks at any point in time, leading to a central question of how the level of investor attention impacts stock returns. However, first we must determine how to measure investor attention.

Technological advancements brought about in recent decades has made it possible for individual investors to easily access information on the Internet at their own discretion. Today, Google's search engine has the majority market share of Internet searches in the world, processing an average of 40,000 search queries every second. While traditional proxies for investor attention have been indirect measures, such as number of news articles, abnormal trading volume or extreme one-day returns, we now arguably have a more direct way to measure investor attention, namely Google search volumes. Moreover, data on search volumes is readily available to the public, in near real-time, through the web platform Google Trends, making stock return prediction based on Google search volumes an interesting topic for research.

Previous studies have seen conflicting results on the relationship between search volumes and stock returns, or lack thereof, and it is not clear from existing literature whether search volumes are significant in explaining returns. Linear methods have been extensively studied, however we find highly complex relationships in the data and argue that linear methods might not be powerful enough to harvest all information from Google search data.

We develop a neural network model for prediction of abnormal stock returns. We evaluate its predictive power, and compare it to that of both linear and semi-parametric models. To test its value in an economic sense we also design a trading strategy using the prediction models, and simulate trading over several years with the use of an online machine learning algorithm. The model is retrained after each trading period, incorporating new data as it becomes available, allowing the model to dynamically adapt to new information. We aim to shed new light on the question of whether Google search volume can significantly predict abnormal stock returns.

The rest of this report is structured as follows. In Chapter 2 we present a brief review of existing literature in the field, highlight potential problem areas and place our research into the body of literature. In Chapter 3 we give a description of the data and calculation of relevant variables, as well as discuss any limitations of the data. Chapter 4 gives a brief introduction to the methodology used and provides the model specifications. The results are presented and discussed in Chapter 5. Finally, Chapter 6 summarizes key findings and gives some recommendations for further work.

Chapter 2

Literature review

Several empirical studies in economic literature have investigated the impact of investor attention on stock returns. As information about investor attention is not straightforward to obtain, many different proxies have been studied. For example, Barber and Odean (2007) use news, abnormal trading volume and extreme one-day returns, Grullon et al. (2004) use advertising expenses and Seasholes and Wu (2007) use price limits as proxies for investor attention. While the traditional proxies mentioned are indirect measures of investor attention, Da et al. (2011) and Joseph et al. (2011) argue that the Google Search Volume Index (SVI) for company names or tickers is a direct measure. The SVI also captures investor attention in a more timely fashion than the traditional proxies, making stock return predictions based on SVI an interesting field of research.

Previous studies find that using SVI to predict volume or volatility is relatively easy, but that the correlation with future price returns is much weaker. Preis et al. (2010) find that SVI on company names is strongly correlated with weekly trading volumes of S&P500 companies. Aouadi et al. (2013) get similar results also in the French stock market, and find that the SVI serves as a decisive factor of the stock market liquidity and volatility as well. Takeda and Takumi (2014) examine the relationship between online search intensity and stock-trading behavior in the Japanese market, and find correlations with SVI on company names that are strongly positive for trading volume, but only weakly positive for stock returns.

The academic literature concerning investor attention and stock returns offers two hypothesis. On the one hand, the investor recognition hypothesis of Merton (1987) states that in a market with incomplete information, stocks with low investor attention will provide higher returns in order to compensate investors for idiosyncratic risk that cannot be diversified. Hence, assuming that the stock market is characterized by incomplete information, one would expect that an increase in investor attention will result in decreased stock returns in the long run. On the other hand, Barber and Odean (2007) state that an increase in investor attention may give rise to an increase in prices, and hence return, in the short run. This is based on the argument that the attention attracted by a stock should affect buying more than selling, as investors can choose from a large set of stocks when buying, and only have a limited choice when selling.

Hence, as the SVI increases we may expect an increase in stock return in the short run, and a decrease in stock return in the long run. This is confirmed by Da et al. (2011), which find that an increase in SVI for Russell 3000 stocks predict higher stock prices in the next two weeks and an eventual price reversal within the year. Bank et al. (2011) also find a positive short-run relationship between changes in SVI and future stock returns, with reversal for longer periods. Da et al. (2011) observe empirically that positive changes in the number of Internet queries push up prices temporarily. Bijl et al. (2016) investigate whether search query data on company names can be used to predict weekly stock returns for individual firms, and find that high levels of SVI predict low future excess returns. Barber and Odean (2007) observe that stocks that retail investors are buying (selling) during one week, have positive (negative) abnormal returns for the next two weeks, while the evidence that retail investors move prices over annual horizons is mixed.

The academic literature also includes studies that find no significant relationship between SVI and stock returns. Preis et al. (2010) investigate the correlation between returns and the SVI, and find no significant correlation. Challet and Bel Hadj Ayed (2013) critically discuss the claims that SVI data contain enough information to predict future financial index returns. By accounting for the many subtle biases that may affect the backtest of a trading strategy, they eliminate several of the biases in the results of Preis et al. (2013). They find that strategies based on financial keywords do not outperform strategies based on completely unrelated keywords.

A common of most of the studies mentioned is that they use either linear regression, quantile regression or correlation methods. All lacking the ability to capture complex, nonlinear relationships in the data, although the academic literature presents indicators that non-linear relationships exist. Da et al. (2011) and Joseph et al. (2011) find that sensitivity of returns to SVI is lowest for easy-to-arbitrage, low volatility stocks and highest for difficult-to-arbitrage, high volatility stocks. Bijl et al. (2016) find that the relationship between SVI and stock returns changes over time. As mentioned, Barber and Odean (2007) argue that the attention attracted by a stock should affect buying more than selling, making the relationship between SVI and trading volume non-linear. Preis et al. (2013) use a method for quantifying complex correlations in time series, and find that there exist complex dependencies between SVI and trading volume. Hence, as trading volume presents the demand of stock purchase, this may indicate that a complex relationship between SVI and stock returns exist. Challet and Bel Hadj Ayed (2014) claim that the use of linear strategies in stock return predictions is unsatisfactory. They criticize the work of Preis et al. (2013), stating that there is no reason why the given linear relationship in their study should hold for the whole period and for all stocks. It is, for instance, easy to find assets with consistently opposite reactions to SVI changes.

We contribute to the literature by implementing artificial neural networks for prediction of abnormal stock returns, able to capture highly complex, nonlinear relationships in the data. We use weekly data for several factors recognized as significant explanatory variables for stock returns by previous research, and evaluate models both with and without SVI data on company tickers in an effort to determine whether Google search volume data can indeed improve predictions.

Chapter 3

Data

3.1 Description

For our analyses we use weekly financial data, as well as weekly data on the Google Search Volume Index (SVI) on company tickers for S&P 1500 companies from 2004 through 2015. The data is obtained from CRSP, Compustat, Kenneth R. French's online data library and Google Trends. The CRSP and Compustat databases are used to obtain all relevant financial information and stock returns for companies of the S&P 1500 index. We obtain weekly values of Fama and French's three factors from French's online data library. Weekly SVI data is obtained from Google Trends. We choose to look at Google searches on company tickers and not company names because we want to capture investor attention, and not necessarily overall consumer attention, although the two are likely to be highly correlated.

We include all companies that have been part of the S&P 1500 index for some time period between the beginning of 2004 and the end of 2015. Only companies with a lack of SVI data, due to low search volumes, had to be excluded. We thus end up with data for 2321 companies in our data set, with a total of 1,012,857 observations. A complete list of the company tickers are given in Appendix A.

In an effort to isolate the true effect of SVI on stock returns, we include in our analyses a set of control variables: previous abnormal returns, volatility, abnormal trading volume and bid-ask spread.

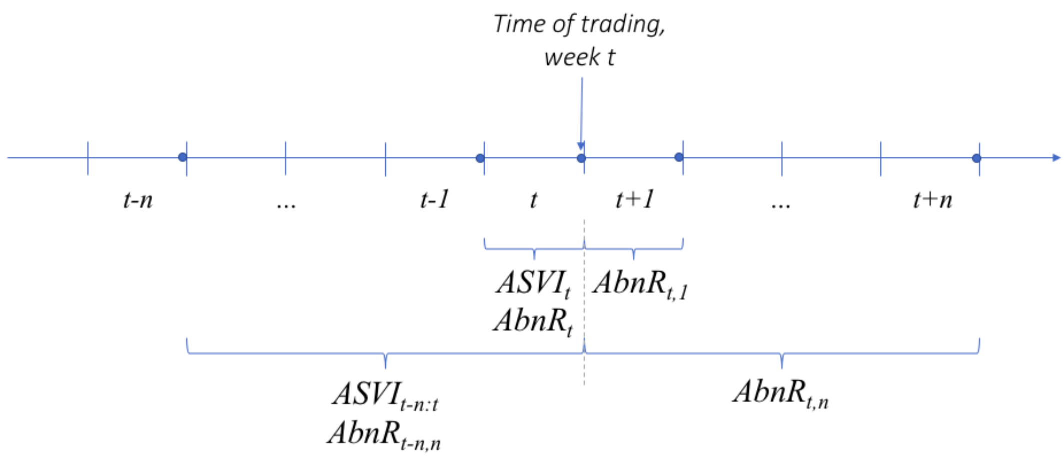

Figure 3.1: Timeline illustrating the notation used, with examples for two key variables, abnormal return, $AbnR$, and abnormal serach volumes, $ASVI$. Time is given by t and the trading horizon by n, both are denoted in weeks.

Figure 3.1 shows an illustration of the notation used and the timing of trades. We define the current week as t and the trading horizon in number of weeks n. Timing of trades is always at the end of the week. At that time all variable data for week t, as well as for previous weeks, is available. The subscript $t-n:t$ represents mean data over the previous n weeks, up to and including the current week t, and this data is also available at the end of week t. The subscript t,n represent cumulative data over the next n weeks, starting at week $t+1$, and this data is not available until the end of week $t+n$.

3.1.1 Google Search Volume Index (SVI)

Google Trends provide an open database of search volume indices constructed from unbiased samples of Google search data. Search volumes are reported as an index over time for each particular search term. The index is normalized within the range 0 and 100, and represents search volume as a proportion of total search volume at that time. The data is normalized given the time frame chosen during download, and thus the value of the index at a point in time is not meaningful in itself, as it can be manipulated to any arbitrary value by choosing the right time frame. Therefore, we standardize the SVI for each particular company by subtracting the mean and dividing by the standard deviation, based on the past year.

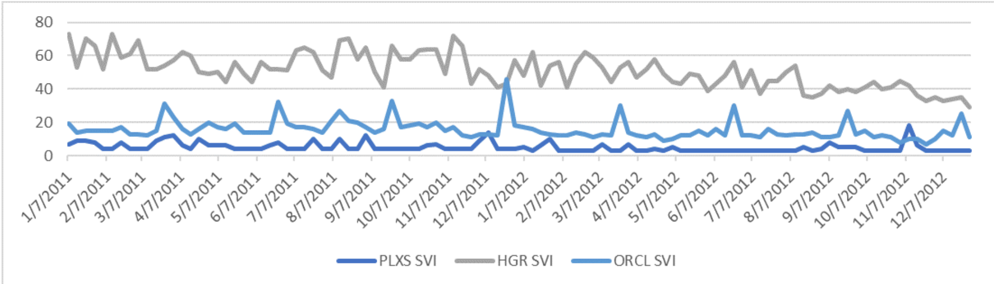

(a) Plain Google search volumes, SVI, for three companies with tickers PLXS, HGR and ORCL

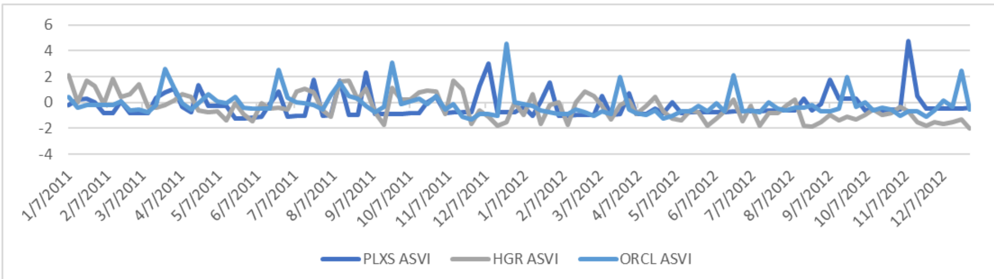

(b) Standardized Google search volumes, ASVI, for three companies with tickers PLXS, HGR and ORCL

Figure 3.2: Search volumes for three different companies from January 2011 to December 2012, before and after standardization

The SVI fetched from Google Trends is standardized according to Equation 3.1 to give the abnormal search volume index (ASVI):

$$ASVI_t = \frac{SVI_t - \frac{1}{52}\sum_{i=0}^{51} SVI_{t-i}}{\sigma_{SVI}} \tag{3.1}$$

where σ_{SVI} is the standard deviation of the SVI during the past year.

Figure 3.2 shows that the search volume indices can differ substantially from one company to another, and that the indices are more comparable after standardization.

We study the ASVI over different time periods, from one week to three months (approximated by 12 weeks). The mean ASVI over the past n weeks is given by:

$$ASVI_{t-n:t} = \frac{1}{n}\sum_{i=0}^{n-1} ASVI_{t-i} \tag{3.2}$$

3.1.2 Abnormal Stock Returns

We calculate weekly log returns according to Equation 3.3:

$$R_t = log(\frac{S_t}{S_{t-1}}) \tag{3.3}$$

where R is total log return, S is the stock price and t is time in weeks.

We calculate firm specific Fama-French beta coefficients from a rolling 1-year regression given by Equation 3.4:

$$R_t = \alpha + \beta_{MKT-Rf} \cdot R_{MKT-Rf,t} + \beta_{SMB} \cdot R_{SMB,t} + \beta_{HML} \cdot R_{HML,t} + \epsilon_t \tag{3.4}$$

where t is the time subscript for the data point and $R_{MKT-Rf,t}$, $R_{SMB,t}$ and $R_{HML,t}$ are the Fama-French factors. We use lagged abnormal returns as an input to our prediction models, and therefore it is important that no future data enters Equation 3.4. For this reason we use a rolling 1-year regression and calculate separate betas for each week, using data from the current week and the previous 51 weeks.

We then calculate the expected return, using the week specific betas:

$$ExpR_t = \alpha + \beta_{MKT-Rf} \cdot R_{MKT-Rf,t} + \beta_{SMB} \cdot R_{SMB,t} + \beta_{HML} \cdot R_{HML,t} \tag{3.5}$$

In order to calculate abnormal return, the actual and expected log returns must be transformed to simple returns and then subtracted. The abnormal log return is thus calculated as follows:

$$AbnR_t = log(e^{R_t} - e^{ExpR_t}) \tag{3.6}$$

We analyze returns over different time periods, from one week to three months. The cumulative abnormal return over the next n weeks is calculated as follows:

$$AbnR_{t,n} = \sum_{i=1}^{n} AbnR_{t+i} \tag{3.7}$$

3.1.3 Volatility

Several previous studies have found that volatility is correlated with stock returns, see for example French et al. (1987) and Balaban and Bayar (2005). We therefore include volatility as a control variable in our model. Specifically, we estimate standard deviation by using the square root of the Parkinson volatility estimator, a simple estimator using only the high and low stock prices during a given week. The estimator used is given by the formula in Equation 3.8:

$$Vol_t = \frac{1}{2\sqrt{log(2)}} \cdot (log(high_t) - log(low_t)) \tag{3.8}$$

where we use the highest ask price of week t, $high_t$, and the lowest bid price of week t, low_t.

We use mean volatility over different horizons in our analyses. The mean volatility over the past n weeks is given by:

$$Vol_{t-n:t} = \frac{1}{n}\sum_{i=0}^{n-1} Vol_{t-i} \tag{3.9}$$

3.1.4 Abnormal trading volume

Previous studies have also found high trading volumes to be associated with higher stock returns, see for example Gervais et al. (2001) and Paital and Sharma (2016). We therefore include trading volume as a control variable in our model. We standardize the trading volume in the same manner, and for the same reasons, as the SVI. Using Equation 3.10 we get the abnormal trading volume:

$$AbnVlm_t = \frac{Vlm_t - \frac{1}{52}\sum_{i=0}^{51} Vlm_{t-i}}{\sigma_{Vlm}} \tag{3.10}$$

where σ_{Vlm} is the standard deviation of volume during the past year.

We use mean abnormal trading volume over different time periods in our analyses. The mean abnormal volume over the past n weeks is given by:

$$AbnVlm_{t-n:t} = \frac{1}{n}\sum_{i=0}^{n-1} AbnVlm_{t-i} \tag{3.11}$$

3.1.5 Bid-Ask spread

Amihud and Mendelson (1986) find an economically and statistically significant positive relationship between bid-ask spread and stock returns, explained by a liquidity risk premium required by investors as the real return is impaired by a high bid-ask spread. We thus include the bid-ask spread also as a control variable in our model, calculated as follows:

$$BidAsk_t = \frac{ask_t - bid_t}{\frac{1}{2}(ask_t + bid_t)} \tag{3.12}$$

where ask_t is the mean ask price during week t and bid_t is the mean bid price during week t. As mean values over a week is used, the bid-ask spread in our data may in some rare cases be negative.

We use mean bid-ask spread over different horizons in our analyses. The mean bid-ask spread over the past n weeks is given by:

$$BidAsk_{t-n:t} = \frac{1}{n}\sum_{i=0}^{n-1} BidAsk_{t-i} \tag{3.13}$$

3.1.6 Summary Statistics

Table 3.1 shows summary statistics for the variables, with a total of 1,012,857 observations when all rows containing N/A values are excluded. For R and $AbnR$, the mean is close to zero, and an augmented Dickey Fuller test confirms stationarity. The low values for the 1st. and 3rd Quar-

Statistic	Mean	St. Dev.	Min	Max	1st. Q	3rd. Q
R	0.001	0.062	−3.277	1.625	-0.024	0.027
$AbnR$	−0.001	0.049	−2.775	1.603	-0.020	0.019
$ASVI$	0.010	0.926	−5.578	7.141	-0.620	0.539
Vol	0.019	0.016	0.000	1.348	0.011	0.023
$AbnVlm$	0.135	0.983	−3.043	6.919	-0.553	0.618
$BidAsk$	0.002	0.005	−0.128	0.932	0.0003	0.001
Number of observations	1,012,857					

Table 3.1: Summary statistics for key variables of weekly data

tile, as well as for Min and Max, indicate that most data points lie close to zero. For $ASVI$ and $AbnVlm$, the mean is close to zero, as expected after standardization. Also, the standard deviations for these variables are slightly larger than one, as expected after standardization when there is slight upward trend in the long run.

3.2 Linearity

We examine the relationship between the response, $AbnR_{t,n}$ and the explanatory variables, $ASVI_{t-n:t}$, $AbnR_{t-n,n}$, $Vol_{t-n:t}$, $AbnVlm_{t-n:t}$ and $BidAsk_{t-n:t}$.

Figure 3.3 to 3.5 show scatter plots and smoothed lines for the respective relationships. The smoothed lines are created by using the qplot() function in R with the geom="smooth" setting, and are essentially smoothed conditional means with confidence bands for the associated scatter plots. The smoothed lines are meant to provide information about the form of any relationship between the response variable and the explanatory variables. The scatter plots show no clear trends, and a high degree of variance. The scales of the scatter plots and the smoothed

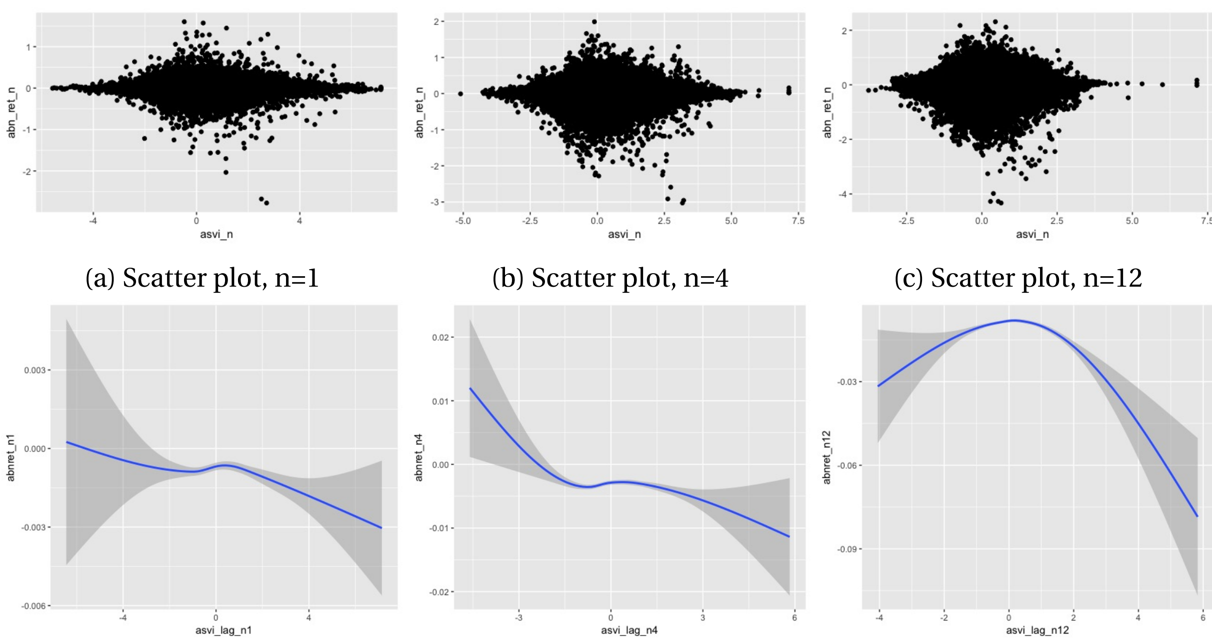

(a) Scatter plot, n=1 (b) Scatter plot, n=4 (c) Scatter plot, n=12

(d) Smoothing line, n=1 (e) Smoothing line, n=4 (f) Smoothing line, n=12

Figure 3.3: Scatter plots and smoothing lines for the relationship between $AbnR_{t,n}$ and $ASVI_{t-n:t}$. Smoothing lines include a 95 % confidence interval.

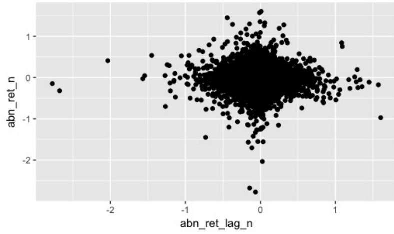

(a) Scatter plot, n=1

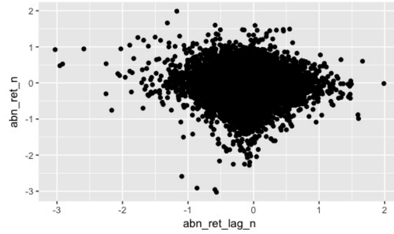

(b) Scatter plot, n=4

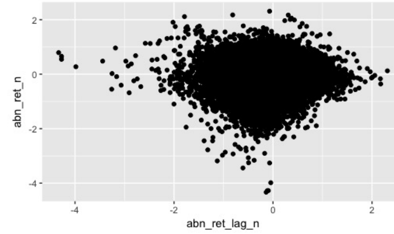

(c) Scatter plot, n=12

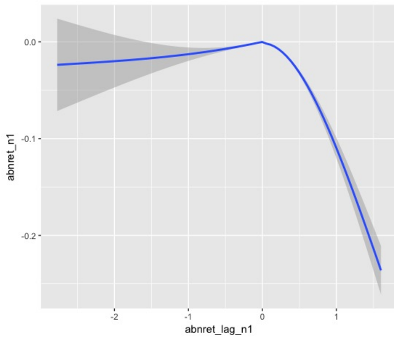

(d) Smoothing line, n=1

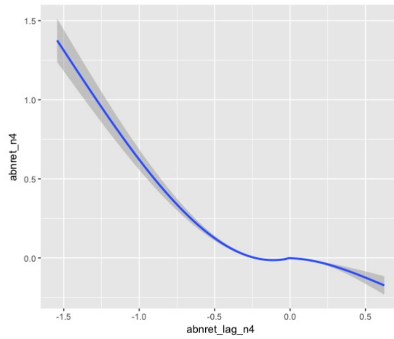

(e) Smoothing line, n=4

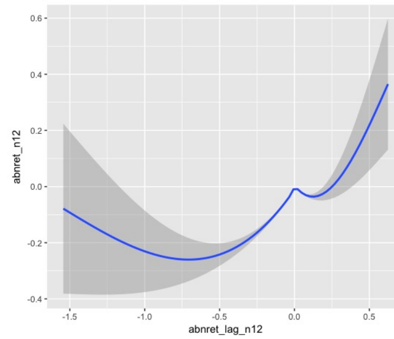

(f) Smoothing line, n=12

Figure 3.4: Scatterplots and smoothing lines for the relationship between $AbnR_{t,n}$ and $AbnR_{t-n,n}$. Smoothing lines include a 95 % confidence interval.

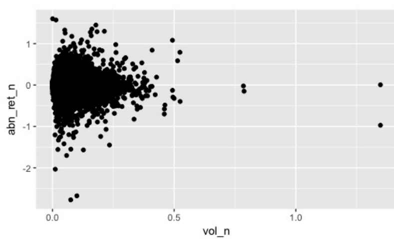

(a) Scatter plot, n=1

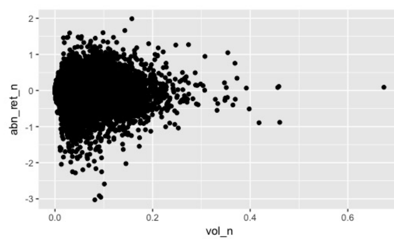

(b) Scatter plot, n=4

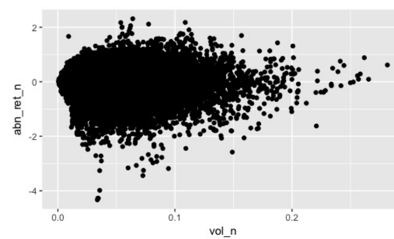

(c) Scatter plot, n=12

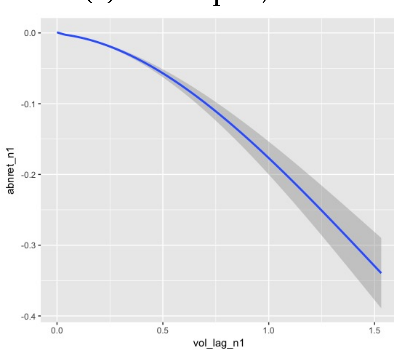

(d) Smoothing line, n=1

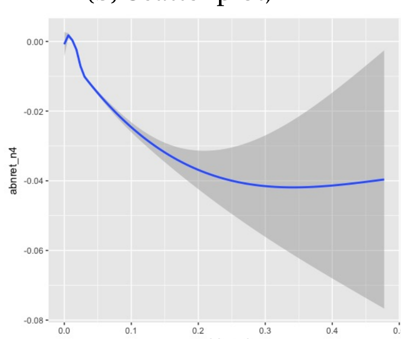

(e) Smoothing line, n=4

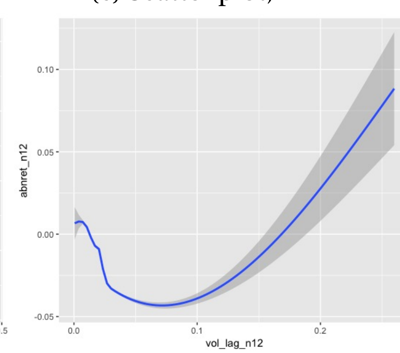

(f) Smoothing line, n=12

Figure 3.5: Scatter plots and smoothing lines for the relationship between $AbnR_{t,n}$ and $Vol_{t-n:t}$. Smoothing lines include a 95 % confidence interval.

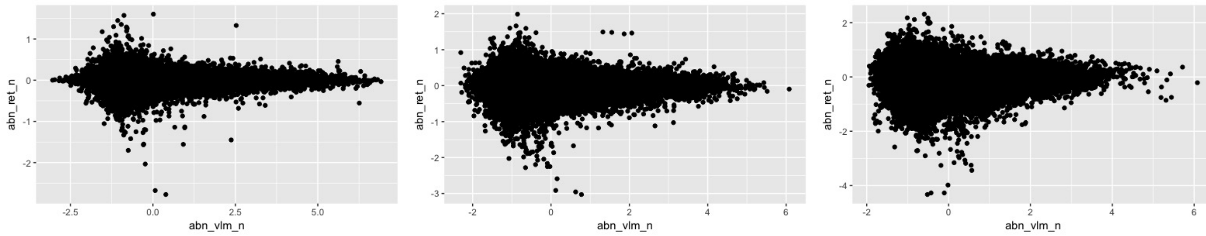

(a) Scatter plot, n=1 (b) Scatter plot, n=4 (c) Scatter plot, n=12

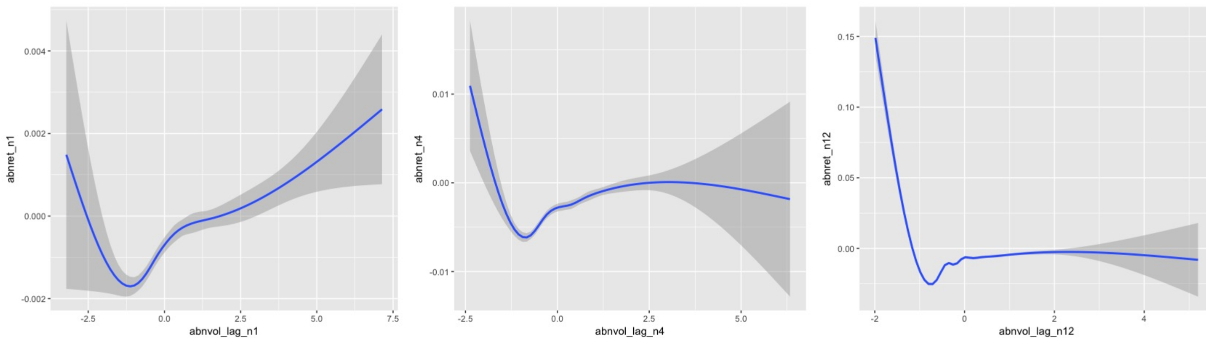

(d) Smoothing line, n=1 (e) Smoothing line, n=4 (f) Smoothing line, n=12

Figure 3.6: Scatter plots and smoothing lines for the relationship between $AbnR_{t,n}$ and $AbnVlm_{t-n:t}$. Smoothing lines include a 95 % confidence interval.

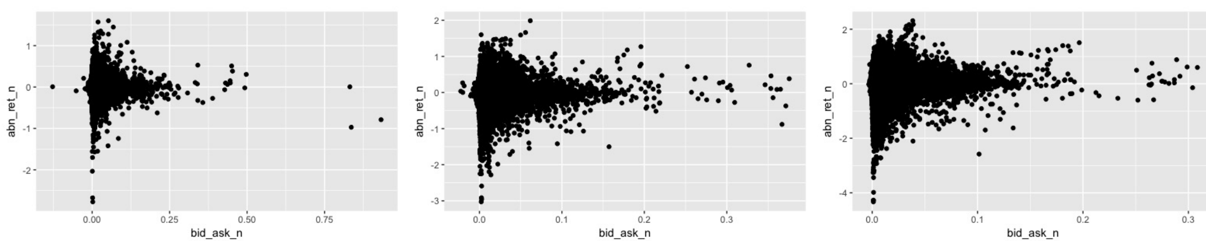

(a) Scatter plot, n=1 (b) Scatter plot, n=4 (c) Scatter plot, n=12

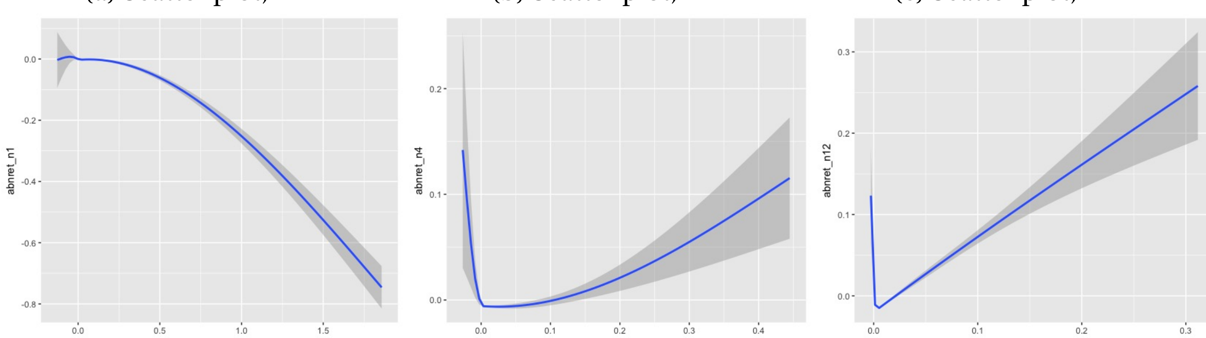

(d) Smoothing line, n=1 (e) Smoothing line, n=4 (f) Smoothing line, n=12

Figure 3.7: Scatter plots and smoothing lines for the relationship between $AbnR_{t,n}$ and $BidAsk_{t-n:t}$. Smoothing lines include a 95 % confidence interval.

lines are different due to the high variability in the scatter plot necessitating a larger range, thus they cannot be interpreted together directly. Nonetheless, the shapes of all of the smoothing lines are clearly non-linear, indicating that the impact of the explanatory variables on the response are complex and the relationships might not be suitably modeled by linear models.

3.3 Limitations

We cannot obtain SVI data for company tickers with low search volumes. This may introduce some survivorship bias in the data, as the remaining data points all meet a minimum threshold of search volume. However, there are only a few hundred companies which do not have SVI data, out of more than 2521 in total. Also, more importantly, we compare our results to a benchmark and a market portfolio consisting of the same selection of stocks, and the exclusion of these low search volume companies will hence not bias the results of the study.

Chapter 4

Methodology

In the following we outline the backtesting procedure. We proceed by introducing the methods used for prediction, the associated model specifications, and the metrics used for model evaluation. Finally, we present the trading strategy used in backtesting.

4.1 Backtesting Procedure

Prediction about the future *in the past* requires a proper backtesting procedure to prevent future information from entering the past, as this would render the results biased and likely invalid. The backtesting procedure outlines principles and methods to be followed in order to assure valid results.

We begin by defining a trading strategy, consisting of a prediction method and a set of rules for trade entries and exits. If the trading strategy is not defined a priori, problems could arise due to backtesting all kinds of strategies until one happens to be a good fit, most likely due to randomness. For the same reason, we choose the assets, the ASVI keywords and the holding periods prior to initiating backtests.

Next, backtesting of the trading strategy is performed. Using the results from backtesting we aim to determine whether non-linear methods are superior for abnormal stock return prediction, and further investigate the predictive power of $ASVI$ data. For the former, we compare trading strategies using different prediction models. For the latter, we compare trading strategies using artificial neural networks both with and without ASVI data. We further perform backtesting for various subsamples in time. Different time periods may be affected by different economical structures, for example the financial crisis in 2008 had a substantial impact on the market. This analysis will indicate how well the trading strategy performs under various market conditions.

Finally, the results from backtesting must be validated to assure robustness. This is done in two steps. Firstly, we perform backtesting for a trading strategy based completely on randomness, and approximate the distribution of the returns achieved by this strategy. This enables us to identify the robustness and significance of the results obtained from the trading strategy. Secondly, we run simulations on a validation set which is not previously used for testing. Trying to optimize tunable parameters, such as neural network specifications, while testing is equivalent to data snooping and may lead to poor out-of-sample performance due to overfitting. Therefore, if the results of the validation set are similar to those of the test set, there is support for validity of the results.

4.2 Prediction Models

For our backtesting analysis we utilize three different types of prediction models for comparison: linear regression models, Generalized Additive Models (GAM) and Artificial Neural Networks (ANN).

4.2.1 Linear Regression Model

Firstly, we adopt a standard linear regression model, as a preliminary analysis and as a benchmark for more complex models. The linear regression model is only able to capture linear relationships in the data, and we expect this model to perform poorly, possibly with inefficient parameter estimates. A linear regression model is on the form:

$$Y_i = \boldsymbol{X}_i \boldsymbol{\beta} \tag{4.1}$$

where Y_i is the response variable, $\boldsymbol{X}_i$ is the vector of predictor variables, and $\boldsymbol{\beta}$ is the parameter vector for the regression.

Model specification

We implement linear regression models according to Equation 4.2, for horizons n=1,4 and 12.

$$\begin{aligned} AbnR_{t,n} =& \beta_0 + \beta_1 AbnR_{t-n,n} + \beta_2 ASVI_{t-n:t} + \beta_3 AbnVlm_{t-n:t} + \beta_4 Vol_{t-n:t} \\ &+ \beta_5 BidAsk_{t-n:t} + \beta_6 ASVI_{t-n:t} * AbnR_{t-n,n} + \beta_7 ASVI_{t-n:t} * AbnVlm_{t-n:t} \\ &+ \beta_8 AbnR_{t-n,n} * AbnVlm_{t-n:t} + \epsilon_{t,n} \end{aligned} \tag{4.2}$$

In an effort to capture some of the complexities of the data we include three interaction terms: $ASVI_{t-n:t} \cdot AbnR_{t-n,n}$, $ASVI_{t-n:t} \cdot AbnVlm_{t-n:t}$ and $AbnR_{t-n,n} \cdot AbnVlm_{t-n:t}$.

Note that, for simplicity, we implement the trading horizon in a symmetric fashion. That is, for a weekly horizon, we model the one-week-ahead abnormal return based on data from the past week, for a monthly horizon we model the one-month-ahead abnormal return based on data from the past month, and so on. Another possible implementation would be to always use data from the past quarter or the past year, even for predictions one week or one month into the future.

4.2.2 Generalized Additive Model (GAM)

Due to indications of non-linear relationships in the data, we choose to also adopt a GAM. The method has previously proven useful in uncovering nonlinear covariate effects (Hastie and Tibshirani (1986)). A GAM is a semi-parametric regression model on the form:

$$g(\mu_i) = \boldsymbol{X}_i \boldsymbol{\beta} + f_1(x_{1i}) + ... + f_m(x_{mi}) + f_{m+1}(x_{1i}, ..., x_{mi}) + ... + f_{m+n}(x_{ni}, ..., x_{mi}) \tag{4.3}$$

$\boldsymbol{X}_i \boldsymbol{\beta}$ constitutes the parametric component of the model. The explanatory variables $\boldsymbol{X}_i$ are included in the model as linear components, and should thus have a linear relationship with the response $g(\mu_i)$. The smooth functions, $f_1, f_2, ..., f_{m+n}$, apply to the explanatory variables, $x_{1i}, x_{2i}, ..., x_{mi}$, and constitute the nonparametric component of the model. As evident from Equation 4.3, the smooth functions $f(\cdot)$ can include several explanatory variables. Also, an explanatory variable entering the non-parametric component of the equation can enter several smooth functions, ex. $f_1(x_{1i})$ and $f_{m+1}(x_{1i}, ..., x_{mi})$.

The $f(\cdot)$'s are smooth functions which allow for rather flexible specifications of the dependence of the response variable on the explanatory variables. We estimate the smooth functions using a cubic regression spline. Cubic polynomials are fitted to the shape in segments and connected at points called knots, such that the function is continuous up to the second derivative.

In Equation 4.3, g is a smooth monotonic link function, $\mu_i = \mathbb{E}(Y_i)$, and Y_i is the response variable which follows some exponential family distribution. The distribution of Y_i must be predetermined together with the function g.

Model specification

We use four different GAM models, one for each of n=1 and 4, and two different for n=12. We use the same set of explanatory variables as for the linear regression model. As a first step in the model development, we let all of the explanatory variables enter the equation as smooth functions and examine their resulting shape. The smooth functions with an approximate linear shape are excluded, and replaced by linear terms. A second model for the trading horizon n=12, excluding multivariate smooth terms, are included in order to illustrate the isolated effect of the smooth function of $ASVI_{t-n:t}$.

We use the mgcv package in R to construct the four GAM models with different model specifications. Equation 4.4 includes all terms that may be present in the models.

$$\begin{aligned} AbnR_{t,n} =& \beta_0 + \beta_1 ASVI_{t-n:t} + \beta_2 AbnVlm_{t-n:t} + \beta_3 BidAsk_{t-n:t} + \beta_4 Vol_{t-n:t} \\ &+ f_1(AbnR_{t-n,n}) + f_2(ASVI_{t-n:t}) + f_3(AbnVlm_{t-n:t}) + f_4(BidAsk_{t-n:t}) \\ &+ f_5(Vol_{t-n:t}) + f_6(ASVI_{t-n:t}, AbnR_{t-n,n}) + f_7(ASVI_{t-n:t}, AbnVlm_{t-n:t}) \\ &+ f_8(AbnR_{t-n,n}, AbnVlm_{t-n:t}) + \epsilon_{t,n} \end{aligned} \tag{4.4}$$

The left-hand side of Equation 4.4 is the response variable, cumulative abnormal return over the next n weeks. We use the identity link function for $g(\cdot)$ in Equation 4.3. The right-hand side is composed of several terms of predictor variables, β_n represent coefficients of the linear predictors and $f_k(\cdot)$ are the non-parametric smooth functions. The three last terms are multivariate smooth functions composed of two variables, resulting in three dimensional curves, and are based on the three interaction variables used in Equation 4.2.

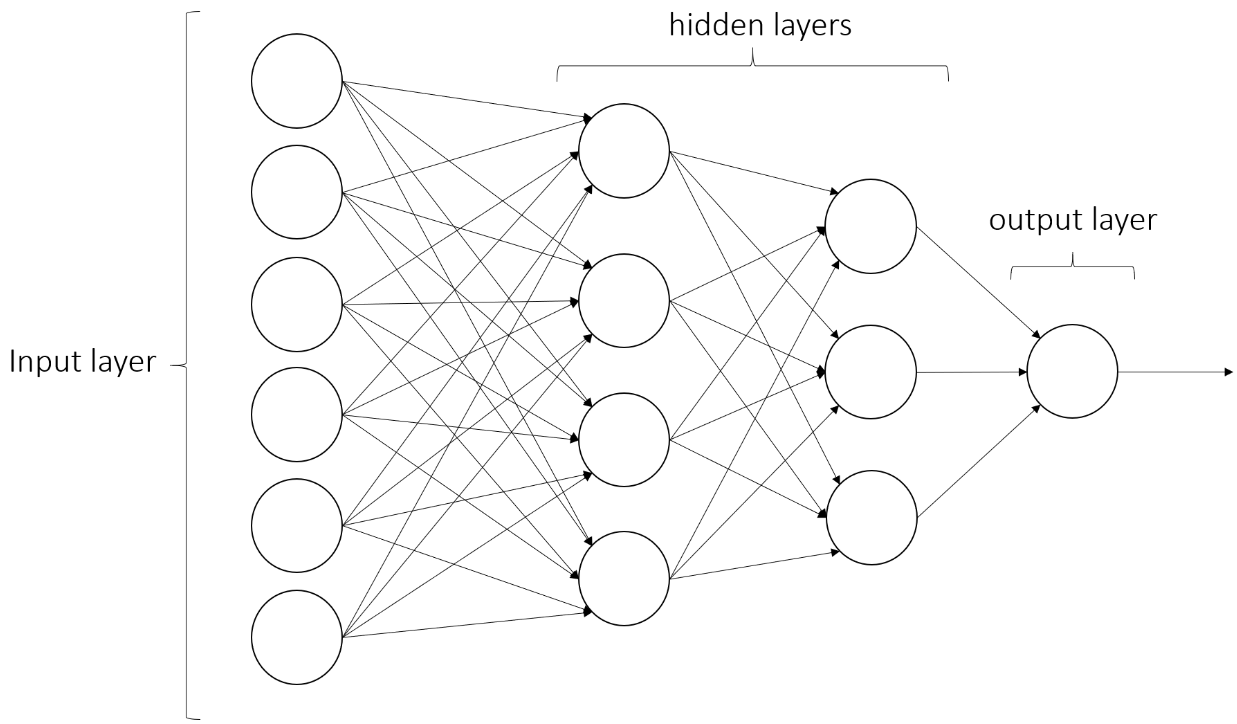

Figure 4.1: Visualization of the structure of an artificial neural network

4.2.3 Artificial Neural Network (ANN)

Finally, we adopt an ANN model. ANNs are a familiy of methods within machine learning, a field of computer science concerned with enabling computers to learn from data without being explicitly programmed. ANNs can be used to perform both regression (continuous output) and classification (discrete output), we will focus here on ANNs in regression.

An ANN is composed of a number of nodes connected by directed links, see Figure 4.1. The nodes are structured in layers: the input layer has one node for each input variable; the output layer has one node for each output value, which is just one for regression; hidden layers are between the input and output layer. Choosing the number of hidden layers and the number of nodes in each hidden layer is an important design choice when constructing an ANN. Between each layer are directed links, enabling information to "flow" through the network. Our focus will be on feed-forward ANNs, in which all links are directed forward in the network and all nodes in a layer thus have links to all nodes in the next layer. Figure 4.1 shows the structure of a feed-forward ANN with six input variables, one output node and two hidden layers.

Each link in the ANN has a weight w_{ij} associated with it. At each node an activation function transforms the weighted sum of the signals from the previous layer into the output that is sent to the next layer. To find optimal values for the weights, w_{ij}, the ANN is trained on existing observations, the training data. During training the weights are gradually adjusted, for example using gradient descent, to minimize the error of the output. We will not go into further technical or mathematical detail of ANNs, refer to for example Russel and Norvig (2010) for a rigorous derivation of the techniques used in training an ANN.

Deep learning is a a branch within machine learning concerned with learning data representations in stead of learning specific data. Thus deep learning models should be able to generalize better to out-of-sample data, and better learn the relationships in the data even if it is noisy, as is often the case with real-world applications. Deep neural networks are ANNs with multiple hidden layers. The goal is that each layer is then allowed to learn one representation of the data, and together the layers will be able to give accurate predictions for the output variable. Deep learning is especially useful when it can be trained on a very large data set, like we have for this study. We refer to LeCun et al. (2015) for more details on deep learning.

In contrast to the linear regression model and GAM described above, ANNs make no assumptions about the input data, its distributions or correlations. It is a highly flexible method capable of capturing complex relations in data. This comes, of course, at a cost. Machine learning methods are by definition *black boxes*, due to the fact that they need no explicit programming. It is usually difficult to find meaningful interpretations of the relationships between variables in an ANN. Thus, while an ANN can be a powerful prediction model capable of high precision in out-of-sample predictions, it is more difficult to provide economic interpretations that can be used directly to describe or analyze a market. Nonetheless, our main goal in this part of the study is to achieve accurate predictions.

Model specification

Specification of the model parameters used for the ANN are shown in Table 4.1. We use the MLPRegressor from the scikit-learn package in Python to construct the neural network. After testing the model with different numbers of hidden layers and numbers of neurons we find that the best performance (within a reasonable time complexity) is achieved with 3 hidden layers, with 512, 256 and 128 neurons for each respective layer. As the ANN has multiple hidden layers,

Model parameter	**Value**
Number of hidden layers	3
Number of neurons in hidden layers	(512, 256, 128)
Activation function	Rectified linear unit (ReLU)
Solver	ADAM (a form of stochastic gradient descent)
Learning rate	Set by the ADAM solver
Max. iterations	500
L2 regularization term	0.0001

Table 4.1: Specification of model parameters for Artificial Neural Network

it is classified as a deep ANN.

For training the weights in the ANN we use the ADAM solver. This is a variation of stochastic gradient descent developed by Kingma and Ba (2014). The method computes individual adaptive learning rates for different parameters from estimates of first and second moments of the gradients; the name ADAM is derived from adaptive moment estimation. For the activation function of the nodes we use the rectified linear unit (ReLU), given by:

$$f(x) = max(0, x). \tag{4.5}$$

This is a very common activation function in ANNs as it makes training easy. Refer to Zeiler et al. (2013) for more details on the advantages of the ReLU. For regularization, i.e. penalization of more complex models in an effort to reduce overfitting, we use the default value of 0.0001.

4.2.4 Model Evaluation

The prediction models are evaluated and compared based on the following metrics: mean squared error (MSE), hit rate, precision and recall.

MSE of predictions gives information about the overall goodness of fit. It is not a meaningful metric on its own, only in comparison to other models. The MSE is calculated as follows:

$$MSE_n = \frac{1}{N}\sum_{i=1}^{N}\left(AbnR_{t_i,n} - AbnR_{t_i,n}^{pred}\right)^2 \tag{4.6}$$

where N is the size of the test data, the subscript t_i is the time in weeks of the i-th data point, $AbnR$ is the actual abnormal return and $AbnR^{pred}$ is the predicted abnormal return. The smaller the MSE, the better the fit of the model.

The other metrics we use are more meaningful stand-alone than the MSE, and are concerned with the direction of abnormal returns, i.e. whether they are positive or negative. For these metrics some new terminology is introduced. True positives are correctly predicted positive $AbnR$, and false positives are negative $AbnR$ that are predicted as positive. Likewise, true negatives are correctly predicted negative $AbnR$, and false negatives are positive $AbnR$ that are predicted as negative.

The hit rate is the percentage of test observations that are correctly predicted as giving either positive or negative abnormal return, i.e. the sum of true positives and true negatives divided by the total number of observations. This score represents the model's overall ability to predict the direction of abnormal returns. The hit rate is calculated according to the following formula:

$$HIT\_RATE_n = \frac{1}{N}\sum_{i=1}^{N}\left\{AbnR_{t_i,n}^{pred} > 0\right\}\cdot\left\{AbnR_{t_i,n} > 0\right\} + \left\{AbnR_{t_i,n}^{pred} < 0\right\}\cdot\left\{AbnR_{t_i,n} < 0\right\} \quad (4.7)$$

The larger the hit rate, the better the model is at predicting the direction of abnormal returns. The expected value of the abnormal return should be zero. Therefore a model score over 0.5 would entail an improvement over a random model.

In order to get more detailed information about the performance of the models we calculate positive and negative precision, as well as positive and negative recall. Precision of positive predictions is defined as the ratio of true positives to the total number of positive predictions, which is the sum of true positives and false positives. Precision of negative predictions is defined in the same manner, see Equations 4.8 and 4.9. The precision provides information about how accurate a positive or negative prediction is, i.e. how likely it is to be true. Recall of positive predictions is defined as the ratio of true positives to the total number of positive observations, which is the sum of true positives and false negatives. Recall of negative predictions is defined in the same manner, see Equations 4.10 and 4.11. The recall provides information about how good the model is at identifying positive or negative $AbnR$.

$$PRECISION_n^{pos} = \frac{\sum_{i=1}^{N}\{AbnR_{t_i,n}^{pred} > 0\}\cdot\{AbnR_{t_i,n} > 0\}}{\sum_{i=1}^{N}\{AbnR_{t_i,n}^{pred} > 0\}} \quad (4.8)$$

$$PRECISION_n^{neg} = \frac{\sum_{i=1}^{N} \{AbnR_{t_i,n}^{pred} < 0\} \cdot \{AbnR_{t_i,n} < 0\}}{\sum_{i=1}^{N} \{AbnR_{t_i,n}^{pred} < 0\}} \tag{4.9}$$

$$RECALL_n^{pos} = \frac{\sum_{i=1}^{N} \{AbnR_{t_i,n}^{pred} > 0\} \cdot \{AbnR_{t_i,n} > 0\}}{\sum_{i=1}^{N} \{AbnR_{t_i,n} > 0\}} \tag{4.10}$$

$$RECALL_n^{neg} = \frac{\sum_{i=1}^{N} \{AbnR_{t_i,n}^{pred} < 0\} \cdot \{AbnR_{t_i,n} < 0\}}{\sum_{i=1}^{N} \{AbnR_{t_i,n} < 0\}} \tag{4.11}$$

4.3 Trading Strategy

In order to evaluate the prediction models in an economic sense we implement a trading strategy with each of the models. We simulate trading over several years, recording the returns earned and compare this to a simple equally weighted portfolio containing all stocks. Table 4.2 presents a definition of the trading strategy. The strategy involves buying the 50% of stocks with highest predicted abnormal return in each period, and selling those stocks among the 50% with lowest predicted abnormal return that are currently held in the portfolio. For simplicity there is no shorting and the stocks in the portfolio are equally weighted. Note that since the

	Strategy
Buy	Buy the 50% highest ranking stocks from the predictive model (i.e. with highest predicted abnormal return) which are not already in the portfolio
Sell	Sell all stocks in the portfolio which are among the 50% lowest ranking stocks
Shorting	No shorting
Weighting	All stocks in the portfolio are always equally weighted

Table 4.2: Trading strategy definition

same amount of capital is always invested the strategy is solely concerned with stock picking, and not timing the market.

We use data from 2004 to the end of 2010 as the training set, data from 2011 to the end of 2014 as the test set, and data from 2015 as the validation set. For the simulation we retrain the model with the realized returns after each trading period, enabling online learning over time.

Trading costs are set to 0.10% of the traded value, and is composed of half the average bid-ask spread in our sample, 0.085%, and a brokerage fee of 0.015%. The brokerage fee may seem small, but we argue that this type of sophisticated trading strategy based on complex prediction models would not be relevant for individual investors. We assume that a substantial amount of capital is to be invested and that institutional investor rates would be available, and in this case the brokerage fee of 0.015% is not unreasonable.

Chapter 5

Results

In the following we first present a brief descriptive analysis in an effort to gain insight into the relationship between abnormal returns and the ASVI. We construct models for abnormal stock returns using both linear regression and GAM, over the full period of 2004–2015.

Secondly, we present our main results, which is a predictive analysis. We construct an ANN for prediction, and obtain model evaluation scores through out-of-sample testing. We also evaluate the linear regression model and GAM for prediction. Further, we use the prediction models with the trading strategy outlined in Section 4.3. The trading strategy enables us to test the value of the models in an economic sense. Moreover, in order to assess the predictive power of the ASVI data, we compare the results of our ANN model to a benchmark ANN model which does not use ASVI data. Further, we analyze the models for different time periods to examine the performance under different market conditions. Finally, we aim to validate the results, firstly by comparing with a randomized trading strategy, and secondly by running the models on the validation set, consisting of data from 2015.

5.1 Descriptive analysis

5.1.1 Linear Regression

We run a linear pooled regression, using data from the whole period of 2004–2015. As abnormal return by definition should not have fixed effects, we find that a panel data regression with fixed

	Dependent variable:		
	$AbnR_{t,n}$		
	n=1	n=4	n=12
$AbnR_{t-n,n}$	-0.018^{***}	-0.008^{*}	0.288^{***}
	(0.001)	(0.004)	(0.012)
$ASVI_{t-n:t}$	-0.0001	-0.0001	-0.0004
	(0.00005)	(0.0001)	(0.0003)
$AbnVol_{t-n:t}$	0.001^{***}	0.002^{***}	0.003^{***}
	(0.00004)	(0.0001)	(0.0003)
$BidAsk_{t-n:t}$	-0.058^{***}	0.339^{***}	1.764^{***}
	(0.008)	(0.022)	(0.042)
$Vol_{t-n:t}$	-0.080^{***}	-0.401^{***}	-1.195^{***}
	(0.003)	(0.007)	(0.015)
$AbnR_{t-n,n}*ASVI_{t-n:t}$	0.003^{***}	-0.019^{***}	-0.044^{***}
	(0.001)	(0.004)	(0.015)
$ASVI_{t-n:t}*AbnVol_{t-n:t}$	-0.0002^{***}	-0.0004^{***}	-0.002^{***}
	(0.00003)	(0.0001)	(0.0004)
$AbnR_{t-n,n}*AbnVol_{t-n:t}$	0.005^{***}	0.069^{***}	0.225^{***}
	(0.001)	(0.004)	(0.013)
Constant	0.001^{***}	0.004^{***}	0.010^{***}
	(0.0001)	(0.0002)	(0.0003)
Observations	1,015,152	1,008,157	989,127
R^2	0.001	0.003	0.009
Adjusted R^2	0.001	0.003	0.009
Residual Std. Error	0.049 (df = 1015143)	0.096 (df = 1008148)	0.166 (df = 989118)
F Statistic	185.944^{***} (df = 8; 1015143)	442.168^{***} (df = 8; 1008148)	$1,068.992^{***}$ (df = 8; 989118)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 5.1: Regression results for the linear regression model given by Equation 4.2, for n=1,4 and 12, including interaction terms.

effects is not necessary, and will give very similar results. The results from the pooled regression are given in Table 5.1. We observe that R^2 is low for the respective models, ranging from 0.1% to 0.9%, indicating that the independent variables have little explanatory power of abnormal return in the linear models. Nonetheless, R^2 is increasing with n, indicating increasing explanatory power for longer holding periods.

The coefficient of $ASVI_{t-n:t}$ is not significant at the 10% level for any of the respective trading horizons. On the other hand, the three interaction terms are significant at the 1% level, indicating that some dependencies between the explanatory variables may be evident. However, a Breusch-Pagan test indicates the presence of heteroscedasticity in the residuals of all three models, possibly invalidating statistical tests of significance. Thus, there is reason to believe that more complex models are needed to gain further insight.

5.1.2 GAM

The results for the four GAM models implemented are given in Table 5.2. The overall performance is slightly better than for the linear models. However, the Adjusted R^2 is still quite low, ranging from 0.2% to 1.4% with higher values for longer trading horizons.

Again, the coefficient of $ASVI_{t-n:t}$ is not significant at the 10% level for n=1 and n=4. However, for n=12, the $ASVI_{t-n:t}$ coefficient of model (4) is significant at the 1% level and slightly negative, supporting the findings of Bijl et al. (2016), Da et al. (2011) and Bank et al. (2011) for longer holding periods. The $ASVI_{t-n:t}$ smooth function of model (3) is also significant and shown in Figure 5.1a. The shape is slightly concave with a maximum at approximately (0,0), indicating that abnormal search volume in any direction will affect abnormal returns, $AbnR_{t,n}$, negatively in the long run. However, when we include the multivariate smooth terms in model (4), we observe that the shape of the $ASVI_{t-n:t}$ smooth function becomes flat and approximately zero. Thus, the shape of the $ASVI_{t-n:t}$ smooth function in model (3) may be caused by non-linear dependencies with other explanatory variables.

The multivariate smooth terms are significant in all the models in which they are included. Figures 5.1b - 5.1d show the multivariate smooth functions estimated for n=12. The smooth functions are highly complex, indicating quite complicated relationships between the explanatory variables and the response. The combined effect of $ASVI_{t-n:t}$ and other explanatory vari-

	Dependent variable:			
	$AbnR_{t,n}$			
	(1) n=1	(2) n=4	(3) n=12	(4) n=12
$ASVI_{t-n:t}$	−0.00004	−0.0002	–	-0.001^{***}
	(0.0001)	(0.0001)		(0.0003)
$AbnVol_{t-n:t}$	0.001^{***}	0.002^{***}	–	0.004^{***}
	(0.0001)	(0.0001)		(0.0003)
$BidAsk_{t-n:t}$	–	0.401^{***}	–	1.881^{***}
		(0.022)		(0.042)
$Vol_{t-n:t}$	-0.102^{***}	-0.433^{***}	–	-1.183^{***}
	(0.003)	(0.008)		(0.017)
s($AbnR_{t-n,n}$)	***	***	***	***
s($ASVI_{t-n:t}$)	–	–	***	–
s($AbnVol_{t-n:t}$)	–	–	***	–
s($BidAsk_{t-n:t}$)	***	–	***	–
s($Vol_{t-n:t}$)	–	–	***	–
ti($ASVI_{t-n:t}$,$AbnR_{t-n,n}$)	***	***	–	***
ti($ASVI_{t-n:t}$,$AbnVol_{t-n:t}$)	***	***	–	***
ti($AbnR_{t-n,n}$,$AbnVol_{t-n:t}$)	***	***	–	***
Constant	0.001^{***}	0.004^{***}	-0.009^{***}	0.009^{***}
	(0.0001)	(0.0002)	(0.0002)	(0.0004)
Observations	1,012,857	1,008,781	990,389	990,389
Adjusted R^2	0.002	0.006	0.014	0.013
Log Likelihood	1,620,111.000	931,256.800	371,040.600	370,437.400
UBRE	0.002	0.009	0.028	0.028
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 5.2: Regression results for the GAM given by Equation 4.4, for trading horizons of n=1,4 and 12. The coefficients of the strictly linear terms are given at the top of the table. The significance of the smooth functions are given next, denoted as $s()$ and $ti()$. The symbol – denotes that the particular term is not part of the given model.

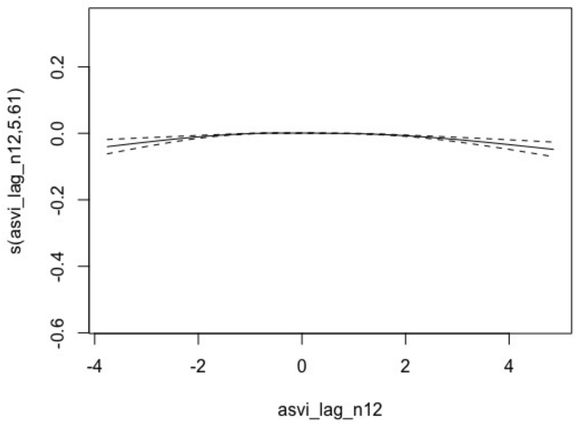

(a) f($ASVI_{t-n:t}$), GAM model (3)

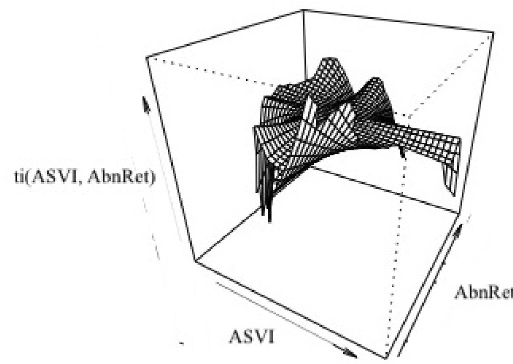

(b) f($ASVI_{t-n:t}$, $AbnR_{t-n,n}$), GAM model (4)

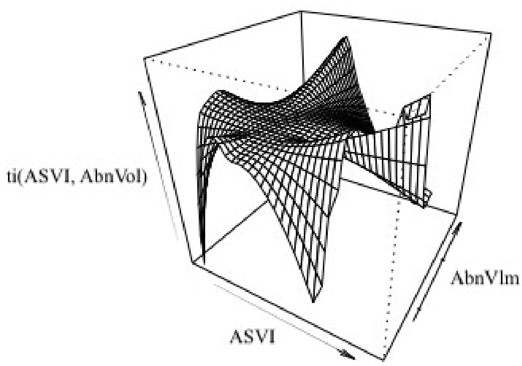

(c) f($ASVI_{t-n:t}$, $AbnVol_{t-n:t}$), GAM model (4)

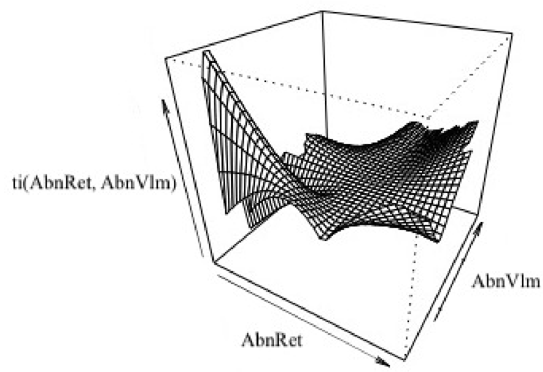

(d) f($AbnR_{t-n,n}$, $AbnVol_{t-n:t}$), GAM model (4)

Figure 5.1: GAM Smooth functions for the given variables and GAM models

ables seem to have a non-linear effect on the response, although the magnitude may be low.

The residuals of all the GAM models show signs of heteroscedasticity, which may invalidate statistical tests of significance. We find it most appropriate to proceed by utlizing non-linear machine learning techniques for prediction of abnormal return, and use the prediction models to attempt to gain more insight into the relationship between ASVI and abnormal return.

5.2 Prediction model evaluations

In this section we use ANN models to predict abnormal stock returns, and compare the performance to prediction models based on linear regression and GAM. Results from evaluation of

Model	Linear regression			GAM				ANN		
Trading horizon	n=1	n=4	n=12	n=1 (1)	n=4 (2)	n=12 (3)	n=12 (4)	n=1	n=4	n=12
MSE	0.0017	0.0068	0.0211	0.0017	0.0069	0.0211	0.0211	0.0016	0.0064	0.0192
HIT_RATE	0.5102	0.5164	0.5264	0.5107	0.5122	0.5225	0.5244	0.5058	0.5264	0.5555
$PRECISION^{pos}$	0.5020	0.5070	0.5160	0.5030	0.5020	0.5080	0.5110	0.4964	0.5251	0.5585
$PRECISION^{neg}$	0.5170	0.5240	0.5340	0.5180	0.5280	0.5390	0.5380	0.5110	0.5277	0.5517
$RECALL^{pos}$	0.4780	0.4830	0.4450	0.4980	0.6280	0.5630	0.5230	0.4146	0.3913	0.4072
$RECALL^{neg}$	0.5420	0.5490	0.6040	0.5230	0.4010	0.4840	0.5250	0.5850	0.6581	0.6923

Table 5.3: Model evaluation scores for linear regression model, GAM and ANN model, with trading horizons of one week, n=1, one month, n=4, and one quarter, n=12

the predictive models are presented in Table 5.3. For monthly and quarterly trading we run the model for each possible starting week in the period and take average scores in order to decrease noise in the results. Furthermore, due to the inherent randomness in the training process of the neural network we also run each ANN model 50 times and take averages.

The mean squared error, MSE, is increasing with the trading horizon for all models, however, this is to be expected as the standard deviation in abnormal returns increase with the holding period. All of the other metrics suggest that the models are better at predicting longer-term abnormal returns, although the increase is not substantial for linear regression and GAM. The MSE is slightly better for the ANN model than the linear regression model and GAM for all horizons. Common to most of the models is a weak recall for positive predictions and a strong recall for negative predictions. Thus, the models are generally good at identifying negative samples, but not so good at identifying positive samples. This might indicate that the signals of subpar returns are more prominent in the data than signals of above par returns.

For weekly trading, the linear regression model and GAM give better scores than the ANN for most metrics. This is quite surprising, and the reason is not altogether clear. However, the margins are small, and given the low R^2 for the linear regression model and GAM it is plausible that the differences are not significant. Nonetheless, neither of the weekly models provide very compelling results. The hit rate for the ANN is just marginally above 50%, meaning it is not much better than guessing. Moreover, the precision of positive predictions is not strong for any of the models, indicating around a 50% probability of a positive prediction being correct. We suspect that there is not enough information in weekly data to make valuable predictions.

For monthly trading, the hit rate and precision metrics are higher than for weekly, although the improvement is very slight for the linear regression model and GAM. As expected, the ANN model outperforms the other models, and with a hit rate, as well as positive and negative precision of almost 53%, the model is likely significantly better than random guessing.

For quarterly trading, we obtain the best results of all trading horizons tested. The ANN model gives substantially better results than the linear regression model and GAM on most metrics, with a hit rate, as well as precision of both positive and negative predictions of more than 55%. The linear and semi-parametric methods have minimal improvements in model scores for longer horizons. On the other hand, the ANN model provides substantial increases in most model scores for longer trading horizons. It seems that there are relationships in the data for longer horizons that the ANN is able to capture, while the linear and semi-parametric models are not.

Considering the model evaluations, there is reason to believe that predictions of significant value can only be made for longer trading horizons. The model evaluations also show that the ANN outperforms the other models, with the GAM being slightly better than the linear regression model. We find support for our hypothesis that non-linear, machine learning methods are superior to linear and semi-parametric methods and may be necessary for satisfactory abnormal stock return prediction.

5.3 Trading strategy

Annualized results from simulations with the trading strategy, where the 50% of stocks with highest predicted abnormal return are held, are presented in Table 5.4. The Sharpe ratios are calculated using as the risk free rate the average US 1-year Treasury rate over the test period, from 2011 through 2014, which is 0.165%. The results verify that the linear regression model and GAM are unsatisfactory, giving lower returns and a lower Sharpe ratio than the equally weighted portfolio for all trading horizons. Interestingly, the GAM achieves equal returns as the linear regression model, for all trading horizons. Thus, although the GAM has slightly better model evaluation scores, it is not able to exploit this with the trading strategy. On the other hand, the ANN model outperforms the equally weighted portfolio for all trading horizons before accounting for

		Excl. transaction costs			Incl. transactions costs		
Horizon	Model	Return	Volatility	Sharpe Ratio	Return	Volatility	Sharpe Ratio
Weekly	Linear regression	9.8%	16.3%	0.56	7.2%	16.3%	0.40
	GAM	9.8%	16.3%	0.56	7.2%	16.3%	0.40
	ANN	15.3%	16.1%	0.94	12.0%	16.1%	0.74
	Equally weighted	14.8%	17.9%	0.82	14.8%	17.9%	0.82
Monthly	Linear regression	8.1%	8.1%	0.98	7.2%	8.1%	0.87
	GAM	8.1%	8.1%	0.98	7.2%	8.1%	0.87
	ANN	17.8%	11.9%	1.48	17.0%	11.9%	1.41
	Equally weighted	13.2%	12.9%	1.01	13.2%	12.9%	1.01
Quarterly	Linear regression	9.0%	7.7%	1.15	13.5%	7.7%	1.12
	GAM (3)	9.0%	7.7%	1.15	13.5%	7.7%	1.12
	GAM (4)	9.0%	7.7%	1.15	13.5%	7.7%	1.12
	ANN	18.3%	9.1%	1.99	18.0%	9.1%	1.96
	Equally Weighted	13.0%	10.5%	1.23	13.0%	10.5%	1.22

Table 5.4: Annualized results with the trading strategy, where the 50% of stocks with highest predicted abnormal return are held. Results for linear regression model, GAM and ANN are presented, as well as an equally weighted portfolio for comparison. Trading horizons of one week, n=1, one month, n=4, and one quarter, n=12. Trading costs are 0.10% of the traded value. The Sharpe ratio is calculated using the average US 1-year Treasury rate over the test period, which is 0.165%, as the risk free rate.

transaction costs. There is also a substantial increase in performance of the ANN with longer trading horizons, and monthly and quarterly horizons provide positive excess returns over the equally weighted portfolio also after accounting for transaction costs. Moreover, the volatility is below the equally weighted portfolio for all trading horizons, with substantial improvements for longer horizons.

Figure 5.2 shows the cumulative returns of the ANN model using the trading strategy, with and without transaction costs, and the normal return given by the equally weighted portfolio, for weekly, monthly and quarterly trading horizons. The returns for the ANN trading strategies closely follow the movements of the equally weighted portfolio, but lie consistently above, except for weekly trading horizon when accounting for trading costs.

The cumulative excess returns for the ANN over the equally weighted portfolio excluding transaction costs are shown in Figure 5.3. For a weekly trading horizon, the returns are about the same as for the equally weighted portfolio, only slightly better. For monthly and quarterly

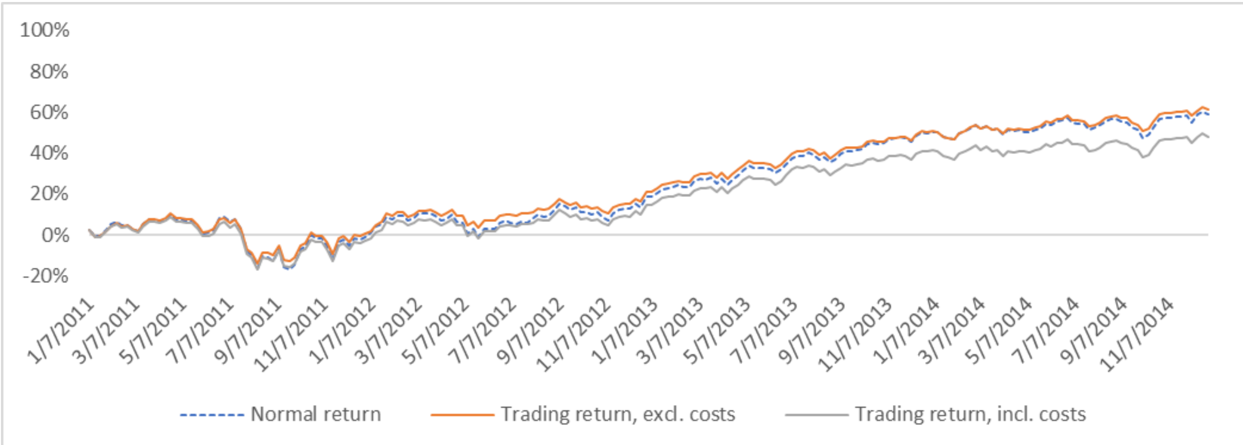

(a) Weekly trading horizon, n=1

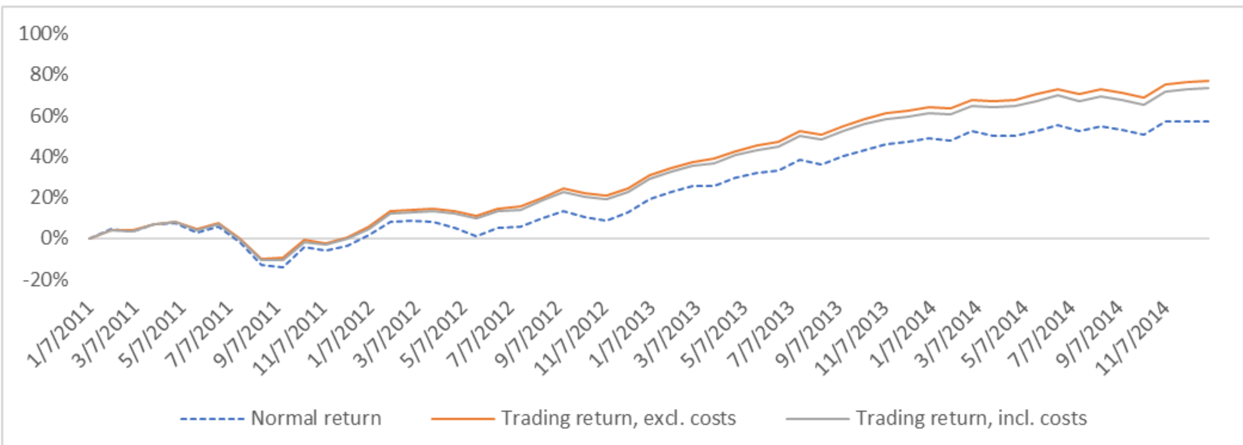

(b) Monthly trading horizon, n=4

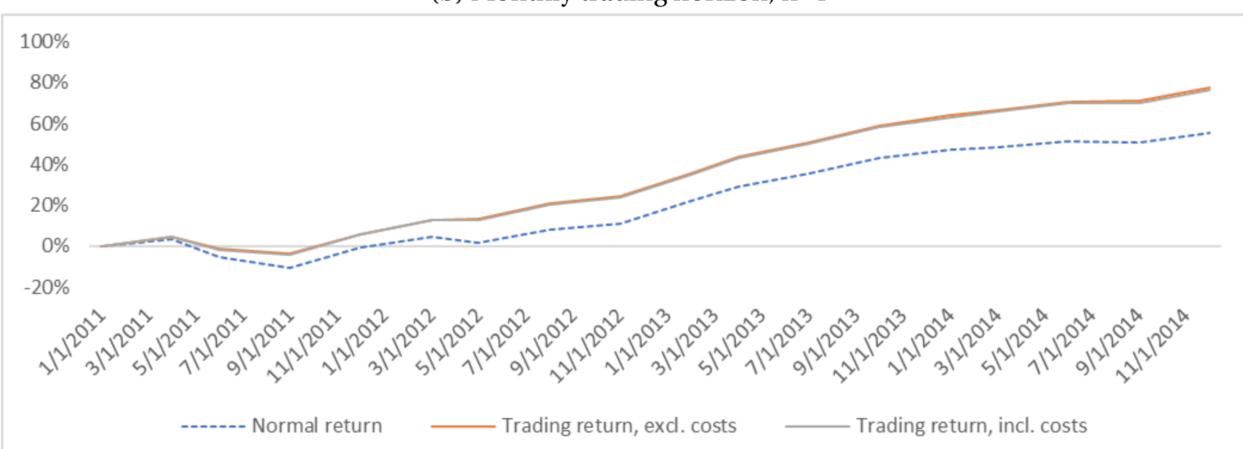

(c) Quarterly trading horizon, n=12

Figure 5.2: Cumulative returns for trading strategy where the 50% of stocks with highest predicted abnormal return are held, using the ANN prediction model, from 2011 through 2014

Figure 5.3: Cumulative excess returns over the equally weighted portfolio, excluding transaction costs, for trading strategy using ANN prediction model from 2011 through 2014, where the 50% of stocks with highest predicted abnormal return are held. Trading costs of 0.10% of traded value are included. Weekly, monthly and quarterly trading horizons.

Figure 5.4: Cumulative excess returns over the equally weighted portfolio, including transaction costs for trading strategy using ANN prediction model from 2011 through 2014, where the 50% of stocks with highest predicted abnormal return are held. Trading costs of 0.10% of traded value are included. Weekly, monthly and quarterly trading horizons.

horizons, the return is consistently above the equally weighted portfolio throughout the entire test period. It is clear that the prediction model performs consistently better for longer trading horizons than for shorter horizons.

Further, the cumulative excess returns for the ANN over the equally weighted portfolio including transaction costs are shown in Figure 5.4. For a weekly trading horizon the return is now below the return of the equally weighted portfolio throughout the period. For monthly and quarterly horizons the excess returns are still positive, indicating significantly better performance than the equally weighted portfolio, even with transaction costs.

5.3.1 Volatility

It is interesting to note that the portfolios based on the prediction models all provide lower volatility than the equally weighted portfolio, as evident from Table 5.4. Thus, although the linear regression model and GAM provide lower returns, the portfolios are more robust than the equally weighted portfolio. The Sharpe ratio captures the strong performance of the ANN models, with both higher return and lower volatility than the equally weighted portfolio, giv-

		Excl. transaction costs			Incl. transactions costs		
Horizon	Strategy	Return	Volatility	Sharpe Ratio	Return	Volatility	Sharpe Ratio
Weekly	Inverse	14.2%	20.0%	0.70	10.7%	20.0%	0.53
	Original	15.3%	16.1%	0.94	12.0%	16.1%	0.74
	Equally weighted	14.8%	17.9%	0.82	14.8%	17.9%	0.82
Monthly	Inverse	8.5%	14.1%	0.59	7.6%	14.1%	0.53
	Original	17.8%	11.9%	1.48	17.0%	11.9%	1.41
	Equally weighted	13.2%	12.9%	1.01	13.2%	12.9%	1.01
Quarterly	Inverse	7.6%	12.1%	0.61	7.3%	12.1%	0.59
	Original	18.3%	9.1%	1.99	18.0%	9.1%	1.96
	Equally Weighted	13.0%	10.5%	1.23	13.0%	10.5%	1.22

Table 5.5: Annualized results with inverted versions of the trading strategy, where the 50% of stocks with lowest predicted abnormal return are held. Results for the original trading strategies are also presented along with the equally weighted portfolio for reference. Trading horizons of one week, n=1, one month, n=4, and one quarter, n=12. Trading costs are 0.10% of the traded value. The Sharpe ratio is calculated using the average US 1-year Treasury rate over the test period, which is 0.165%, as the risk free rate.

ing remarkable risk-adjusted returns. This is quite unexpected as financial theory tells us that a more diversified portfolio will, on average, be less risky and hence the equally weighted portfolio consisting of all stocks in the market is expected to have lower volatility than a portfolio consisting of a subset of all stocks.

In order to further investigate portfolio volatility of the trading strategies, we run the ANN models with an inverse trading strategy, where the 50% of stocks with lowest predicted abnormal return is held. The results from the inverse trading strategy is presented in Table 5.5, along with the results from the original trading strategy and the equally weighted portfolio for reference. Interestingly, the volatility for the inverted strategies is above the equally weighted portfolio for all horizons.

In the descriptive analysis we found evidence for a negative impact of stock volatility on abnormal returns. This effect may also be evident in the ANN models, possibly leading the original trading strategies to, on average, buy less volatile stocks.

		Excl. transaction costs			Incl. transactions costs		
Horizon	Model	Return	Volatility	Sharpe Ratio	Return	Volatility	Sharpe Ratio
Weekly	Quarterly input	15.2%	15.5%	0.97	13.0%	15.5%	0.83
	Original	15.3%	16.1%	0.94	12.0%	16.1%	0.74
	Equally weighted	14.8%	17.9%	0.82	14.8%	17.9%	0.82
Monthly	Quarterly input	18.9%	18.3%	1.02	17.3%	18.3%	0.93
	Original	17.8%	11.9%	1.48	17.0%	11.9%	1.41
	Equally weighted	13.2%	12.9%	1.01	13.2%	12.9%	1.01
Quarterly	Quarterly input	18.3%	9.1%	1.99	18.0%	9.1%	1.96
	Original	18.3%	9.1%	1.99	18.0%	9.1%	1.96
	Equally Weighted	13.0%	10.5%	1.23	13.0%	10.5%	1.22

Table 5.6: Annualized results using ANN models with quarterly input data for all trading horizons. The return for the original models and equally weighted portfolio is included for reference. The trading strategy involves holding the 50% of stocks with highest predicted abnormal return. Trading horizons of one week, n=1, one month, n=4, and one quarter, n=12. Trading costs are 0.10% of the traded value. The Sharpe ratio is calculated using the average US 1-year Treasury rate over the test period, which is 0.165%, as the risk free rate.

5.3.2 Input data horizon

It is not clear whether the quarterly model is superior due to the longer horizon for predictions or the longer horizon for input data, as we have implemented the models in a symmetric fashion, with weekly predictions based on weekly input data and so on. In an effort to identify the impact of each, we implement ANN models with quarterly input data for all prediction horizons, and run the trading strategies using these models.

The results from the trading strategy with quarterly inputs are presented in Table 5.6. For a weekly trading horizon, there is no significant change in return, however the volatility is decreased with quarterly input data. For a monthly prediction horizon, the return is considerably improved, and is actually tantamount to that of a quarterly trading horizon. On the other hand the volatility is slightly higher than for the model with monthly input data. It is evident that both the prediction horizon and the horizon of input data impact the performance of the model. However, the weekly abnormal return seems to be difficult to predict no matter the horizon of the input data.

5.4 Google search volumes

In order to examine whether the ASVI significantly contributes to better and more valuable predictions for the ANN model, we compare the results to a benchmark model, which does not use ASVI data. The results from model evaluation of the two ANN models are presented in Table 5.7.

Trading horizon	n=1		n=4		n=12	
Model	ASVI	Benchmark	ASVI	Benchmark	ASVI	Benchmark
MSE	0.0016	0.0016	0.0064	0.0064	0.0192	0.0193
HIT_RATE	0.5058	0.5041	0.5264	0.5185	0.5442	0.5333
$PRECISION^{pos}$	0.4964	0.4990	0.5251	0.5210	0.5619	0.5622
$PRECISION^{neg}$	0.5110	0.5119	0.5272	0.5226	0.5436	0.5417
$RECALL^{pos}$	0.4146	0.3911	0.3913	0.3820	0.3519	0.3401
$RECALL^{neg}$	0.5850	0.6112	0.6581	0.6570	0.7373	0.7468

Table 5.7: Model evaluation scores for the the ANN model with ASVI data and the ANN benchmark model, with trading horizons of one week, n=1, one month, n=4, and one quarter, n=12

Horizon	Model	Excl. transaction costs Return	Volatility	Sharpe Ratio	Incl. transactions costs Return	Volatility	Sharpe Ratio
Weekly	ANN with ASVI	15.3%	16.1%	0.94	12.0%	16.1%	0.74
	ANN Benchmark	15.5%	15.8%	0.97	12.2%	15.8%	0.76
	Equally weighted	14.8%	17.9%	0.82	14.8%	17.9%	0.82
Monthly	ANN with ASVI	17.8%	11.9%	1.48	17.0%	11.9%	1.41
	ANN Benchmark	17.7%	12.0%	1.46	16.9%	12.0%	1.39
	Equally weighted	13.2%	12.9%	1.01	13.2%	12.9%	1.01
Quarterly	ANN with ASVI	18.3%	9.1%	1.99	18.0%	9.1%	1.96
	ANN Benchmark	17.7%	9.1%	1.92	17.4%	9.2%	1.89
	Equally Weighted	13.0%	10.5%	1.23	13.0%	10.5%	1.22

Table 5.8: Annualized results for the trading strategy, where the 50% of stocks with highest predicted abnormal return are held, using the ANN model with ASVI data and an ANN benchmark model. Equally weighted return is also included for reference. Trading horizons of one week, n=1, one month, n=4, and one quarter, n=12. Trading costs are 0.10% of the traded value. The Sharpe ratio is calculated using the average US 1-year Treasury rate over the test period, which is 0.165%, as the risk free rate.

The differences between the models are slight, and somewhat mixed. There is no difference in the MSE for the weekly and monthly horizons, but marginally better with the ASVI model for a quarterly horizon. The hit rate exhibits a clear trend of enhanced performance with ASVI data, for all trading horizons, with bigger differences for longer horizons. Precision is overall slightly better with ASVI for monthly and quarterly horizons, but this cannot be seen for a weekly horizon. The recall metrics give somewhat mixed indications. Overall, it seems that the ASVI models may have a slight advantage over the benchmark models, and more so for longer horizons.

We further investigate the predictive value of the ASVI by running the trading strategy with the benchmark ANN model, and comparing the results to that of the ASVI ANN model. The results are reported in Table 5.8. Comparing the returns achieved by the ASVI models to the benchmarks, it is evident that with weekly and monthly trading horizons, the ANN model is not able to improve the returns earned by utilizing ASVI data. However, for a quarterly trading horizon, the return is indeed higher for the ASVI model.

Figure 5.5 shows the cumulative excess returns of the ASVI model over the benchmark model, which does not use ASVI data, for the different trading horizons. For a weekly trading horizon, there is no clear trend in the cumulative excess returns of the ASVI model, and as both the re-

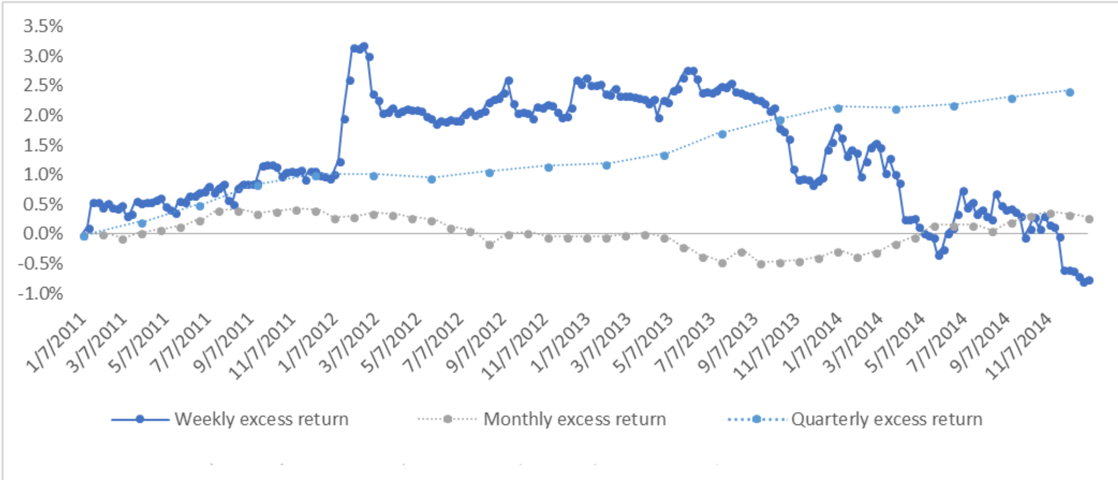

Figure 5.5: Cumulative excess returns for the ANN with ASVI data over the benchmark trading strategy without ASVI data, using the trading strategy from 2011 through 2014 where the 50% of stocks with highest predicted abnormal return are held. For weekly, monthly and quarterly trading horizons.

turn and volatility of the benchmark are better than with ASVI data, it is plausible that the weekly ASVI data is not providing any useful information and actually making predictions less accurate. For a monthly trading horizon, the cumulative excess returns of the ASVI model lie close to zero for the entire period, indicating that either the ASVI data is not providing any useful information or at least that this information is not able to enhance performance for this particular trading strategy. For a quarterly trading horizon we see a clear trend of positive excess returns for the ASVI model over the benchmark. It seems clear that quarterly ASVI data is able to enhance predictions to the extent of significantly improving the returns of the trading strategy.

Next, we examine in more detail the effect of the ASVI identified by the ANN. Figure 5.6 shows the estimated average impact of the ASVI on predicted abnormal returns for weekly, monthly and quarterly trading horizons. By holding all other variables constant and varying the value of the ASVI we record the impact on predicted abnormal return using the ANN. The curves are estimated by smoothed conditional means. The simulation data is generated using a random subsample of 1000 data points. For each sample 1,250 synthetic observations are generated, with ASVI values uniformly distributed in the range $[-5.50, 7.00]$, which is approximately the range of the ASVI in our population. It is important to note that we are solely interested in the shape of the curve and the magnitude of the change in predicted abnormal return for different values of the ASVI. The overall level of predicted abnormal return will be determined by the

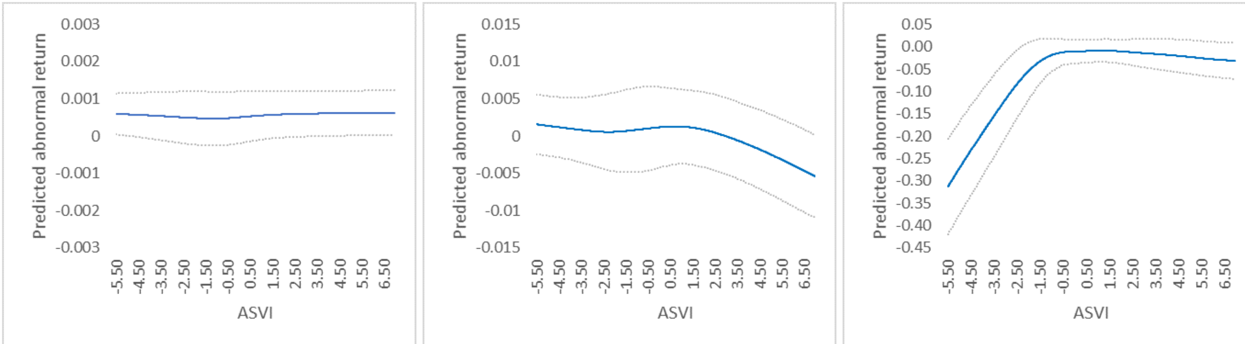

(a) Weekly impact (b) Monthly impact (c) Quarterly impact

Figure 5.6: Average impact of ASVI on predicted abnormal return, for different trading horizons. A randomized subsample of 1000 data points are used to generate simulation data. Grey dashed lines represent confidence intervals of one standard deviation.

value of the other variables.

From Figure 5.6a it is clear that for a weekly horizon the ASVI on average has no significant impact on predicted abnormal return. Thus, we find evidence that the ASVI over the past week has no impact on abnormal returns over the next week. From this we can conclude that for a weekly model ASVI data probably should not be included as an input variable, possibly explaining the poor results compared to the benchmark model. This is also in line with the descriptive analysis, where the ASVI was found to be not significant for a weekly trading horizon.

On the other hand, for longer horizons there seems to be some impact of the ASVI. Figure 5.6b shows a slight downward trend of predicted abnormal returns with increasing ASVI for a monthly horizon, where the largest decrease is for ASVI values larger than zero. Thus it seems that increasing values of ASVI in the past month leads to decreasing abnormal return, although the magnitude of the change is not massive. Although the monthly model was not able to achieve improved returns with ASVI data, there is reason to believe that the ASVI does provide some useful information for a monthly horizon, but that the magnitude is too small to make an impact with the trading strategy used.

For a quarterly horizon, Figure 5.6c shows that there is a substantial impact of the ASVI. This is also supported by the results of the trading strategy for a quarterly horizon, where the return was significantly improved with ASVI data. For negative values of the ASVI, a decrease is associated with decreased predicted abnormal return. However, for positive values of the ASVI the curve flattens and there seems to be little impact on predicted returns.

Period	2006–2007	2008–2009	2010–2011	2012–2014
MSE	0.0218	0.0530	0.0209	0.0188
HIT_RATE	0.5427	0.5564	0.5375	0.5512
$PRECISION^{pos}$	0.5577	0.5506	0.5486	0.5596
$PRECISION^{neg}$	0.5293	0.5642	0.5403	0.5490
$RECALL^{pos}$	0.6989	0.3243	0.3161	0.3999
$RECALL^{neg}$	0.3844	0.7693	0.7552	0.6986

Table 5.9: Model evaluation scores for the the ANN model for different time periods, with a trading horizon of one quarter, n=12. The time periods are: Before the financial crisis (Jan 2006 – Dec 2007); during the financial crisis (Jan 2008 – Dec 2009); first part of recovery (Jan 2010 – Dec 2011); market normalization (Jan 2012 – Dec 2014).

5.5 Time periods

We also analyze the performance of the ANN for four different time periods: Before the financial crisis (Jan 2006 – Dec 2007); during the financial crisis (Jan 2008 – Dec 2009); first part of recovery (Jan 2010 – Dec 2011); and market normalization (Jan 2012 – Dec 2014).

Results from model evaluation for the different time periods are presented in Table 5.9. The performance seems to be best in the periods 2008–2009 and 2012–2014. Interestingly, the precision and recall of negative predictions are considerably better during the financial crisis. The metrics for the period 2008–2009 indicate that the model generally predicts many samples as negative in this period, and although only 32% of positive samples are identified, the precision of the positive predictions is quite good. This is likely a valuable feature in a financial crisis.

Results from the trading strategy for the different time periods are presented in Table 5.10. The best results are achieved during the financial crisis, for the time period 2008–2009, almost doubling the return of the equally weighted portfolio, in addition to lowering the volatility. The performance is lower for the time periods after the financial crisis than the periods before and during. However, the performance is still satisfactory for all of the time periods, with consistently higher returns and lower volatility than the equally weighted portfolio. Moreover, the ASVI model outperforms the benchmark model for all periods, further underpinning the significant impact of ASVI on predictions for a quarterly horizon.

		Excl. transaction costs			Incl. transactions costs		
Horizon	Model	Return	Volatility	Sharpe Ratio	Return	Volatility	Sharpe Ratio
2006–2007	ASVI ANN	11.3%	9.8%	0.67	11.0%	9.8%	0.63
	Benchmark ANN	10.1%	9.9%	0.54	9.4%	9.9%	0.47
	Equally weighted	3.0%	10.8%	−0.16	3.0%	10.8%	−0.16
2008–2009	ASVI ANN	16.8%	33.1%	0.46	16.4%	33.1%	0.45
	Benchmark ANN	16.0%	33.0%	0.44	15.3%	33.0%	0.42
	Equally weighted	9.6%	36.3%	0.22	9.6%	36.3%	0.22
2010–2011	ASVI ANN	14.2%	17.2%	0.80	13.8%	17.2%	0.78
	Benchmark ANN	13.6%	16.8%	0.79	12.9%	16.8%	0.75
	Equally Weighted	9.5%	16.8%	0.55	9.5%	16.8%	0.55
2012–2014	ASVI ANN	19.0%	5.6%	3.40	18.7%	5.6%	3.33
	Benchmark ANN	18.3%	5.2%	3.48	17.7%	5.2%	3.36
	Equally Weighted	17.8%	5.8%	3.05	17.8%	5.8%	3.05

Table 5.10: Annualized results for different time periods, using the trading strategy where the 50% of stocks with highest predicted abnormal return are held. The time periods are: Before the financial crisis (Jan 2006 – Dec 2007); during the financial crisis (Jan 2008 – Dec 2009); first part of recovery (Jan 2010 – Dec 2011); market normalization (Jan 2012 – Dec 2014). Trading horizons of one quarter, n=12. Trading costs are 0.10% of the traded value. The Sharpe ratio is calculated using the average US 1-year Treasury rate over each time period as the risk free rate, which is 4.755%, 1.575%, 0.310% and 0.130%, respectively for each time period.

5.6 Validation

The final step of the backtesting procedure is validation of the results. We do this in two parts. Firstly we run a randomized trading strategy and use this to evaluate the robustness and significance of the results from the prediction model trading strategies. Secondly we run our analyses on the validation set, consisting of data from 2015, which has been held out from testing. The results from the validation set will affirm the out-of-sample performance and reveal any overfitting problems.

5.6.1 Robustness

In order to examine the robustness of the models, we compare the results to that of a randomized trading strategy, in which 50% of the stocks are chosen randomly for purchase each period, and the remaining 50% are sold. We run the strategy 100 times and approximate the distribution

(a) Weekly trading horizon, n=1

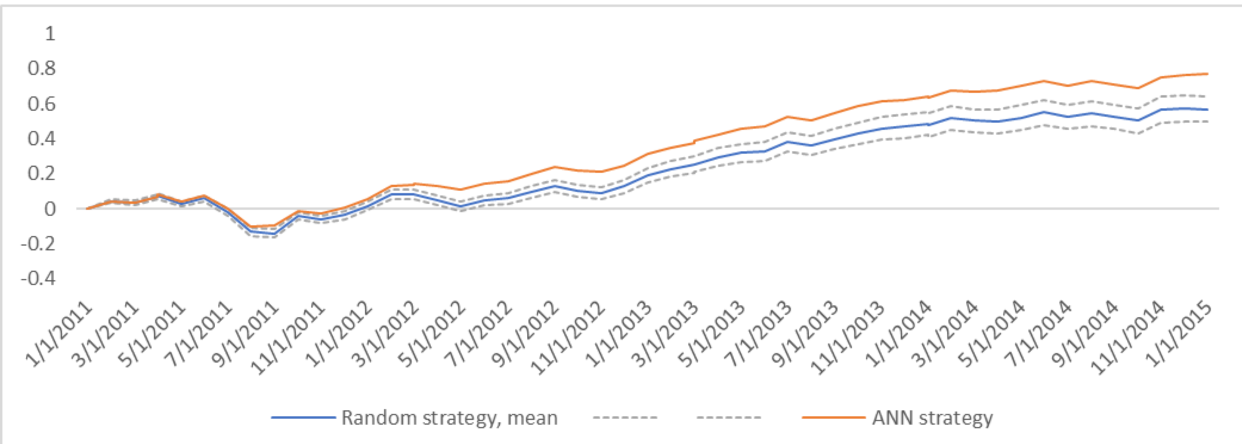

(b) Monthly trading horizon, n=4

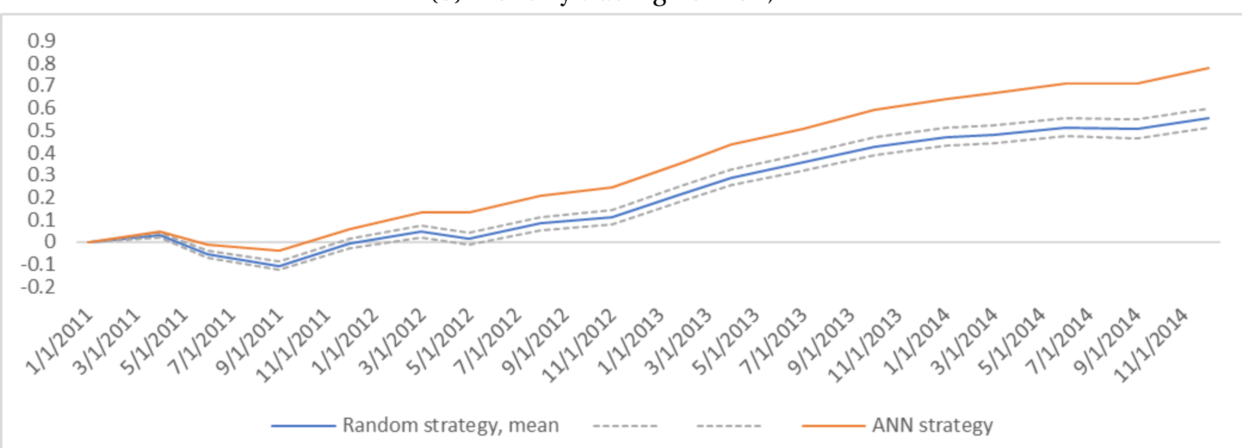

(c) Quarterly trading horizon, n=12

Figure 5.7: Cumulative returns for the ANN strategy, in which the 50% of stocks with highest predicted abnormal return are held, and the randomized strategy, in which 50% of stocks are chosen randomly to be held. The returns are reported excluding transaction costs. The orange line represents the mean return of the random strategy, and the grey dashed lines represent a confidence interval of three standard deviations.

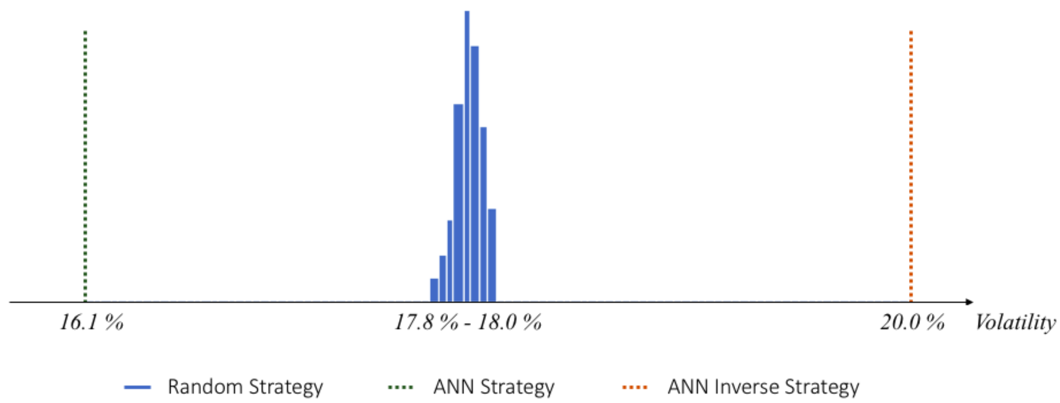

(a) Weekly trading

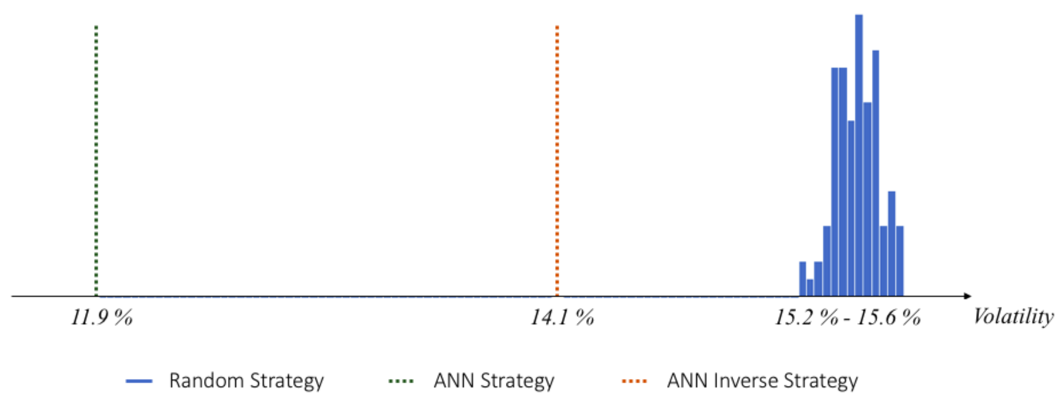

(b) Monthly trading

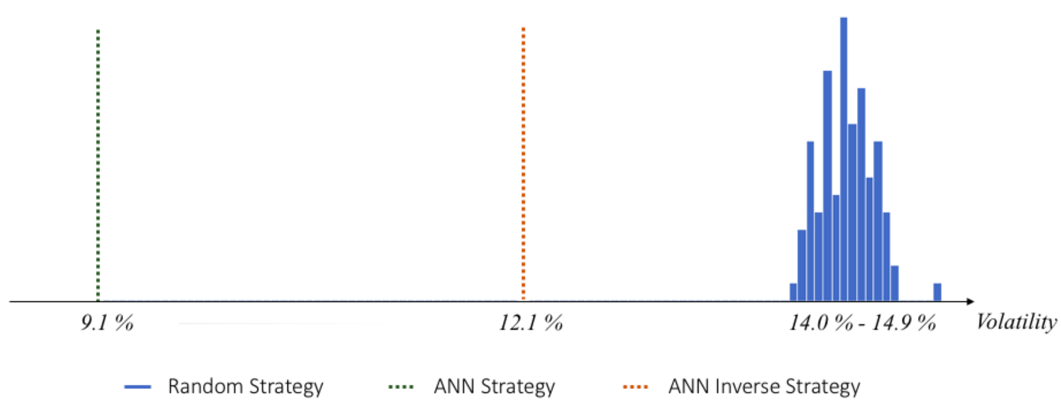

(c) Quarterly trading

Figure 5.8: Volatility distribution for the random strategy, along with the volatility of the ANN strategy and the inverse ANN strategy.

of returns and volatility.

Figure 5.7 shows the cumulative return of the ANN models plotted against the average cumulative return for the randomized trading strategies, with confidence bands of three standard deviations. As expected, the mean return of the randomized trading strategy is the same as the equally weighted portfolio. The weekly trading model is less than two standard deviations from the mean of the random strategy, and thus the weekly results are not very robust. In fact, the probability of achieving equal returns or higher with a randomized strategy is approximately 9%. However, for monthly and quarterly trading models, the return is well beyond three standard deviations from the mean of the randomized strategy, meaning that the probability of achieving comparable results by luck is vanishingly small.

Figure 5.8 shows the distribution of the volatility for a randomized portfolio, with lines representing the volatility of the ANN strategy and the ANN inverse strategy. For volatility, we do not necessarily expect the randomized portfolio to resemble the equally weighted portfolio. Since there is less diversification in the randomized portfolio, it should on average have somewhat higher volatility than the equally weighted portfolio, and this effect is evident for the monthly and quarterly trading horizons. The ANN strategy is well below three standard deviations from the mean volatility of the randomized strategy, for all trading horizons, and it is clear that the probability of achieving such low volatility by luck is practically zero. Furthermore, it is interesting to note that for the monthly and quarterly trading horizons, even the inverse strategies achieve lower volatility than the randomized portfolios.

We conclude that the results are robust and significant, and that while the return of the weekly model is not significantly different from random stock picking, the volatility is decreased.

5.6.2 Validation set

Lastly, we run the trading strategy on the validation set, consisting of data from 2015. Results from the trading strategy with both the ANN model with ASVI data and the benchmark ANN are presented in Table 5.11. We find the same trends as for our testing set. There are increasing returns for longer horizons with both models. Also, the volatility is lower for almost all of the ANN models than for the equally weighted portfolio across trading horizons.

The difference between the ASVI model and the benchmark for weekly trading is reversed

		Excl. transaction costs			Incl. transactions costs		
Horizon	Model	Return	Volatility	Sharpe Ratio	Return	Volatility	Sharpe Ratio
Weekly	ASVI ANN	3.1%	12.5%	0.23	0.0%	12.5%	−0.02
	Benchmark ANN	−2.3%	15.3%	−0.16	−11.9%	15.3%	−0.79
	Equally weighted	−3.4%	14.9%	−0.24	−3.5%	14.9%	−0.25
Monthly	ASVI ANN	3.4%	7.0%	0.46	2.4%	7.0%	0.32
	Benchmark ANN	1.0%	6.8%	0.12	−0.1%	3.3%	−0.13
	Equally weighted	−3.3%	8.6%	−0.41	−3.4%	8.5%	−0.43
Quarterly	ASVI ANN	3.9%	4.9%	0.75	3.5%	4.9%	0.68
	Benchmark ANN	0.1%	5.0%	0.12	0.0%	5.0%	−0.01
	Equally Weighted	−5.0%	7.4%	−0.71	−5.1%	7.3%	−0.73

Table 5.11: Annualized results for the validation period, 2015. Using the trading strategy, where the 50% of stocks with highest predicted abnormal return are held, for both the ANN with ASVI data and the benchmark ANN model. Trading horizons of one week, n=1, one month, n=4, and one quarter, n=12. Trading costs are 0.10% of the traded value. The Sharpe ratio is calculated using the 2015 US 1-year Treasury rate, which is 0.20%, as the risk free rate.

from that of the test set. However, this supports our conclusion that the weekly models are in general highly inconsistent, providing returns resembling that of a randomized trading strategy. Although the ASVI model provides higher returns than the benchmark for the weekly trading horizon on 2015 data, there is reason to believe that the returns are mainly a result of randomness, and that any conclusion that ASVI data improves the weekly model cannot be safely drawn. For monthly trading, the improvement with ASVI data is more substantial than with the test set. For quarterly trading, the improvement with ASVI data is again more drastic than with the test set, but indicates the same trends: Higher returns and equal volatility for the ASVI model as the benchmark. We conclude that the results of the predictive analysis are valid and robust.

Chapter 6

Conclusion

In this report, we investigate the predictability of stock returns, and examine whether Google search volumes can improve such predictions. We use weekly data for S&P 1500 companies from 2004 through 2015. Our aim is to shed new light on the relationship between Google search volume and stock returns, by utilizing Artificial Neural Networks (ANN) and a robust backtesting procedure for trading simulations.

We find that the ANN approach is superior to linear and semi-parametric regression in predicting abnormal stock returns. Various evaluation metrics indicate that the ANN prediction model provides the best performance. Trading strategies also reveal substantially higher returns for the ANN. For monthly and quarterly horizons, the ANN trading strategy is able to achieve both substantially higher returns and lower volatility than an equally weighted portfolio including all stocks. Further, we find that both the horizon of the input data and the prediction horizon impacts the goodness of the models, with more accurate predictions for longer prediction horizons and for model input that includes aggregate data over longer periods. However, for a weekly prediction horizon, we find that accurate predictions cannot be made even with quarterly input data.

By comparing the ANN trading strategy including Google search volume data with an ANN benchmark strategy, we find that Google search volume has significant predictive power of abnormal stock returns only for the quarterly trading horizon. The return for the strategy including Google search volume data outperforms the benchmark by 60 basis for a quarterly trading horizon. We further analyze the impact of Google search volume by holding all other variables constant and varying the value of Google search volumes to record the impact on predicted abnormal return. For weekly and monthly trading horizons, we find that the Google search volume does not have a substantial impact on predicted abnormal returns. However, for the quarterly trading horizon, we find that a decrease in search volume leads to lower predicted abnormal return, but that an increase in search volume does not impact the predicted abnormal return substantially.

The ANN trading strategy is further tested on different time periods to evaluate its performance under various market conditions. We find the performance to be satisfactory for all periods, with consistently higher returns and lower volatility than the equally weighted portfolio.

Finally, model validation of the ANN trading strategy is performed. First, we assure robustness by running a trading strategy based on buying and selling random stocks, and compare its performance against the ANN trading strategy. Using the resulting distribution of abnormal returns for the random strategy, we find that the performance of the monthly and quarterly ANN strategies are significantly better than the random strategy, however, the weekly ANN strategy is not. Furthermore, the volatilities of the ANN strategies for all trading horizons are significantly lower than the volatilities of the randomized strategies. Second, results are obtained for the validation set, consisting of data from 2015, confirming the validity of the test results and assuring there is no problem of overfitting.

For future research we recommend further analysis of input data horizons, in an effort to identify a horizon that optimizes predictions. Longer prediction horizons may also be analyzed, to determine whether performance can be enhanced beyond that of a quarterly trading horizon. It would also be of interest to analyze the impact of Google search volumes for longer horizons using artificial neural networks.

References

Amihud, Y. and Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17:223–249.

Aouadi, A., Arouri, M., and Teulon, F. (2013). Investor Attention and Stock Market Activity: Evidence from France. *Economic Modelling*, 35(C):674–681.

Balaban, E. and Bayar, A. (2005). Stock returns and volatility: Empirical evidence from fourteen countries. *Applied Economics Letters*, 12:603–611.

Bank, M., Larch, M., and Peter, G. (2011). Google search volume and its influence on liquidity and returns of german stocks. *Financial Markets and Portfolio Management*, 25(239).

Barber, B. M. and Odean, T. (2007). All that glitters: The effect of attention and news on buying behavior of individual and institutional investors. *The Review of Financial Studies*, 21(2):785–818.

Bijl, L., Kringhaug, G., Molnár, P., and Sandvik, E. (2016). Google searches and stock returns. *International Review of Financial Analysis*, 45:150–156.

Challet, D. and Bel Hadj Ayed, A. (2013). Predicting financial markets with Google Trends and not so random keywords.

Challet, D. and Bel Hadj Ayed, A. (2014). Do Google Trend data contain more predictability than price returns?

Da, Z., Engelberg, J., and Gao, P. (2011). In search of attention. *The Journal of Finance*, LXVI(5):1461–1499.

Freeman, J. D. (1992). Behind the smoke and mirrors: Gauging the integrity of investment simulations. *Financial Analysts Journal*, 48(6):26–31.

French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, 19:3–29.

Gervais, S., Kaniel, R., and Mingelgrin, D. H. (2001). The high-volume return premium. *Journal of Finance*, 56:877–919.

Grullon, G., Kanatas, G., and Weston, J. P. (2004). Advertising, breadth of ownership, and liquidity. *The Review of Financial Studies*, 17(2):439–461.

Hastie, T. and Tibshirani, R. (1986). Generalized additive models. *Statistical Science*, 1(3):197–318.

Joseph, K., Wintoki, M. B., and Zhang, Z. (2011). Forecasting abnormal stock returns and trading volume using investor sentiment: Evidence from online search. *International Journal of Forecasting*, 27:1116–1127.

Kingma, D. P. and Ba, J. (2014). ADAM: A method for stochastic optimization. *CoRR*, abs/1412.6980.

Kristoufek, L. (2013). Can Google Trends search queries contribute to risk diversification? *Scientific Reports*, 3(2713).

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521:436–444.

Leinweber, D. (2007). Stupid data miner tricks: Overfitting the S&P 500. *The Journal of Investing*, 16(1):15–22.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, XLII(3):483–510.

Ng, A. Y. (2004). Feature selection, L1 vs. L2 regularization, and rotational invariance. In *Proceedings of the Twenty-first International Conference on Machine Learning*, ICML '04, pages 78–, New York, NY, USA. ACM.

Paital, R. R. and Sharma, N. K. (2016). Do trading volume and bid-ask spread contain information to predict stock returns? intraday evidence from india. *Asian Economic and Financial Review*, pages 135–150.

Preis, T., Moat, H. S., and Stanley, H. E. (2013). Quantifying trading behavior in financial markets using Google Trends. *Scientific Reports*, 3:1684.

Preis, T., Reith, D., and Stanley, H. E. (2010). Complex dynamics of our economic life on different scales: Insights from search engine query data. *Philosophical Transactions, The Royal Society A*, 368:5707–5719.

Russel, S. J. and Norvig, P. (2010). *Artificial Intelligence: A Modern Approach*. Pearson Education.

Seasholes, M. S. and Wu, G. (2007). Predictabe behavior, profits, and attention. *Journal of Empirical Finance*, 14:590–610.

Takeda, F. and Takumi, W. (2014). Google search intensity and its relationships with return and trading volume of japanese stocks. *Pacific-Basin Finance Journal*, 21:1–18.

Wood, S. N. (2006). *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC.

Zeiler, M. D., Ranzato, M., Monga, R., Mao, M. Z., Yang, K., Le, Q. V., Nguyen, P., Senior, A. W., Vanhoucke, V., Dean, J., and Hinton, G. E. (2013). On rectified linear units for speech processing. In *ICASSP*, pages 3517–3521. IEEE.

Appendix A

Company tickers used as Google search terms

JJSF; PLXS; HGR; SUNW; UCI; KDE; ORCL; MSFT; SDS; AYE; TROW; OSIP; HON; EMC; XRIT; TECD; RGEN; FO; INDB; SIGM; TKLC; HVT; LLTC; CY; TMP; BCP; GENZ; IMDC; ADPT; TCB; ASTE; NAVG; WERN; T; IRN; SKYW; AMWD; ADM; RNT; VIVO; SOV; WTS; FNB; TGX; ACET; AD; BRLI; ASCL; FISV; HTLD; TSFG; BWS; BC; UIS; CERN; WBS; CMVT; CDR; FBP; NWK; STSA; DELL; AVD; PLAB; MESA; OFG; CTR; KO; SAFM; UBSI; CCC; SKS; WLM; CDN; ED; BRL; CHCO; EBIX; HRH; FILE; FIC; CNMD; CELG; ATU; ACO; XRAY; FAST; CPF; RBNC; AIRM; CRK; GDOT; DTE; ALEX; DD; CBM; AIN; IIVI; EK; SIVB; CLHB; PBKS; XOM; IPAR; MXIM; PTEC; RESP; WMI; APCC; BMC; LENS; DY; GMT; GD; GE; LH; NVLS; ONB; PKY; PBCT; PDE; SIGY; GM; COR; GR; GNCMA; LNN; MSS; NOVN; EYE; EXAM; LYB; IR; FXCM; FRC; MMI.3; WD; IBM; NAV; AAT; INN; ITT; FF; HII; KSU; SGNT; ECOL; LG; RATE; MPC; AMCX; SXC; SYRG; FRAN; TNGO; FBHS; XLS; XYL; LMOS; AM; VAC; MAY; BCEI; MYG; WPX; KORS; TRIP; POST; SLCA; BNNY; PSX; TUMI; WAGE; FB; SUPN; ALEX; TROX; EGL; FIVE; ADT; KRFT; WWAV; PCG; UIHC; ABBV; ERA; TPH; BCC; ZTS; AVIV; PGL; ENTA; CST; PEP; MO; AMBC; COP; FNBC; PVH; NWSA; AMGN; MNK; MUSA; SAIC; ALOG; FTD; SLB; ALLE; S; ANDW; OGS; INGN; IBP; KN; APOG; CVX; AAPL; PAHC; CVEO; DNOW; TIME; WPG; CTRE; SUN; RYAM; TMST; MIK; AMAT; CTLT; SYF; PGN; TKR; HQY; TR; RYI; UCL; CDK; HYH; ENVA; ABCW; CRC; MRO; UST; VMC; ASBC; FL; WWY; ATG; RSH; TXN; WYE; EIX; ATW; AZZ; G; GT; BOH; UVV; HSY; KR; BN; CVS; BSET; GIS; MHP; BER; KMB; PD; UTX; UGI; BGG; HPC; SVNT; CW; BMET; PG; JCP; SO; CAT; BOBE; BCO; CL; BRC; FMC; DE; BMY; WAG;

BA; VVI; TRY; LABS; SNV; ACXM; CCBL; CR; BDK; CPY; CVBF; LUK; ABT; CAI; CALM; CRS; DOW; CAMP; GY; AMR; GCO; FOE; LMT; MWV; NEM; WGL; CAH; IP; CAE; CASY; EXC; CNP; IDA; PFE; CSK; CBE; CBSS; EMR; JNJ; GLW; CRDN; CHG; PPG; PPL; MMM; MZ; MRK; CHRS; MOT; CMN; SLE; DPL; CHTT; CKP; CIN; FE; HNZ; SCG; PGN; CMS; EAT; HLT; CHIR; CHZ; CHD; CINF; POM; WEC; FJC; TXT; CTAS; PEG; HAL; MDU; CZN; GP; CYN; XEL; CLC; ROH; ETR; CCU; WR; EAS; AEP; FPL; CEG; ALE; ASH; EQT; KSE; GXP; OGE; PBI; COHR; TXU; CNB; AA; NOC; RTN; AVA; AEE; SGP; DQE; CMA; WPS; CBSH; OKE; HSC; CPB; CTCO; WHR; PSD; CD; HRS; NFG; KMG; CA; F; DOV; ATSN; CMTL; MEE; DIS; EGN; BOL; ILA; SNS; SJI; GAP; LTR; K; VAR; STL; CSL; CLF; CTB; CUZ; CBRL; R; AVT; IRF; ABI; STR; SRR; HPQ; BAX; CFR; DUK; XRX; PNW; APD; MUR; VNO; AHC; CTS; ALK; DSCP; NBR; KTO; ARW; AIT; DP; DMN; BF; DNEX; CAS; COHU; CMC; OMC; DBRN; FLS; CUB; EV; EDO; ESIO; ELMG; GLK; HP; WFT; HOC; HRL; HUB; ATO; ENZ; MAS; NPK; NR; FITB; NUE; OXY; OXM; RF; PLL; PKE; FAF; FMER; PBY; TRMK; MTB; FFBC; FINB; FMBI; ROG; ROL; FTN; SHW; WB; SCL; TE; LDR; FST; TSO; FELE; BEN; FMT; VAL; FUL; AJG; AGI; WMB; CK; TNB; SRP; GNTX; DNY; WFC; GEOI; NI; SKY; CEN; PAS; FON; SMP; APA; MAT; GGG; BDX; B; WY; TEK; CSC; IFF; AVP; DBD; TJX; TXI; CMI; HE; HWKN; HCSG; PH; CHB; AT; HELE; HELX; HIB; CNF; BCC; ACV; PNK; PKI; WWW; NX; HNI; HH; JCI; EY; SJM; MDP; JBHT; HBAN; ATI; SWK; MCD; SXI; VFC; LZ; BMS; KMT; ELK; NU; INSU; TFX; IDTI; TNL; AVY; INTL; ADP; IFSIA; IMGC; DIOD; SVU; TII; FRX; IGT; IDCC; TYC; RDC; AGL; CDI; SFA; IVC; MMC; SII; SWKS; AXE; JLG; FLE; ITG; GFF; CLX; GPC; CBT; HMT; BCR; KLAC; KAMNA; CKR; CC; KDN; KEA; KELYA; NYT; RML; BNE; CBB; FSS; ABM; GRB; JPM; GCI; KLIC; LSI; GAS; LZB; UNS; JP; LRCX; MCO; LANC; LNCE; HUM; TTN; UNP; KRI; LAWS; LNC; TGT; HRB; DDS; LXU; BK; DHR; PCH; LIZ; PNY; RRC; ABS; LFB; BNI; URS; LEE; TDW; MTSC; MRD; PEI; ESL; LLY; JH; ADCT; SMTC; MEG; FNM; WGO; MTW; TER; NSM; MCS; OSG; KMI; SRV; VVC; PIR; MI; SYY; RHI; JEC; THC; SUP; EFX; GWW; HB; MNT; MENT; WTR; MER; MCY; HAS; IPG; BMI; METH; WPO; GDW; MCRS; ATX; LDG; MU; MSCC; UIL; CTX; FDO; MLHR; PHM; MIL; MSA; DSL; AGE; AIR; MODI; MOGN; RDK; MOLX; TRN; SPW; RI; PNM; BW; BDG; WMT; NAFC; FTO; POP; LPX; NCC; CAG; NDC; ITW; NPBC; COA; FRK; NTY; NATR; CGC; TNM; GLT; KWD; NEWP; BLL; NKE; VSH; NDSN; JWN; AFL; NFB; NTRS; TEX; NWN; FRT; STK; ASE; LUV; BEZ; RBC; LNT; OII; DJ; GPS; KEX; OLG; OCAS; AXP; BUD; CB; UDR; INTC; MEL; ORI; BAC; SAFC; SPC; YELL; SXT; LDL; MDT; OMI; SNA; PNC; PCAR; TTC; CTL; FDX; AFG; PPP; ADI; LEG; NWL; TOY; PRX; SWX; AMD; DCI; LOW; BIO;

PAYX; AOC; DLX; MOG; NBL; BRE; BKH; PVA; PNR; PENX; TMO; CSX; TMK; PETD; GTY; RYL; WST; ARXX; ASO; EC; SENEB; WEN; AGYS; PDC; BRO; CV; POSS; OCR; POWL; SWN; UNT; PPD; PCP; GMP; PLFE; WMS; CI; RRR; LTD; NSC; PGR; NJR; ASN; PL; CNL; CFC; PSA; KWR; D; QSII; KEY; ANSI; DRS; HOV; TRA; RLI; UNF; RPM; LM; MYE; AOS; BLC; MBG; COO; TRB; ALO; VZ; BLS; HU; HCN; RUS; SBC; TIN; BOW; USB; HD; MDC; SLM; WSO; WDC; HUG; SFY; REY; AIG; RIGS; MSC; CORS; BCF; RBN; CLE; BLI; HCP; SIE; RGLD; NDE; CMO; RYAN; STI; WRI; SEIC; SIGI; BSC; NHP; VRX; STJ; RSC; CVC; MWD; PR; SHLM; MYL; RJF; STZ; HDI; KBH; TOL; ARG; APC; CCE; C; SIAL; ECL; SYMM; SFD; ENC; UNM; SON; Y; AWR; BBT; CHP; FNF; SPAR; SMSC; SR; SPF; STT; SQA; UHT; SVC; STC; SF; SYK; RGR; SUPX; SUSQ; SBL; TBCC; VOXX; BHI; HKF; INT; PXR; TIF; TBL; VIA; ABMD; CCL; FED; MBI; PII; SCH; PWN; GGC; HAR; IVX; PXD; TLAB; NSS; MAN; BID; ACI; BR; DRE; MTH; KE; BEC; CEC; LYO; TNC; SPLS; ADBE; MRTN; AFAM; ZQK; IMR; ODP; ALTR; IEX; PCL; MAG; CRUS; SYMC; AGN; TG; CGNX; TECH; FRE; VTR; EOG; ERTS; DGII; ZIXI; LSCC; SLR; ACF; LSCP; WGR; AZR; GIII; PMTC; MAFB; NEOG; CSCO; THO; COG; RDA; HET; SUG; HOLX; DFG; THOR; VRC; MATR; TTI; VICR; COMS; GVA; UTR; BEAV; ORB; RTI; SWY; HCA; TYL; KSWS; XLNX; ACAT; PKS; SWFT; CAND; WTSLA; TRH; CHUX; BHE; TRMB; BJS; EXPO; VITL; VPI; MTRX; AN; ASHW; ISYS; MRX; MMSI; TFS; ATK; ESE; LUFK; CATY; KRB; HMA; SONC; MLI; RBIN; ATML; MVK; PSB; AZO; IO; REGN; CVH; PMCS; CEPH; JNY; TSS; X; AHG; MEDI; ANN; HBHC; PNRA; HNT; OSI; DV; IDXX; OSTE; RGS; AES; RHB; CURE; CRVL; IHP; VRTX; MNRO; PRGS; VTS; AMHC; SY; ZBRA; BMHC; EZPW; MTG; HAE; PRA; BIIB; SEPR; FCFS; AAON; AW; THQI; AXYS; LFG; PSTI; ODFL; SFP; SWS; WTNY; HCR; SKO; YRK; ANK; IAAI; WRLD; WNC; HMN; KIM; FSH; OI; QCOM; PRGO; TTEK; BDY; OFC; SBSE; GILD; ABAX; WFMI; NXTL; PDLI; SMG; ROP; USFC; SCHL; SNPS; LIFC; ELY; PLMD; STAR; NBTB; TWX; FRED; LNCR; BSG; ICUI; RARE; JBX; VCI; FD; POS; CBK; EP; NCS; KOPN; MW; OPTN; TRST; AG; STBA; SMRT; FDC; REM; BKE; ABFS; BSX; KSS; JNC; FIF; USPH; FCF; STE; BBBY; DHI; KRON; ESRX; FINL; COLB; SBUX; TSN; SPN; PX; WIRE; ZOLL; GTK; HITK; FBN; OHI; LTC; CAKE; ASGN; LFUS; CSAR; EFII; NET; BRKT; KEM; EPIC; SBIB; MTX; HCC; PDCO; RDN; UFI; QSC; BLTI; LGND; TCO; CKH; CPWR; HMSY; SM; BBOX; PQUE; PMTI; RSCR; SHFL; TDSC; UFCS; APPB; UMBF; IACI; AGY; CREE; CRY; CHK; CYBX; UEIC; WPI; DDR; APSG; AVID; CHS; GYMB; INTU; JAH; LSTR; MCHP; PSUN; ETH; ACE; ATN; FOSL; JBL; ORLY; SANM; WIND; GGP; ATR; RCL; ANIK; CDWC; PERY; TWI; RGA; XTO; HR; AKR; RIG; SKT; ARGN; BELM; DLP; FLIR; PZZA; USTR; ALL; DFS; LBY; SGMS; ACAI;

PETM; NYFX; SGY; PPS; CPT; ACTL; LNY; MCRI; PFGC; SPSS; NXL; BWA; EQR; BFS; TCT; ARRS; GMCR; MNC; MOV; PHLY; UTEK; WLV; UHS; CLP; BWC; ZLC; ATVI; CMOS; DECK; IT; MERQ; OMG; TGIC; ULTE; BYD; LXP; REG; NVR; AF; ATMI; JILL; FFIN; ITRI; FWRD; MATK; PTEN; NYB; ROCK; UFPI; JDSU; URBN; AEC; HHS; MHO; NFX; ALSC; CGI; GTW; MACR; PLNR; SGR; SHOO; TQNT; MED; VRTS; VLY; CYT; EMN; SPG; CACB; CLDN; FFFL; VMI; HARB; SFN; HAIN; GBP; MAA; PLT; PYX; ALB; BZH; MLM; RYN; BPFH; DAKT; DSPG; MAPS; MINI; QLGC; SNIC; TSCO; EMMS; AKS; CBR; CPRT; RKT; SCSC; SHU; TRBS; ZEUS; AVB; MAC; ADC; AEOS; CELL; GDI; MROI; ABCB; FRNT; NKTR; JOSB; PSSI; SSD; BTH; LEH; CGX; AAI; ESS; HIW; LRY; AIV; GBX; HME; MATW; GGI; CLI; ADTN; AMED; MPS; ASVI; MOVI; TWR; RS; ACS; DAR; KNGT; SSYS; O; COF; MCK; ALAB; NWSB; PCLE; VECO; GDT; ESI; APOL; CAPA; MCRL; SWC; RSAS; VOL; NSIT; RCII; BRKS; IAAC; CMX; TSAI; ASD; SOL; MVL; WDFC; DLTR; NWRE; NATI; PMI; ABC; MBFI; WFSL; WM; HLIT; BGP; DRI; HA; SSS; WAB; LECO; DRTE; WMO; FEIC; RMD; FCX; WFR; CAM; AVZ; HPT; WABC; IART; DNR; CKFR; PDX; SMSI; AFC; CPC; DO; RE; DVA; ESST; POOL; RSYS; VRTY; ADVS; AEIS; HYSL; BANR; HSIC; IFIN; IVAC; NTAP; PHTN; PRXL; SCUR; SNDK; SAM; BKI; EL; LXK; MIG; SWM; WAT; WON; CTXS; EME; EAGL; IPXL; KNSY; HIG; MSM; GLB; BFT; SRCP; LCAV; PPDI; EE; JW; WSM; WL; ALXN; ANDE; ARTC; CNCT; CSGS; APN; IRM; RCI; DST; AH; BBX; CENX; CYTC; GEF; HUBG; JDAS; PRGX; LU; PSS; ANST; CPWM; CRR; PLCM; DF; SYKE; YHOO; BRK; NDN; POG; TUP; DIGE; FCN; JAKK; MGAM; MLNM; SRZ; EDS; FDS; WOR; IMN; AFFX; ANSS; BCGI; DPMI; GWR; KEYS; LRW; SEBL; TIE; RNR; STRA; VMSI; DCOM; GES; ET; LAMR; MTON; SRCL; TTEC; CUNO; CYMI; HOTT; LTBG; SPPI; RSTI; ANF; UVN; USNA; RX; MWY; NOI; TSG; TGI; ARQL; CBST; CHGO; ZNT; HIBB; MBRS; NCI; ZION; TALX; IM; LIN; SSP; FORR; ISCA; PLUS; NCOG; STLD; UNFI; VPHM; ATAC; LAD; IGTE; PWAV; MNST; VSAT; CVD; DEL; NCR; DGX; ROK; ASF; KRC; SKP; WGOV; AMH; CWTR; SWBT; BSTE; CIEN; EPIQ; ERES; RADS; SLGN; RAH; GTRC; MVSN; WTFC; FBC; GIFI; MEAD; TTWO; ARE; APH; BGC; AMZN; RMBS; NCEN; Q; RCRC; RFMD; BXP; FIX; MMS; RL; PVN; SRT; EGP; CDIS; CSTR; DXPE; EGHT; OZRK; CTV; BJ; CXW; EOP; JLL; CPS; AME; SLG; VLO; ACAS; CRZO; KNDL; PEGS; YUM; ARM; FARO; PLCE; PLB; DRQ; GPI; UNA; CCRD; CHRW; TRAD; OSIS; PHCC; PSEM; SCMM; PWER; AMCC; BFAM; OYOG; AMB; AMG; EPR; MTD; SAH; ADSK; CPO; URI; AMSG; DEPO; POWI; CECO; MRCY; NARA; VRSN; MDS; BXS; PWR; CARS; GYI; BRKL; IBOC; ICBC; ISSX; SRDX; MAR; BBY; SEE; WDR; PRV; LHO; BRCM; FTBK; LVLT; MANH; FNFG; SONO; UMPQ; EQY; LLL; OTL; HDWR; CANI; SBSI; WCN;

FII; AMT; HZO; SRE; CTSH; GISX; KG; TRDO; SCHS; ULTI; MSTR; RSG; CFNL; ELBO; CVG; UBA; XL; DRIV; ECLP; CCI; EBAY; PRSP; UCBH; WSBK; CZR; SCSS; CNQR; PFCB; INSP; ARJ; ETM; NVDA; DPH; KFY; AMRI; CATT; COCO; EPAY; EWBC; EXBD; IVGN; MTEX; SRNA; CTCI; PBG; KWK; MKSI; PCLN; VARI; VSEA; CNX; SFG; ACDO; HSII; INFA; CBH; CFB; TUES; GS; NLS; CBU; LPNT; ROIA; TRI; TWTC; RJR; SKX; CYBS; DITC; FFIV; JNPR; STMP; UTHR; PVTB; LII; DSS; COST; EGOV; HCBK; JCOM; ECPG; MDRX; PRFT; TIBX; DVN; FCS; LAB; NST; TOO; NTCT; QSFT; RHAT; OMN; FDRY; SKE; VTIV; UNTD; ININ; BLK; CIR; NMG; SNH; PTV; AKAM; CHIC; DGIN; IWOV; PCTI; BCSI; RETK; RTEC; DW; A; TDY; UPS; HLEX; NTG; PKG; EASI; EW; ENR; GRP; EXAR; EXPD; DTAS; GTIV; LMNX; PLMO; SLAB; UTSI; WBSN; MET; GBCI; HANS; CCMP; FWHT; HSTM; KKD; LPSN; FULT; VDSI; VCLK; ICST; CRL; CYH; JNS; VC; EFD; ISRG; NNN; ACLS; AVCT; ENDP; ILMN; VASC; WEBX; AMMD; CPKI; MDCO; OPNT; SPTN; AV; DNB; POL; GB; HYDL; TTMI; COH; JKHY; MON; SFCC; OPLK; WHQ; XXIA; BKMU; HTCH; NVTL; SLXP; UTIW; SYD; ACAP; GRMN; RECN; AET; FLR; ALGN; HAFC; BLUD; PEET; MCF; GPN; OIS; IBCP; EAC; FLO; ROXI; BTU; ADS; CEY; FTI; KFT; COL; TASR; USPI; BABY; MDTH; ZMH; ACN; JOYG; OMCL; MKC; CCRN; ODSY; ATH; PFG; AGP; KIND; WOOF; AHS; AYI; AAP; DJO; IGI; IFC; VAS; PRU; CNC; STGS; SYNA; GME; CG; MANT; UCBI; ANT; KEI; JBLU; LCI; ATRS; CPSI; KYPH; NFLX; PNFP; ARO; NPO; SRX; BGFV; AVO; HEW; CVGW; KIRK; RRGB; CIT; PSYS; ARB; GPRO; SCST; KMX; XEC; CMCSA; TAYC; WYNN; DKS; MRH; CEVA; PRAA; SAFT; SINT; EQIX; CME; PXP; STX; PFS; LOJN; LSS; IPCC; KOMG; WRNC; ENH; LEN; IPMT; KMRT; MOH; DTSI; NTGR; MHS; PRSC; TCBI; AMIS; CNO; NFP; LKQX; BPA; CRI; JAS; SNX; BWLD; GHCI; OPEN; TSRA; AEL; CMP; HS; PJC; TPX; UTI; AHL; MGLN; PETS; AIZ; 9566B; HOS; NRG; SBNY; SBRA; HSP; AINV; NAL; GNW; GHL; ANGO; NILE; NUVA; RJET; KAR.1; CBG; CAB; IRC; LTM; MGI; CRM; MNTA; DPZ; ENS; WCG; BLKB; BUCY; PFWD; PSEC; VLTR; ACC; BMR; EXR; KRG; GOOGL; RNOW; SNTS; DYN; COGT; DWA; NANO; MOS; CLMS; TWGP; TXRH; IWA; NP; FOXA; MKTX; MPWR; NWE; SPOK; BBG; SMA; EDR; NE; OXPS; ANRZQ; FRP; TAP; PBH; SVR.1; WEX; NDAQ; NOVL; BNK.3; BOFI; TDS; PAY; DRH; ZUMZ; GOSHA; LHCG; LINC; MEND; VLCM; FSP; OSK; NSR; THS; FTK; MFB; MPW; CHAP; CNSL; DMND; EXPE; HITT; SWSI; CF; DRC; HPY; MWIV; RUTH; TRLG; PHS; AMP; TLEO; FNF; NCIT; AMSF; CBEY; IRBT; UA; ICE; CORE; CYNO; TRAK; BAS; LYV; VIAB; CMG; WNR; ACOR; CROX; LQDT; NTLS; ROSE; UAL; HS.2; KOP; ME; GPOR; NYX; GPRE; HOT; QSR; EQ; MA; AAWW; JCG; HOMB; PGTI; SNCR; RBK; WIN; WYN; GTLS; KALU; EVR; HBI; WU; CVLT; RVBD; EHTH; EXLS; DEI; LDOS; SXE.1; ROST;

EBS; KBR; CPLA; FSLR; IPHS; WINN; SE; ALGT; IPGP; EIG; HF; AVAV; ORIT; CENTA; BR; UFS; SFNC; TWC; FIRE; HAYN; SMCI; CNK; DAL; PCS; ACM; BGS; SLH; IBKR; KS; STTX; SCOR; DFS; COV; TEL; DHX; PMC; REXX; TYPE; MASI; VRTU; DOLNQ; EXH; PKT; TDC; CTCT; GXDX; CML; MV.2; ROIC; PATAQ; ZEP; LL; MSCI; APEI; AREX; ENSG; NFBK; RBCN; CATM; ENTR; MDAS; ORN; FOR; GFG; IPCM; AHC; CPN; DAN; IRDM; HI; NOG; PM; V; DPS; IPI; NX; UNH; LPS; SNI; JBT; GTATQ; RAX; HSNI; IILG; TKTM; TREE; DISCA; HCI; MYRG; CFL; CLW; VRTS; MJN; SWI; OPEN; MDSO; GOV; GSM; LOGM; CFN; AVGO; HAWKQ; SEM; VSI; ECHO; EDMC; VRSK; AOL; CLD; DG; LEA; STRI; FTNT; RUE; ESV; KRA; CHSP; TW; QNST; SPB; FNGN; SSNC; PRI; EXPR; FAF; RRTS; ONE; QEP; FN; CBOE