Performance Evaluation of Artificial Neural Networks in the Foreign Exchange Market

Harry
Hedman
870526-8574
harryh@kth.se

May 27, 2012

Master's Thesis (TMTHM), Spring 2012
Department of Mathematics, Royal Institute of Technology
Supervisors: Nils Rosendahl and Alexander Wojt
Examiner: Timo Koski

Abstract

This thesis examines the performance of artificial neural networks in the foreign exchange market. The thesis is restricted to comprise two types of network architectures: feedforward and probabilistic neural networks, respectively. The networks' capabilities are evaluated in a trading simulation, where predictions of exchange rate log-returns are backtested using historical data. All G10 currency pairs are considered, 45 in total. The results presented indicate that although several networks generate substantial returns, the average performance is rather modest. The foreign exchange market indeed appears efficient.

Acknowledgements

The author would like to thank Nils Rosendahl and Alexander Wojt at Nordea Markets in Stockholm for the opportunity given to write this thesis, and for the many rewarding discussions regarding the architecture of the artificial neural networks presented in the following. Naturally, the author assumes full responsibility for any errors present in the paper.

Contents

1 Introduction

The foreign exchange (FX) market is by far the largest financial market in the world and is of monumental importance. It operates 24 hours a day, seven days a week, and constitutes the foundation for world trade as we know it. FX trading is facilitated by a network of dealers as opposed to an exchange, and takes place simultaneously over the counter in different parts of the world. As of April 2010 average daily turnover totalled $4.0 trillion, up from $3.3 trillion three years before (von Kleist et al. (2010) [14]).

The FX market consists of a variety of different products. Measured by turnover, FX swaps constitute the largest part (44% in April 2010). However, most of the recent market growth is due to a dramatic increase in spot transactions; average daily turnover has increased 48% between April 2007 and April 2010. In 2010 it totalled $1.5 trillion and constituted 37% of global FX turnover (von Kleist et al. (2010) [14]).

For this reason and others the FX spot market is of particular interest. Since the breakdown of the Bretton Woods system in 1971 spot exchange rates have become increasingly floating, and a lot of theory has been devoted to the determinants of these exchange rates. In the long term, they are in principle determined by expectations of real interest rates. However, in the short term, exchange rates are often volatile and driven more or less by supply and demand.

The depth and liquidity of the spot market indicates that it should be hard to make money by predicting exchange rate changes. This claim is verified by Shmilovici et al. (2009) [12], who find that FX markets are indeed efficient enough to forbid profits being made. Nevertheless, a lot of time and resources are devoted to developing prediction schemes capable of beating the market, challenging the efficient-markets hypothesis while contributing to its validity.

For the purpose of time series prediction, the use of artificial neural networks (ANNs) is becoming increasingly popular. Hornik et al. (1989) [5] show that multilayer feedforward networks can estimate any continuous function arbitrarily well, making them suitable in situations where non-linearities are at hand. Dunis and Williams (2003) [2] find that ANN models outperform a benchmark of other statistical techniques in predicting EUR/USD exchange rate movements, and utilising such a network produce a profitable trading strategy. Zhang and Hu (1998) [16] reach a similar conclusion in a study of the GBP/USD exchange rate, and find that the number of input nodes has greater effect on network performance than the number of hidden nodes.

Probabilistic neural networks (PNNs) are viable candidates for classification purposes, albeit seldom applied to the FX market. Chen, Leung and Daouk (2003) [1] use PNNs to forecast the direction of movement of an index of Taiwanese stocks. They find that a PNN-based trading strategy outperforms trading strategies based on other forecasting methods, and it is quite possible that their results can be extended to also comprise the FX market.

The objective of this paper is to investigate the ability of ANNs to predict movements in spot exchange rates, and to implement and compare different ANN-based trading strategies. The study comprises all G10 currency pairs, 45 in total. However, in-depth analysis is limited to three currency pairs; the EUR/USD, AUD/USD and CHF/JPY. The EUR/USD is the world's largest exchange rate pair as measured by average daily turnover, with a market share of approximately 28% (von Kleist et al. (2010) [14])[1], and is interesting to analyse for this reason alone. The latter two currency pairs are studied in detail because the ANN-based trading models perform particularly well when applied to the AUD/USD and CHF/JPY exchange rate pairs, respectively.

Two different types of ANNs are considered: multilayer feedforward networks, suggested

[1]The market share presented is that of the entire FX market as of April 2010, i.e. including products other than spot exchange rates.

for trading purposes by e.g. Dunis and Williams (2003) [2] and Zhang and Hu (1998) [16], and PNNs, advocated by Chen, Leung and Daouk (2003) [1], albeit in a stock market setting. Contrary to much of the previous research on the topic, this paper analyses in depth the risks associated to trading with the suggested ANN-models, using risk metrics such as Value-at-Risk and Expected Shortfall in addition to standard performance measures such as Information ratio. Furthermore, the trading models developed in this paper are based on ANNs that are retrained on a regular basis, i.e. where a “sliding window” of inputs is employed by the networks to better detect changes in market sentiment.

The paper is organised as follows. First, the data set used in the study is analysed in Section 2. Section 3 gives a brief overview of the theory of neural network modelling, and is followed by Section 4, in which the neural network design procedure is explained. In Section 5, the trading model is introduced, and the main results are presented in Section 6. Finally, Section 7 concludes.

2 The Data Set

In all forecasting applications, the quality of the data set under study plays an important role. In this section the data source is introduced, and the rationale for including the variables presented as inputs to an ANN is explained. Furthermore, the respective EUR/USD, AUD/USD and CHF/JPY exchange rate time series are analysed, and a motivation is given for the transformation of the time series into log-returns.

2.1 Data Retrieval and Choice of Variables

The data used in this study originate from Thomson Reuters EcoWin Pro, a Gothenburg-based provider of financial market data and part of the Thomson Reuters Group. The data set was retrieved on 3 February 2012 and consists of daily data on exchange rates, interest rates, equity indices and oil prices from 3 January 2000 to 2 February 2012.

It should be stressed that since subsequent predictions of exchange rates are made for trading purposes, the time of the day at which the observations are made plays an important role – it would be a fatal mistake to develop a trading model using information that is not available to the market at the time the trades are made. The variables of which the data set is composed are described in Table 12 of Appendix A, where furthermore the sampling times of the respective data series are presented. Note that since the nine exchange rate pairs presented in the table are sampled at the same time of the day, the values of the 36 remaining permutations can be computed.

Although precise model specifications are postponed until Section 4, the variables in Table 12 deserve some explanation already at this point (except, of course, the exchange rates, which are dependent variables). To help decide which variables to use as inputs to a feedforward neural network, Kaastra and Boyd (1996) [8] argue that "economic theory can help in choosing variables which are likely important predictors". This ratifies the inclusion of interest rates in the model. The covered interest rate parity suggests that the forward exchange rate depends on the spread in risk free interest rates between the two currencies. The swap market is deep and liquid which indicates that data on swap rates is reliable, making swap rates a suitable choice of interest rates to be included in the prediction model as explanatory variables.

A country's economic performance affects the relative value of its currency. Data on equity indices is retrieved for this reason, as the performance of equities is closely tied to the state of the economy at large. One could of course argue that there are other, better measures of the real economy, such as industrial production or GDP growth. However, the sampling frequency of such economic indicators is often monthly at best, whereby financial data is preferred. Finally, oil is included because of its strong impact on the world economy, and because changes in the oil price can be assumed to affect different countries in different ways. E.g. a demand driven increase in the price of oil is likely to strengthen the NOK, the currency of oil producing Norway.

Instead of basing variable selection solely on economic rationale, one could argue that e.g. linear regression analysis of potential explanatory variables would be a suitable first step. In this paper no such techniques are employed. One of the strengths of ANNs is that "there are few a priori assumptions about the models under study" (Zhang, Patuwo and Hu (1998) [17]), whereas resorting to linear regression for variable selection, on the contrary, would require such a priori assumptions. Potentially important explanatory variables in the ANN might be left out, which would be very unfortunate.

2.2 Exchange Rate Characteristics

As explained in the introduction, this paper focuses mainly on the EUR/USD, AUD/USD and CHF/JPY exchange rate pairs, respectively. During the sampling period the EUR and the AUD have both appreciated against the USD, and the CHF has appreciated against the JPY. The evolution of the three exchange rates is depicted in Figure 1, and summary statistics are presented in Table 1. All exchange rate pairs are skewed and the respective time series are apparently path-dependent, which implies that some form of transformation is needed. The necessity of such a transformation is elaborated upon in e.g. Mehta (1995) [10], which makes the case that when working with ANNs in the FX market, "price inputs are not a desirable set".

	EUR/USD	AUD/USD	CHF/JPY
N	3152	3152	3152
Mean	1.2172	0.7509	85.4839
Median	1.2638	0.7540	86.8232
Max	1.5987	1.1023	106.6999
Min	0.8270	0.8270	59.1567
Std	0.1982	0.1585	9.8544
Skewness	-0.4066	0.1693	-0.6394
Kurtosis	2.1013	2.1076	3.0909

Table 1: Exchange rate summary statistics.

Log-returns R_t are thus computed, according to

$$R_t = \ln(S_t/S_{t-1}),$$

where S_t denotes the value of the exchange rate at time t. For each exchange rate pair, the computations are carried out using daily as well as weekly sampling frequencies. Since daily data are available, it is natural to compute daily log-returns. The rationale for extending the study to also comprise weekly sampling of exchange rates is that this may help to mitigate the effect of noise in the data set. Drawing on this, it is interesting to investigate potential differences in trading performance between daily and weekly ANN models. We know that transaction costs adversely affect the return of high frequency trading. If, in addition to this, patterns are more easily observable in a framework using weekly data, weekly sampling may lead to superior results. When weekly sampling is employed, observations of exchange rates made on Wednesdays are extracted from the daily data set.

Summary statistics for the log-returns of the EUR/USD, AUD/USD and CHF/JPY exchange rate pairs are presented in Table 2. The table depicts the respective cases of daily and weekly sampling as described above in separate columns. Given the construction of the weekly data sets, it comes as no surprise that the number of observations, N, is five times larger when the sampling is daily compared to when it is weekly. Furthermore, the means are approximately five times larger when the sampling is weekly, which is also expected. This is not, however, the case for median, maximum or minimum returns; for the AUD/USD and CHF/JPY exchange rates, it appears that the maximum weekly log-returns are smaller than the daily counterparts. The implication is that large daily returns likely occur quite randomly and, if anything, are followed by recoils.

The log-returns for the respective exchange rate pairs are depicted in Figures 2 and 3 for the daily and weekly sampling frequencies, respectively. The histograms show that the empirical

	Daily data			Weekly data		
	EUR/USD	AUD/USD	CHF/JPY	EUR/USD	AUD/USD	CHF/JPY
N	3152	3152	3152	630	630	630
Mean [%]	0.0079	0.0154	0.0079	0.0387	0.0770	0.0343
Median [%]	0.0201	0.0572	0.0278	0.0563	0.1074	0.1501
Max [%]	3.7188	8.2218	6.2065	10.1195	6.6858	5.5837
Min [%]	-2.7968	-7.6444	-8.1754	-4.6223	-17.2096	-8.888
Std [%]	0.6741	0.9158	0.7701	1.4827	1.9023	1.6402
Skewness	0.0145	-0.4782	-0.5248	0.3440	-1.4472	-0.6366
Kurtosis	4.2397	13.9203	11.0304	6.3039	14.2363	5.6956

Table 2: Summary statistics for exchange rate log-returns.

distributions have means close to zero around which they are relatively symmetric, and that the EUR/USD log-return series are the least fat-tailed. From the plots of the log-return time series we see that the volatility of the log-returns is non-constant and experiences a peak following the financial crisis in 2008. Finally, the action taken by the Swiss central bank in the summer of 2011 to depreciate the CHF is clearly visible in both the daily and the weekly plots of the CHF/JPY log-returns.

In light of the discussion in the latter two paragraphs, we conclude the data analysis by briefly investigating potential serial correlation in the daily log-return time series. To this end, the log-returns are regressed on lagged log-returns, where the model is arbitrarily restricted to include five lags. I.e. the model

$$R_t = \beta_0 + \beta_1 \cdot R_{t-1} + \cdots + \beta_5 \cdot R_{t-5} + u_t \tag{2.1}$$

is estimated using OLS, where u_t denotes the error term. Statistics from the respective regressions are presented in Table 3, and they provide useful insight regarding the behaviour of the respective log-return time series. More precisely, we can deduce from the p-values that joint insignificance of the explanatory variables can be rejected when the model is estimated using the AUD/USD time series, but not when using data on the EUR/USD or CHF/JPY exchange rate pairs. The implication is that there is no strong evidence of autocorrelation in the log-return time series of the latter two currency pairs, but that we should be very careful not to make similar assumptions for the series of AUD/USD log-returns. To conclude the discussion, the fact that the lagged returns are jointly significant in the AUD/USD setting suggests that this exchange rate pair is particularly well suited for the purpose of prediction.

	R^2	F-statistic	p-value
EUR/USD	0.0017	1.0827	0.36777
AUD/USD	0.0090	5.7284	0.00003
CHF/JPY	0.0023	1.4536	0.20180

Table 3: Regression statistics obtained from estimating the model in equation (2.1). The F-statistics and p-values are those of the full models.

The same log-transformations are performed on the respective equity index and oil price time series, for the same reasons. The time series of interest rates are also transformed, but differences are computed instead of log-returns. Summary statistics for these purely explanatory variables are omitted in this paper, as is graphical representation of their respective untransformed trajectories.

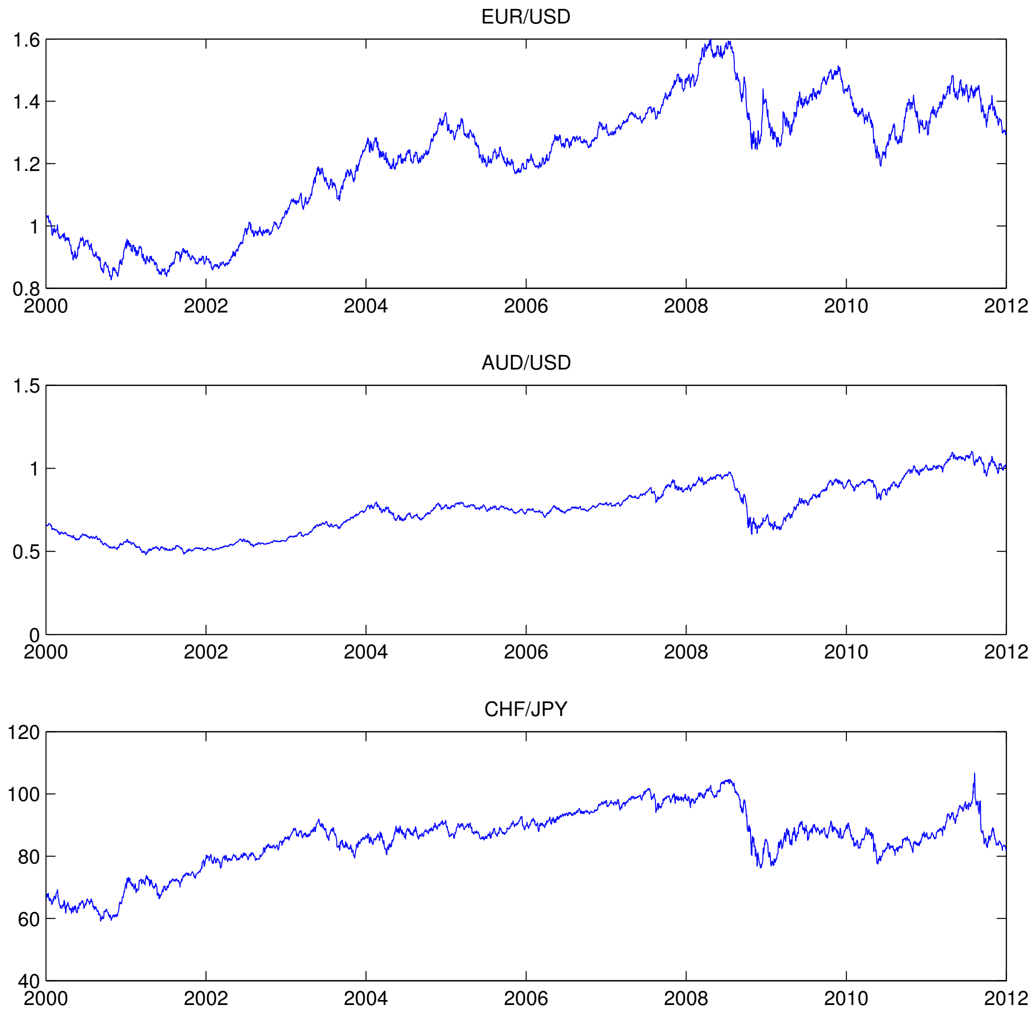

Figure 1: Daily exchange rates during the time period 4 January 2000 to 1 February 2012.

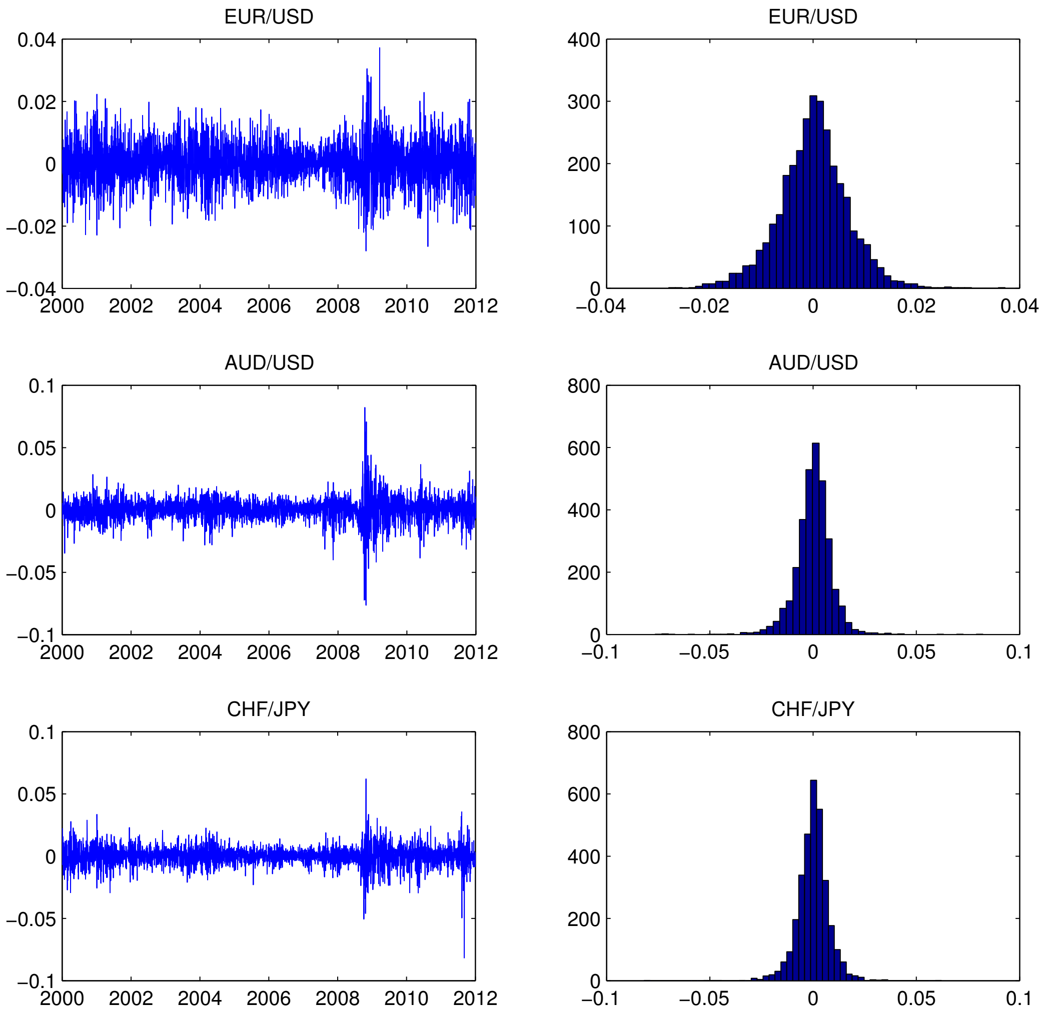

Figure 2: Daily exchange rate log-returns during the time period 4 January 2000 to 1 February 2012 (left column) and histograms of the same log-returns (right column).

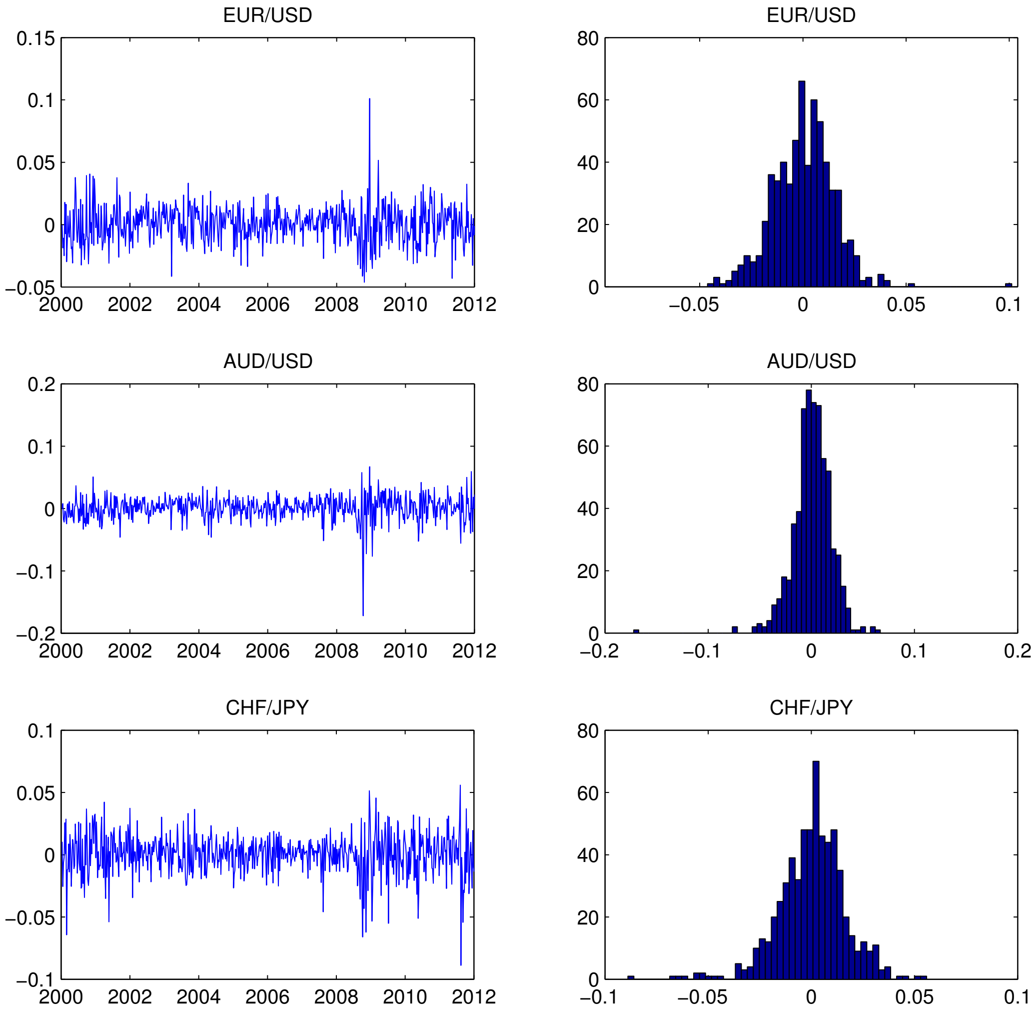

Figure 3: Weekly exchange rate log-returns during the time period 12 January 2000 to 1 February 2012 (left column) and histograms of the same log-returns (right column).

3 Theory of Neural Network Modelling

In this section two important types of ANNs are introduced, the feedforward neural network (FFNN) and the probabilistic neural network (PNN). In subsequent sections of this paper, the former is employed for function approximation and the latter for classification[2].

3.1 Feedforward Neural Networks

FFNNs have become increasingly popular for the purpose of function approximation. Hornik et al. (1989) [5] show that multilayer FFNNs can approximate any continuous function arbitrarily well given that the number of hidden nodes in the network is sufficient[3], without making any other assumptions with regards to functional form. For this reason FFNNs are applicable in situations where unknown nonlinear relationships are at hand, which presumably is the case in the FX market.

A FFNN consists of layers of interconnected nodes and transfer functions. Inputs are weighted and mapped into outputs, and the training of the network corresponds to updating weights and biases in order to minimise a performance function, usually the sum of squared errors. The rest of this section is devoted to describing the different components of the FFNN, the principles of its design and the training procedure. For a more rigorous survey of FFNNs, the reader is referred to Kaastra and Boyd (1996) [8] or Fine (1999) [3].

Figure 4 depicts schematically the structure of a FFNN with one hidden layer, similar to that implemented and used in later sections of this paper. The input layer comprises nodes x_1 to x_4, the inputs to the network. At the hidden nodes, denoted X_1 to X_3 in the figure, values are computed according to

$$a_j = f(b_j + \sum_{i=1}^{4} w_{j,i} \cdot x_i), \quad j = 1, 2, 3,$$

where f denotes a continuous transfer function (usually a sigmoid-type function such as $f(x) = 1/\left(1 + e^{-x}\right)$), b_j is the bias at node X_j and $w_{j,i}$ denotes the weight assigned to input x_i at hidden node X_j. These values are in turn fed forward to the output node Y, where the network output, $\tilde{y}$, is computed according to

$$\tilde{y} = g(b_Y + \sum_{j=1}^{3} w_j \cdot a_j).$$

Here w_j denotes the weight assigned to the hidden node output a_j, b_Y denotes the bias at the output node and g denotes a linear transfer function.

The network can be generalised to comprise an arbitrary number of input nodes, hidden nodes and output nodes. Moreover, the hidden nodes can be organised in layers, where nodes in contiguous layers are completely connected. As long as there is no feedback mechanism present, the network qualifies as a FFNN.

Consider a time series consisting of T different input-output combinations, i.e. a set $A = \{(\mathbf{x}_t, y_t) \,|\, t \in \{1, 2, \ldots, T\}\}$. For a given network architecture (such as the one presented in Figure 4), the modelling procedure commences by randomising the weights and biases of the network. The training of the network can then be carried out using the set A; inputs $\mathbf{x}_t$, $t \in \{1, 2, \ldots, T\}$, are fed through the network, which generates outputs $\tilde{y}_t$, $t \in \{1, 2, \ldots, T\}$. The weights and

[2]Appropriately designed, the feedforward neural network can be used also for classification.
[3]A sufficient condition is that the function is Borel measurable.

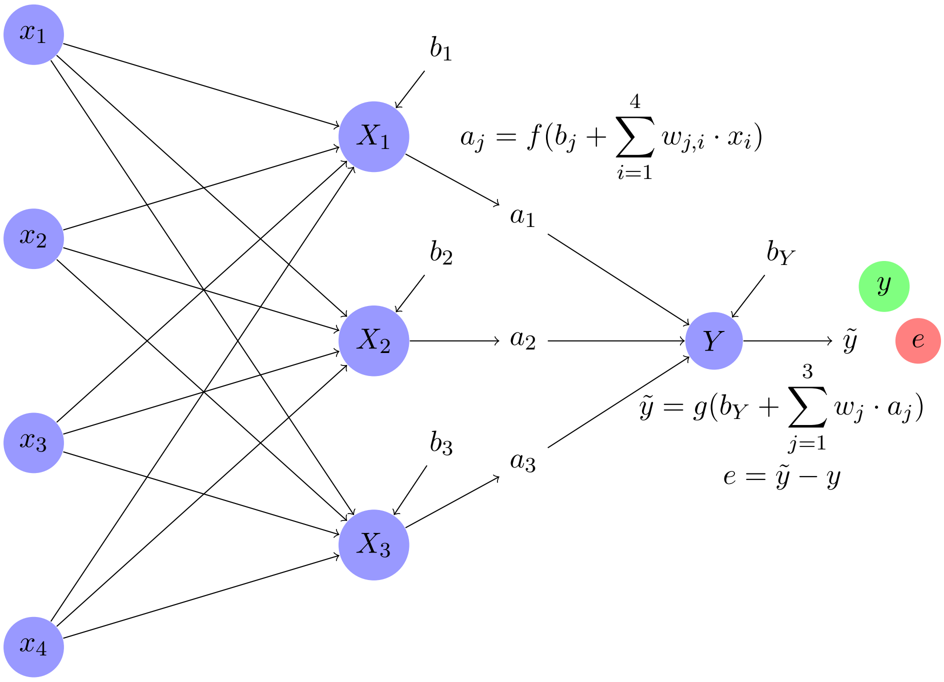

Figure 4: Feedforward neural network with one hidden layer.

biases are then adjusted as to reduce prediction errors. More specifically, the sum of squared errors

$$SE_A = \sum_{t=1}^{T} (\tilde{y}_t - y_t)^2$$

is computed and the weights and biases are updated using a backpropagation algorithm[4] to reduce SE_A. (A thorough survey of the backpropagation of errors is beyond the scope of this paper, and the reader is referred to Marsland (2009) [9] or Fine (1999) [3].) The network is then fed the same input-output set A, a new value of SE_A is computed, and the weights and biases are updated again. The iteration of this process composes the training of the network, and is carried out until a local minimum of SE_A is reached.

Crucial for the performance of the network is its ability to generalise, i.e. to make accurate predictions $\tilde{y}_t$ of y_t also when $(\mathbf{x}_t, y_t) \notin A$. For this reason, the sum of squared errors SE_B is computed on a disjoint set $B = \{(\mathbf{x}_t, y_t) \,|\, t \notin \{1, 2, \ldots, T\}\}$, and an additional stopping criteria is introduced. The training procedure stops if SE_B starts to increase, regardless of the development of SE_A. This prevents the network from overfitting, or learning the "noise" in the data set A. The sets A and B are often referred to as training and validation sets, respectively.

Another useful way to decrease the risk of overfitting is to keep the network as simple as possible with regards to the number of hidden layers and nodes. Only if sufficiently complex can the network learn to recognise all anomalies of the training data set at hand. A good overview of the problem of overfitting and potential remedies is provided in Kaastra and Boyd (1996) [8].

3.2 Probabilistic Neural Networks

For the purpose of developing a profitable trading model, it may be reasonable to lower the ambition from predicting the exact exchange rate in the future to simply deciding if it will be higher or lower than today. This approach gives rise to a classification problem. Chen, Leung

[4] A requirement for backpropagation to be feasible is that the transfer functions (f and g in Figure 4) are differentiable.

and Daouk (2003) [1] successfully use a probabilistic neural network (PNN) to forecast the direction of movement of an index of Taiwanese stocks, and the following section is devoted to the theory of PNNs.

PNNs can be used for solving classification problems. In this setting, a training set consisting of known input variables and corresponding outputs, is used to estimate a probability density function (PDF). Each output in the training set belongs to some class. When evaluated on data outside the training set, the PNN then classifies the the input variables using the estimated PDF. A class is assigned, corresponding to that with the highest probability of occurrence.

This Bayesian classification technique requires that a PDF is estimated for each class (Wasserman (1993) [15]). The PNN achieves this using Parzen windows; for a rigorous discussion of the topic the reader is referred to Parzen (1962) [11]. If each training vector consists of n components, a unit volume n-dimensional Gaussian is centered around each training vector in each class. The estimate of the n-dimensional Gaussian on the training set is then obtained by summing all such Gaussians in $\mathbb{R}^n$.

For the objective of this paper, two category classification is sufficient; the classes corresponding to up and down movements of the exchange rate. An extensive survey of two category classification using PNNs can be found in Wasserman (1999) [15], and is, albeit somewhat altered, reproduced below. The main result concerns the derivation of the PDF and how it is used for classification. Let $\mathbf{X} \in \mathbb{R}^n$ denote some n-component test vector to be classified. Two PDFs, one for each class $k \in \{1,2\}$, assign values $f_k(\mathbf{X})$ to the test vector according to the following equation:

$$f_k(\mathbf{X}) = \frac{1}{(2\pi)^{n/2}}\sigma^n \frac{1}{N_k}\sum_{i=1}^{N_k} \exp\left(-(\mathbf{X}-\mathbf{Y}_{ki})^T(\mathbf{X}-\mathbf{Y}_{ki})/(2\sigma^2)\right), \tag{3.1}$$

where

$$\begin{aligned}
f_k(\mathbf{X}) &= \text{the value of the PDF of class } k \text{ at point } \mathbf{X},\\
i &= \text{training vector number},\\
n &= \text{number of components in the training vector},\\
\sigma &= \text{standard deviation (to be chosen)},\\
N_k &= \text{number of training vectors in class } k,\\
\mathbf{X} &= \text{the vector to be classified, and}\\
\mathbf{Y}_{ki} &= i\text{th training vector from class } k.
\end{aligned}$$

The vector $\mathbf{X}$ is then classified according to the following decision rule, assuming that the "cost" of misclassification is the same regardless of to which class $\mathbf{X}$ does in fact belong:

$$\begin{aligned}
d(\mathbf{X}) &= \theta_1 \quad \text{if } h_1 f_1(\mathbf{X}) > h_2 f_2(\mathbf{X}),\\
d(\mathbf{X}) &= \theta_2 \quad \text{if } h_1 f_1(\mathbf{X}) < h_2 f_2(\mathbf{X}),
\end{aligned}$$

where

$$\begin{aligned} d(\mathbf{X}) &= \text{the decision on test vector } \mathbf{X}, \\ \theta_k &= \text{class } k, \\ h_k &= \text{the probability of occurrence of training vectors from class } k, \\ f_k(\mathbf{X}) &= \text{the probability density function for class } k \text{, evaluated at } \mathbf{X}. \end{aligned}$$

It remains for the user to decide on which standard deviation σ to assign to the Gaussians. The value of σ affects how the PNN classifies inputs, and should be chosen with care. However, it is shown in Specht (1967) [13] that as $\sigma \to 0$, nearest neighbour classification is obtained, adding to the robustness of the PNN. Nonetheless, the choice of σ is the main concern when implementing a PNN to solve a classification problem.

4 Network Architecture

In this section the architecture of the different ANNs used in subsequent trading simulations is described. The networks are constructed in MATLAB, where the implementation is facilitated by using the MATLAB Neural Network Toolbox. An overview of the many capabilities of the toolbox is given in [6]. The scope of this thesis is limited to the analysis of FFNNs and PNNs, respectively. E.g. recurrent neural networks and radial basis networks are likely powerful tools for the purpose of time series prediction, however they are not discussed in this paper.

4.1 Network Inputs

It is chosen, for the sake of simplicity, to use the same types of input variables for all networks. Drawing on the discussion in Section 2.1, the set of input variables is chosen to comprise equity index log-returns, changes in swap rates and oil price log-returns. In addition, lags of the respective variables are included in the models, as are lags of the dependent variable, i.e. the exchange rate log-returns. The variables are lagged up to the arbitrary number of five time periods.

In theory it would be possible, and perhaps advantageous, to include all equity index log-returns and swap rate changes (from all ten G10 countries) as inputs regardless of which exchange rate pair is modelled. However, in order to curb the computational time required the model complexity is reduced. Only variables stemming from the two most relevant countries or regions[5] are included in each model, and in addition these values are differenced. The precise definitions of the network inputs are given below.

To predict the log-return from time $t-1$ to time t of exchange rate pair X/Y,[6] denoted $R_{t,X/Y}$, the inputs presented to the respective networks are

$$
\begin{aligned}
&R_{k,X/Y}, && k \in \{t-5,\dots,t-1\}, && \text{Lagged exchange rate log-returns,}\\
&R^e_{k,X} - R^e_{k,Y}, && k \in \{t-5,\dots,t-1\}, && \text{Lagged differences in equity index log-returns,}\\
&\Delta_{k,X} - \Delta_{k,Y}, && k \in \{t-5,\dots,t-1\}, && \text{Lagged differences in swap rate changes,}\\
&R^o_k, && k \in \{t-5,\dots,t-1\}, && \text{Lagged oil price log-returns.}
\end{aligned}
$$

Here $R^e_{k,C}$ denotes the log-return from time $k-1$ to time k of an equity index in the country corresponding to currency C, $\Delta_{k,C}$ is the change from time $k-1$ to time k in the swap rate offered in currency C, and R^o_k is the log-return of the price of oil from time $k-1$ to time k. The time interval is one day and one week in the cases of daily and weekly data sampling, respectively. A more detailed description of the variables can be found in Table 12 of Appendix A.

4.2 Feedforward Neural Networks

FFNNs with 20 input nodes (the number of inputs to the network), and one hidden layer comprising ten hidden nodes, are implemented in MATLAB using the Neural Network Toolbox. The number of hidden nodes affects the network's ability to approximate potential nonlinear relationships between the input and output variables, and is an important parameter in the network design process. The decision in this paper, to use ten such nodes, is based on the findings of Zhang and Hu (1998) [16], and serves to keep computational times on a reasonable level. Furthermore, a relatively simple network is less prone to overfitting the data.

[5] I.e. the countries from which the respective currencies of the exchange rate pair originate.
[6] The PNN models instead classify the log-returns as being either positive or negative.

A sliding window methodology is employed, which implies that the network is reconfigured and retrained on a daily basis in the case of daily data sampling and on a weekly basis in the case of weekly data sampling. The window lengths for the daily and weekly models are 250 data points and 125 data points, respectively. I.e. 250 or 125 input-output combinations are used to train the network, for the task of predicting the subsequent (251st or 126th) exchange rate log-return. The window lengths are chosen as to capture the current market sentiment while still providing sufficient amounts of data to enable the training of the respective networks.

All networks use a log-sigmoid transfer function $f(x) = 1/(1+e^{-x})$ in the hidden layer and a linear transfer function $g(x) = x$ in the output layer. Prior to the start of the training, biases and weights are randomly selected. The networks are trained using the Levenberg-Marquardt algorithm, and the performance function to be minimised is the sum of squared errors.[7] The data set is randomly divided into training and validation sets, where the respective sets contain 70 and 30 per cent of the data. The training process stops when the magnitude of the gradient of the performance function is smaller than 10^{-5}, or if the performance function evaluated on the validation set fails to decrease for six consecutive iterations.

4.3 Probabilistic Neural Networks

Having already settled on the network inputs described in Section 4.1, the process of designing the PNNs is nearly finished. It remains to decide which standard deviation σ to assign to the PDFs in equation 3.1, as well as to choose appropriate window lengths. As is the case with the FFNNs, the PNNs are implemented in MATLAB using the Neural Network Toolbox.

The sliding windows used to implement the PNNs in this thesis comprise 200 and 125 input-output combinations in the respective cases of daily and weekly data sampling. The rationale for choosing these particular values is, as in the FFNN case, that information on relatively recent market behaviour is supposedly preferable to train a network, since the market sentiment can change rather quickly. At the same time, too short windows are not advantageous for the purpose of learning.

The process of selecting an appropriate σ is more rigorous. Let N denote the number of days in the training window. A validation window of length $k < N$ is selected, on which the performance of several PNNs, constructed using different values of σ, is evaluated with respect to the proportion of correct classifications. The process can be described as follows. Days 1 to $N-k$ of the training window are used to predict the direction of change of the exchange rate at day $N-k+1$, days 2 to $N-k+1$ are used to predict the direction of change at day $N-k+2$, etc. The procedure continues until days k to $N-1$ have been used for predicting the directional change of the exchange rate at the Nth and final day of the training window. This gives a total of k predictions for each value of σ, which are used to compute the proportion of correct directional predictions for the respective values of σ.

The value of σ that generates the best performance is then selected for the purpose of predicting the directional change of the exchange rate at time point $N+1$, i.e. outside the training sample. Using this value of σ, the PDFs in equation 3.1 are estimated using N input-output combinations, i.e. data corresponding to the entire training window is used for this final estimation. The entire procedure is iterated as the window slides forward to comprise the time interval 2 to $N+1$, and so on. In the implementation, N and k are assigned the values 200 and 50, and 125 and 25, in the respective cases of daily and weekly data sampling.

[7] A survey of the Levenberg-Marquardt algorithm is well beyond the scope of this paper. The algorithm is described in e.g. Hagan and Menhaj (1994) [4].

4 NETWORK ARCHITECTURE

A heuristic approach is taken to determine which values of σ should be considered in the above procedure. Let $\mathbf{Y}_t \in \mathbb{R}^{20}$ denote a vector consisting of network inputs at some time point t (the inputs are the ones suggested in Section 4.1). It is found upon investigation of the data set that a typical distance between two input vectors $\mathbf{Y}_k$ and $\mathbf{Y}_{k+1}$ is $||\mathbf{Y}_k - \mathbf{Y}_{k+1}||_2 \approx 0.1$.[8] Based on this observation, the set of possible values of σ is taken to be $\{0.002, 0.01, 0.05, 0.1, 0.2, 0.3, 0.5\}$.

[8]Somewhat surprisingly, the distance is roughly the same in the respective cases of daily and weekly data sampling.

5 The Trading Model

In this section the trading model used for backtesting the respective ANNs is described. As such, the model is not particularly sophisticated, and it is included in the paper mainly for the sake of clarity. Only one exchange rate pair at a time is considered by the trading model, and a position must at all times be taken in one of the two currencies.

The following trading rule is implemented. In the case where FFNNs are used to predict exchange rate log-returns, buy and sell signals are generated when the predicted log-return is positive and negative, respectively. If the sign of the predicted log-return is the same as that of the prediction of the previous day (or week), no action is taken; the position remains unchanged. The trading rule is completely analogous when PNNs are instead used. Buy and sell signals are generated when the log-return for the next day (or week) is classified by the network as being positive or negative, respectively. The position remains unchanged until a different classification is encountered.

It is important to note that over the course of the backtesting period, holdings in the respective currencies are not invested at the corresponding risk free rates. Although common in the previous research, this oversimplification indeed raises questions regarding the reliability of the trading results. An important driver of exchange rates is the interest rate spread between the respective currencies, which ideally should be accounted for. When the trading intensity (number of trades/number of trading days) is high, the problem is mitigated somewhat, as this is an indication that both currencies are held for similar proportions of the trading period. Moreover, if this is not the case, the results obtained using weekly data sampling are particularly unreliable.

As it turns out, the FFNN models typically conduct trades roughly every other trading day, whereas the PNN models trade on about one third to one fourth of the possible occasions. In light of this observation, it may be the case that the weekly PNN models are evaluated with a substantial disadvantage.

Transaction costs are accounted for in the following manner. The EUR/USD is the world's most frequently traded exchange rate pair. For the purpose of computing the transaction costs that arise from trading the EUR/USD, typical bid and ask prices are assumed to be 1.2999 \$/€ and 1.3001 \$/€, respectively. In relative terms, the bid-ask spread is $0.0002/1.3 \approx 0.00015$. When backtesting the EUR/USD models, transaction costs are accounted for by subtracting half this value (≈ 0.000075) from the corresponding log-return each time a trade is made. This approximation is likely quite accurate, and if anything, on the conservative side.

Other exchange rate pairs are less liquid, which implies that the relative bid-ask spreads associated to trading these currencies are higher. For the sake of simplicity, transaction costs are doubled for all other exchange rate pairs. I.e. the number 0.00015 is subtracted from the corresponding log-return each time a trade is made.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Beginning of trading period	25 Dec 2000	3 Jul 2002	16 Oct 2000	3 Jul 2002
End of trading period	1 Feb 2012	1 Feb 2012	1 Feb 2012	1 Feb 2012
Number of trading days	2897	500	2947	500
Window length	250	125	200	125

Table 4: Window lengths and trading periods.

The time periods during which the trading performances of the respective ANNs are evaluated are not the same for all networks. The time periods depend on the data sampling frequency (daily or weekly) as well as the window lengths of the respective networks, i.e. the amount of historical data that is used for training the networks. However, for each network architecture, the backtesting period is the same for all 45 exchange rate pairs. The trading periods for the respective models are presented in Table 4.

6 Results

In this section the results obtained from backtesting the respective trading models explained in Section 5 are presented. Special attention is given to the EUR/USD, AUD/USD and CHF/JPY exchange rate pairs – only the results corresponding to these currency pairs are scrutinised in detail in this paper. An overview of the results stemming from other exchange rate pairs is given in Tables 13 to 54 of Appendix A.

The section is organised as follows. First, the performance measures used in the subsequent model evaluations are defined. This is followed by an overview of the performance of the respective trading models based on results from all 45 G10 exchange rate pairs, before a study of the respective EUR/USD, AUD/USD and CHF/JPY results is conducted. Next, two top performing models are analysed in greater detail, where metrics similar to empirical Value-at-Risk and empirical Expected Shortfall are computed for the respective models. The section concludes by summarising the findings.

6.1 Performance Measures

When evaluating the performance of any trading model, two metrics are of particular importance: return and volatility. A potential investor will judge different investment opportunities based primarily on these two measures, in combination with his or her personal appetite for risk. However, for the sake of model evaluation, other measures are also of interest. The performance measures included in this paper are similar to those used by Dunis and Williams (2003) [2], and are presented in the following.

"Cumulative return" is simply the return achieved by the trading model during the entire backttesting period, net of approximated transaction costs as described in Section 5. Since the backtesting period is not the same for all models under study, "Annualised return" is computed for the sake of comparability. "Annualised volatility" denotes the sample standard deviation of log-returns, scaled to an annual figure by multiplying it by $\sqrt{252}$ and $\sqrt{52}$ in the respective cases of daily and weekly data sampling. The ratio of "Annualised return" and "Annualised volatility" is denoted "Information ratio", and is perhaps the most important measure to take into consideration when evaluating an FX trading model. FX trading implies frequently taking new positions in different currencies, and hence Information ratio is preferred over the more famous Sharpe ratio (as the choice of benchmark risk free rate is not entirely clear).

"Annualised transaction costs" is the adverse effect of transaction costs on return, cumulated over the entire backtesting period and scaled to an annual figure. "Number of trades" is the number of trades undertaken during the backtesting period, and should be compared to the total number of possible trading days for the respective models, depicted in Table 4. "Winning trades" and "Losing trades" are the respective ratios of profitable and non-profitable trades conducted, after accounting for transaction costs. "Average gain" and "Average loss" each dentote the sample mean of log-returns achieved during the backtesting period, conditional on the log-returns being greater than zero or less than or equal to zero, respectively. It follows from this that the figures are daily averages in the case of daily data sampling, and weekly averages in the case of weekly data sampling. The ratio "Average gain/Average loss" is included to facilitate the comparison of the respective measures. Finally, "Winning up periods" and "Winning down periods" denote the proportion of time periods during which positive returns are generated, when the exchange rate appreciates and depriciates, respectively.

6.2 Backtesting of the Trading Models

6.2.1 Comparison of Feed-Forward and Probabilistic Neural Networks

To begin the analysis, the overall trading performance of FFNNs is compared with the performance of PNNs. The results are striking. By looking at the average annualised return taken over all 45 G10 exchange rate pairs, it is evident that FFNNs generate better results than PNNs when the respective models are backtested on the data set used in this thesis. The qualitative difference in performance is rather stunning; after accounting for transaction costs, FFNNs on average generate positive (albeit small) returns whereas PNNs on average generate negative returns. Table 5 depicts arithmetic averages of annualised return, annualised volatility and Information ratio for the respective models, computed using figures from Tables 6, 7 and 8, as well as Tables 13 to 54 of Appendix A.

A word of caution is warranted at this point. Based on the results in Table 5, it is tempting to arrive at the conclusion that PNNs are ill-suited for the purpose of predicting FX movements. One should, however, be careful not to make such generalisations. The results obtained are dependent on the network design as proposed by the author, which can likely be improved. In addition, it appears that PNNs trade less frequently than FFNNs. Drawing on the arguments in Section 5, the fact that the proceeds of the trading are not invested might in part explain the adverse average performance of the PNN trading models.

Table 5 indicates that daily trading on average is more successful than weekly trading when using FFNNs, whereas the opposite holds for PNNs. Looking at the table, it is furthermore interesting to see that the annualised volatility for a given trading frequency is the same regardless of which type of network is used. Finally, when positive, the average returns are very small in magnitude; an indication that FX markets in general are highly efficient.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Annualised return	0.0068	0.0037	-0.0101	-0.0020
Annualised volatility	0.1181	0.1157	0.1181	0.1157
Information ratio	0.0684	0.0252	-0.0863	-0.0220

Table 5: Arithmetic averages of key performance measures, computed using the results obtained from all 45 G10 exchange rate pairs.

6.2.2 The EUR/USD, AUD/USD and CHF/JPY Exchange Rate Pairs

The main scope of this paper is on the results related to trading the respective EUR/USD, AUD/USD and CHF/JPY exchange rate pairs. The EUR/USD is scrutinised in detail because it is the world's most traded exchange rate pair. An implication of this high liquidity is that the EUR/USD market is likely very efficient, making profits hard to achieve. Indeed, this hypothesis cannot be rejected based on the results presented in the following. None of the networks under consideration produce significant trading profits when trading the EUR/USD. For the latter two exchange rate pairs, the AUD/USD and the CHF/JPY, the story is different. They are included in the study because appealing trading results are generated when backtesting the corresponding models.

Tables 6, 7 and 8 depict the results obtained from backtesting the different trading models. As mentioned, the EUR/USD trading returns are not particularly good. Despite low transaction costs, only one model generates a positive return: a very modest annualised return of no more than 0.36%. Figure 5 depicts the accumulation of total return over time, and it can be seen that no single market event can explain the adverse returns of the EUR/USD models. Instead a stable but poor overall performance is observed. Contrary to Dunis and Williamson (2003) [2], this paper does not advocate trading the EUR/USD using ANNs.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.0421	-0.3067	-0.2525	-0.1327
Annualised return	0.0036	-0.0374	-0.0246	-0.0147
Annualised volatility	0.1059	0.1046	0.1064	0.1047
Information ratio	0.0339	-0.3575	-0.2311	-0.1403
Annualised transaction costs	0.0092	0.0016	0.0061	0.0007
Number of trades	1414	210	959	85
Winning trades	51.91%	49.05%	50.89%	48.24%
Losing trades	48.09%	50.95%	49.11%	51.76%
Average gain	0.0050	0.0110	0.0050	0.0103
Average loss	0.0050	0.0109	0.0051	0.0117
Average gain/Average loss	1.0064	1.0104	0.9701	0.8834
Winning up periods	52.39%	50.18%	49.31%	58.91%
Winning down periods	47.45%	41.78%	50.28%	43.11%

Table 6: EUR/USD trading results.

The results obtained from backtesting the respective AUD/USD models are presented in Table 7. It can be seen that all four trading models, including the PNN models, generate positive annualised returns. The fact that the models employing a daily sampling frequency work well comes as no surprise in light of the data analysis performed in Section 2.2, where it is concluded that autocorrelation is present in the time series of daily exchange rate log-returns. No corresponding analysis was carried out on the weekly sample, but it turns out that the weekly models outperform their daily counterparts regardless of if the network design is feed-forward or probabilistic.

The weekly FFNN model performs best, generating an annualised return of 8.18% and an Information ratio of 0.5754 after accounting for transaction costs. The relatively high return can be explained by a sufficient proportion of winning trades (52.46%) in combiantion with the fact that average gains are 11.14% larger than average losses. This is achieved despite that the weekly FFNN is rather bad at predicting correctly depreciations of the AUD/USD exchange rate; the fraction of winning down periods is less than 50%.

Figure 6 depicts "Correct directional change" over time for the respective trading models. The metric is defined as the proportion of correct directional predictions achieved by the networks, and is computed annually.[9] The steep increase in cumulative return for the weekly AUD/USD PNN model observed during 2009, depicted in Figure 5, is in part explained by a high proportion of correct directional change predictions during the same time period (> 60%). However, by looking at Figures 1 and 2 in Section 2.2, an indication is given that the large magnitude of the returns is rather driven by high volatility in the AUD/USD exchange rate during the same time period.

The results obtained when trading CHF/JPY are depicted in Table 8. They are, contrary to when trading AUD/USD, highly dependent on which of the four trading models is being considered. Whereas the weekly FFNN model generates the highest annualised return and

[9] "Correct directional change" differs from "Winning trades" since the latter measure is conditional on a trade taking place.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.8572	1.1297	0.4584	0.6963
Annualised return	0.0553	0.0818	0.0328	0.0565
Annualised volatility	0.1473	0.1421	0.1471	0.1424
Information ratio	0.3756	0.5754	0.2229	0.3968
Annualised transaction costs	0.0184	0.0038	0.0099	0.0022
Number of trades	1424	244	778	144
Winning trades	49.23%	52.46%	48.84%	51.39%
Losing trades	50.77%	47.54%	51.16%	48.61%
Average gain	0.0065	0.0146	0.0064	0.0142
Average loss	0.0062	0.0132	0.0063	0.0137
Average gain/Average loss	1.0645	1.1114	1.0181	1.0406
Winning up periods	53.51%	56.20%	55.77%	55.11%
Winning down periods	46.21%	48.67%	44.56%	50.00%

Table 7: AUD/USD trading results.

Information ratio of all models studied in this paper (10.54% and 0.9235, respectively), other models generate more modest, or even negative, values. Moreover, all four models behave similarly until the end of 2008, which is likely explained by a rather stable CHF/JPY exchange rate. Following the bankruptcy of Lehman Brothers in September 2008, volatility increased, and it appears that the weekly FFNN model best anticipated the market movements during the subsequent financial crisis. From Figure 6, it is clear that the weekly FFNN model is outstanding at predicting correctly the directional changes during 2010 (nearly 70% of the predictions are correct), which coincides with a rather steep increase in cumulative return depicted in Figure 5. However, the black-box nature of ANNs makes it hard infer potential underlying economic explanations of the success of the trading model.

With regards to the success of the weekly FFNN model, the main points to take away from Table 8 are the following. It is not the proportion of winning trades (48.37%) that drives the accumulation of returns, but rather the ratio of average gains to average losses (1.2014). Quite naturally, it is the big deviations that are most important to predict.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.2152	1.6198	0.5833	-0.2132
Annualised return	-0.0209	0.1054	0.0401	-0.0246
Annualised volatility	0.1202	0.1141	0.1202	0.1149
Information ratio	-0.1736	0.9235	0.3335	-0.2143
Annualised transaction costs	0.0184	0.0033	0.0110	0.0023
Number of trades	1427	215	863	150
Winning trades	47.93%	48.37%	50.87%	50.67%
Losing trades	52.07%	51.63%	49.13%	49.33%
Average gain	0.0053	0.0126	0.0053	0.0109
Average loss	0.0054	0.0104	0.0054	0.0123
Average gain/Average loss	0.9846	1.2014	0.9751	0.8844
Winning up periods	51.16%	58.18%	60.49%	62.55%
Winning down periods	47.90%	48.44%	42.84%	36.89%

Table 8: CHF/JPY trading results.

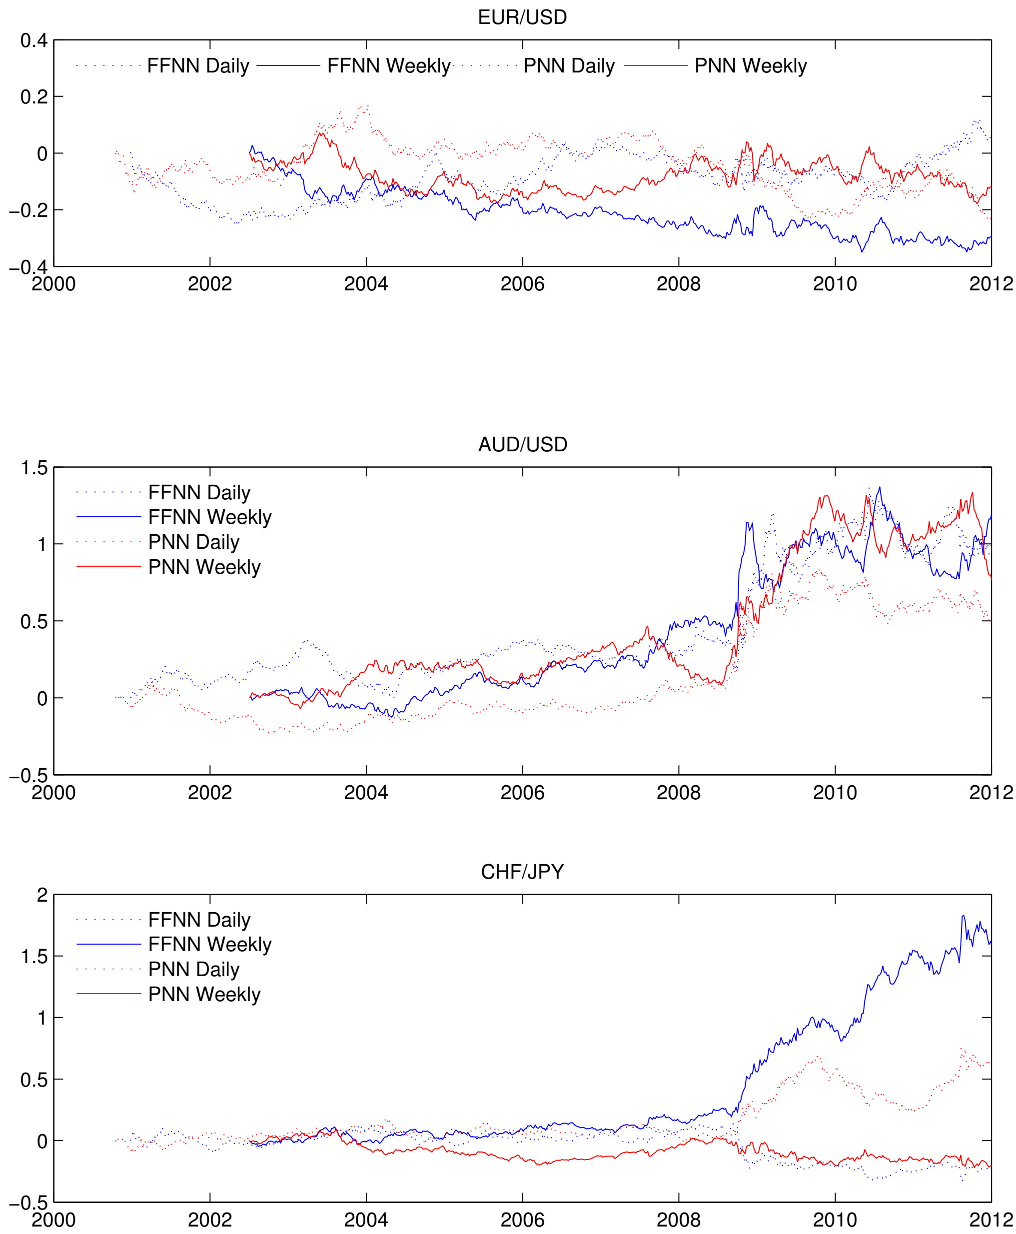

Figure 5: Cumulative returns.

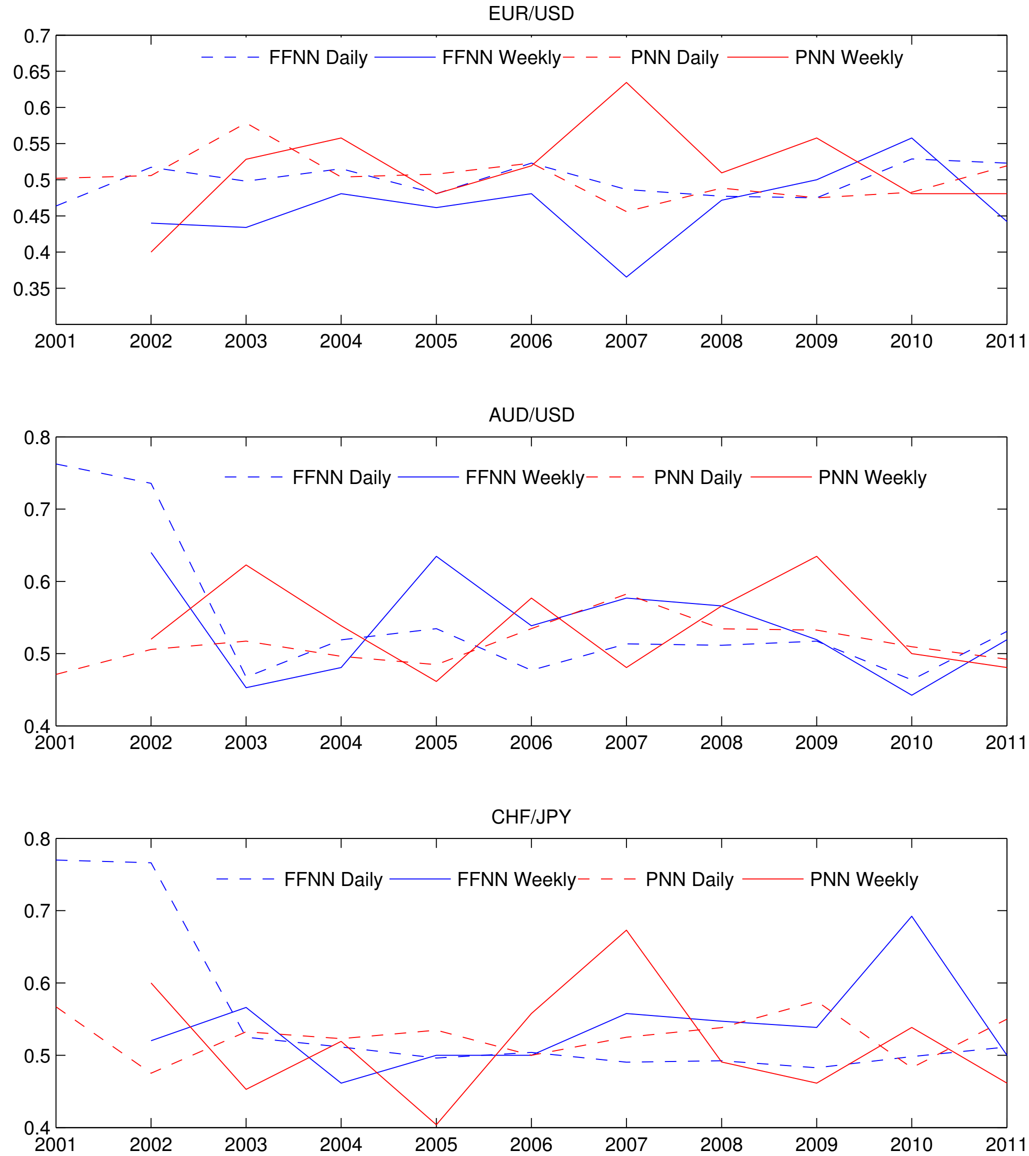

Figure 6: Correct directional change over time.

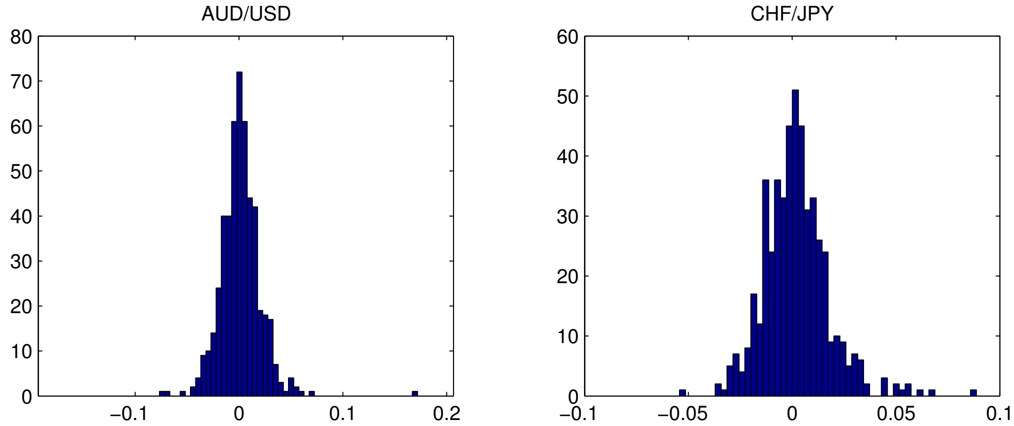

Figure 7: Histogram over log-returns generated by the weekly FFNN models.

6.2.3 Risk Analysis of Two Top Performing Models

The following analysis is restricted to the two top performing trading models encountered in this paper, namely the respective AUD/USD and CHF/JPY weekly FFNN models. Table 9 presents summary statistics for the (weekly) log-returns generated by the respective models, and Figure 7 depicts corresponding log-return histograms. It is instructive to compare the summary statistcs with those presented for the exchange rate log-returns in Table 2 of Section 2.2. Although the statsistics are not perfectly comparable since the sampling periods are somewhat different, the success of the respective trading models can be inferred from the larger mean, median, minimum and maximum values of the log-returns. Moreover, the log-returns are skewed to the right when generated by the trading models, as opposed to the exchange rate log-returns which exhibit negative skewness. This change in the character of the returns is clearly visible when comparing the respective histograms. It is clear that the models are performing well.

However, the risk averse investor has reason to be concerned. The standard deviation of the log-returns are of roughly the same size in the respective tables, which implies that the underlying volatility of the FX market is still inherent in the trading returns.[10] The following analysis aims to quantify the risks faced when trading with the respective models by computing metrics similar to Value-at-Risk and Expected Shortfall.

	AUD/USD	CHF/JPY
N	500	500
Mean [%]	0.1512	0.1926
Median [%]	0.0798	0.1312
Max [%]	17.1946	8.8730
Min [%]	-7.6661	-5.4247
Std [%]	1.9711	1.5819
Skewness	1.2859	0.8236
Kurtosis	14.5000	6.0726

Table 9: Summary statistics for log-returns generated by the weekly FFNN models.

The histograms in Figure 7 give a hint that the log-returns generated by the respective weekly FFNN models are non-normal. To further investigate the distributional properties of

[10]The turbulence following the financial crisis in combination with the smaller data set used for trading makes the standard deviation of the log-returns generated by the trading models unfairly magnified.

the respective log-return samples, quantile-quantile-plots (qq-plots) are presented in Figure 8. Drawing on the notation used in [7], a qq-plot is the plot of an ordered sample $z_1 \geq, \ldots, \geq z_n$ of observations belonging to some common distribution, against the quantiles of some reference distribution F. I.e., a qq-plot depicts the points

$$\left\{ \left(\mathrm{F}^{-1}\left(\frac{n-k+1}{n+1}\right), z_k \right) : k = 1, \ldots, n \right\}.$$

It follows that a qq-plot is linear if the observed data is generated by the reference distribution. With this in mind, it is clear from the upper two plots in Figure 8 that neither the AUD/USD nor the CHF/JPY trading models generate normally distributed log-returns; the respective empirical samples are too fat-tailed.

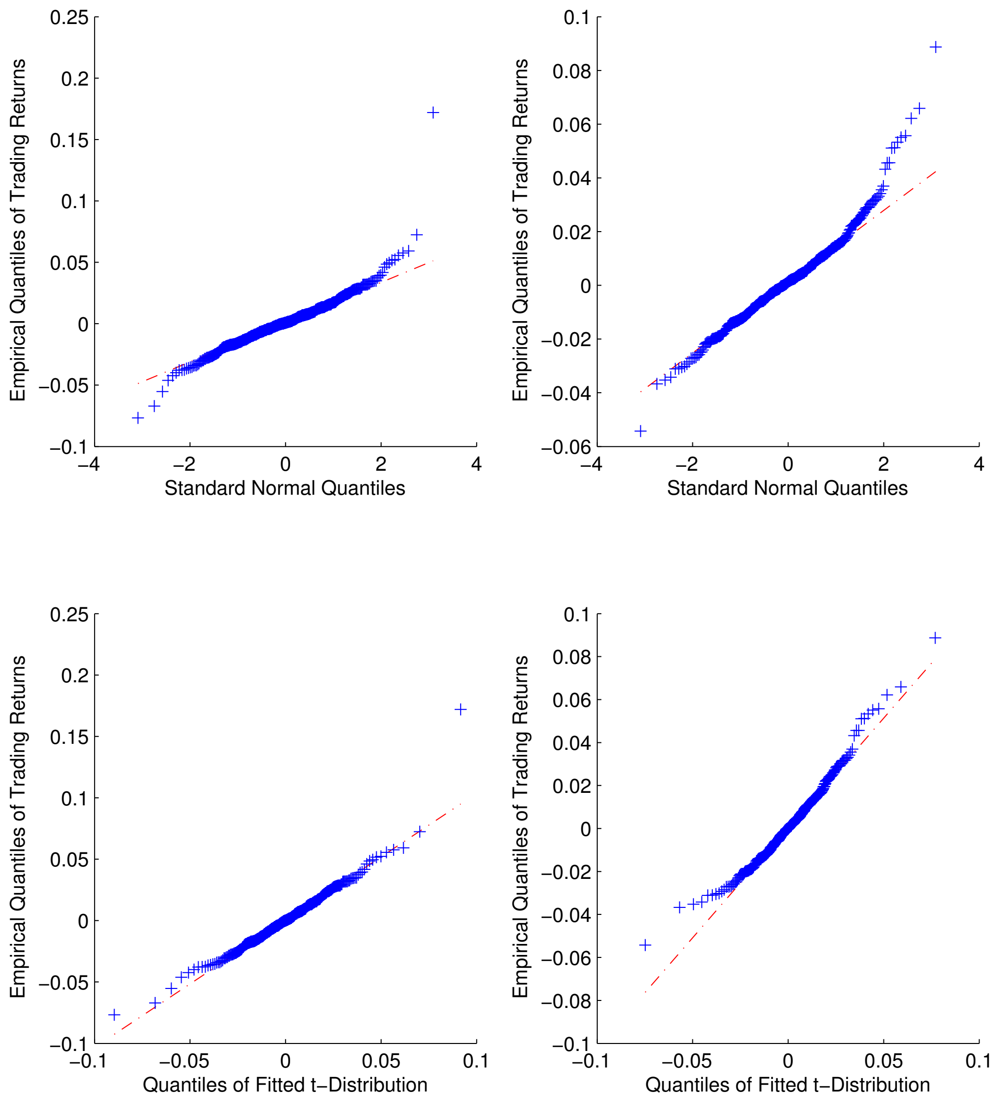

Figure 8: Qq-plots of the log-returns generated by the weekly FFNN trading models. The plots to the left depict the AUD/USD exchange rate pair, whereas the plots to the right depict the CHF/JPY exchange rate pair.

The Student's t-distribution has heavier tails than the normal distribution. In light of the findings so far, t-distributions are fitted to the empirical samples of log-returns using maximum likelihood estimation. The obtained parameter estimates are presented in Table 10, and the respective estimated t-distributions are used as reference distributions in the bottom two plots of Figure 8.

	AUD/USD	CHF/JPY
$\hat{\mu}$	0.0010	0.0011
$\hat{\sigma}$	0.0144	0.0120
$\hat{\nu}$	4.6152	4.5623

Table 10: Parameters of t-distributions, fitted using maximum likelihood estimation.

It is clear from the bottom right qq-plot that the t-distribution is not a suitable choice to approximate the CHF/JPY log-returns. The empirical trading returns exhibit lighter left tails and heavier right tails than the estimated t-distribution. The AUD/USD log-returns are, for the most part, reasonably well approximated by the estimated t-distribution, which is depicted in the bottom left qq-plot. However, the fitted distribution completely fails to capture one of the outlying observations, which can have serious implications from a risk management perspective. One could of course argue that there is no real risk associated to underestimating the upside of a trading strategy, but given the nature of the trading strategy, the problem cannot be overlooked. One wrong prediction can transform a significant gain into a significant loss, and given the rather modest ability of the FFNN to predict correctly directional changes of the CHF/JPY exchange rate, I choose to discard the suggested t-distribution for the purpose of estimating the underlying distribution of weekly log-returns.

Although there might exist other, more well-suited distributions that can be used to approximate the empirical distribution of the respective log-return samples, based on the results so far the use of parametric methods for the purpose of computing risk metrics is refrained from in this paper. Instead an empirical approach is chosen.

Empirical approaches are not, however, completely unproblematic. On the contrary, inherent in the following analysis is the assumption that the log-returns contained in the respective empirical samples are independent (within the samples). In Section 2.2 it was concluded that there was no evidence of serial correlation up to order five in the daily sample of exchange rate log-returns for the CHF/JPY, which is a source of some comfort. However, the hypothesis of joint insignificance of lagged returns could be rejected in the AUD/USD case, which unfortunately increases the probability that the weekly AUD/USD trading returns are not independent. Nevertheless, the empirical approach is preferred over the parametric.

In the following Value-at-Risk and Expected Shortfall computations, an addtional simplifying assumption is made: losses are not discounted to present value. The reason for this is that there is some ambiguity regarding which risk free rate to apply to a portfolio which repeatedly changes holdings between two currencies. However, the assumption is conservative in the sense that losses are over- rather than underestimated, and the assumption can hence be deemed acceptable. Moreover, it makes sense not to discount the losses since no money is invested between trades.

Again drawing on the notation in [7], the empirical estimate of Value-at-Risk at level p is

$$\widehat{\mathrm{VaR}}_p = \mathrm{L}_{\lfloor np \rfloor + 1},$$

where $\mathrm{L}_1 \geq \cdots \geq \mathrm{L}_n$ is the ordered sample of losses. In the present context, the losses are computed from the log-returns, i.e.

$$\mathrm{L}_k = -\left(e^{R_k} - 1\right),$$

where R_k is the kth smallest log-return in the sample under study. Furthermore, it is shown in [7] that the empirical Expected Shortfall estimator at level p is computed as

$$\widehat{\mathrm{ES}}_p = \frac{1}{p}\left(\sum_{k=1}^{\lfloor np \rfloor} \frac{\mathrm{L}_k}{n} + \left(p - \frac{\lfloor np \rfloor}{n}\right) \mathrm{L}_{\lfloor np \rfloor + 1}\right).$$

Losses at the longer time horizon of 50 weeks are computed using Monte Carlo simulation. 1000 samples, each consisting of 50 log-returns, are constructed by drawing with replacement from the respective empirical log-return samples. The log-returns are then summed, and the losses are computed as above with the modification that R_k is the kth smallest sum of 50 log-returns. This gives a sample of $n = 1000$ losses from which the risk metrics can be computed.

Empirical Value-at-Risk and Expected Shortfall estimates are computed at levels $p \in \{0.05, 0.01\}$ using the respective samples of AUD/USD and CHF/JPY trading returns, and are presented in Table 11. It follows that the CHF/JPY trading model is preferrable from a risk perspective, which is in line with what one should expect given the model's higher information ratio (0.9235 as compared to 0.5754 for the AUD/USD model).

	AUD/USD		CHF/JPY	
Horizon (weeks)	1	50	1	50
$\widehat{\mathrm{VaR}}_{0.05}$	0.0273	0.1374	0.0209	0.0733
$\widehat{\mathrm{VaR}}_{0.01}$	0.0391	0.2009	0.0306	0.1229
$\widehat{\mathrm{ES}}_{0.05}$	0.0375	0.1789	0.0278	0.1077
$\widehat{\mathrm{ES}}_{0.01}$	0.0558	0.2358	0.0375	0.1675

Table 11: Empirical Value-at-Risk and Expected Shortfall from trading the respective AUD/USD and CHF/JPY exchange rate pairs using weekly FFNN models. The values presented are the relative shares of the portfolio value that is lost.

6.3 Summary of Results

The results presented in this section are of interest for a number of reasons. Firstly, it is indicated that FFNN trading models on average perform better than PNN trading models when applied to the FX market. However, the network design conducted in this paper can most likely be improved, and one should be careful not to reject the capabilities of PNNs based solely on the results presented here. Furthermore, profits appear hard to come around in the highly liquid EUR/USD market, which can be interpreted as to verify the established hypothesis that liquidity increases efficiency in financial markets.

The results presented indicate that the key success factor for an ANN trading model is its ability to predict *large* price movements. This has implications with regards to the neural network training process; perhaps the training data set should only comprise observations of large price movements (this idea is proposed in e.g. Kaastra and Boyd (1996) [8]). Based on the results stemming from the AUD/USD exchange rate pair, it can be argued that serial correlation in exchange rate returns, when present, likely increases the probability of success for ANN trading models. This should apply to other types of forecasting techniques as well, and comes as no surprise.

The overall performance of the ANN trading models developed in this paper is modest, at best. The Information ratios generated by the top performing models are quite high, but by no means outstanding. The returns generated by the weekly FFNN CHF/JPY trading model are high, but the underlying reasons for the success of the model are not elaborated upon in detail in this paper. The black-box nature of ANNs remains problematic in the sense that high returns, when achieved, may appear somewhat arbitrary.

The Value-at-Risk and Expected Shortfall analysis presented in this section is a nice supplement to more commonly used risk measures such as annualised volatility. The results presented are furthermore in line with expectations (given the annualised volatility figures), and can be used to compare the performance of the two most successful ANN trading models presented in this paper with investment opportunities available in other asset classes.

7 Concluding Remarks

This paper comprises performance evaluations of various ANN-based trading models applied to the FX market. By backtesting the respective models on all 45 permutations of the G10 exchange rate pairs, it is found that the average performance of the networks is modest, at best. On a few occasions quite high Information ratios are generated by the models, but the underlying economics are not analysed in detail. As such, the success of these models may simply be a result of data mining. Furthermore, it is found that FFNNs on average outperform PNNs with regards to the returns accumulated during the backtesting period.

There are, however, good reasons not to overgeneralise the conclusions arrived at in this paper. Inherent in the neural network design process is an infinite number of degrees of freedom, and a good portion of arbitrariness. In addition, the trading model could be refined in a number of ways to better exploit the full potential of the ANNs. Possible improvements with respect to these two issues are suggested below.

The inputs to the neural networks are of great importance for their performance, both in the case of FFNNs and PNNs. It would probably be a good idea to analyse in greater depth the economic drivers of the respective exchange rate pairs before specifying the inputs to the respective networks. However, this implies a significant workload, and to perform such analysis for all 45 permutations of exchange rate pairs is beyond the scope of this thesis. Regardless, refining the use nation specific inputs would likely improve the performance of the networks. Another aspect of the design process that applies to FFNNs and PNNs alike is the selection of window lengths, i.e. the question of how much historical data to use for the training of the networks. The windows used in this thesis were chosen quite arbitrarily, and in this regard there might well be room for improvement.

The list of design parameters is seemingly endless. Regarding the architecture of the FFNNs used in this paper, the configuration of the hidden nodes could probably be improved. The respective numbers of hidden layers and nodes in each layer affect the networks' ability to learn. However, there is no way to know a priori how the nodes should be organised as to maximise the performance of the respective networks, and in the end, this too boils down to a matter of trial and error. The same type of reasoning can be applied regarding the assignment of standard deviations to the Gaussians in the PNN setting. Although some measures are taken in this paper as to find a suitable σ, the procedure is by no means perfect.

As discussed in Section 5, the trading model implemented in this paper does not take into account the carry effect of investing the cash held between trades. It would be good to include this effect in the model, ideally using data on implied interest rates from the FX forward market. Furthermore, the trading model could be refined as to comprise e.g. stop-loss thresholds or the possibility of trading many exchange rate pairs simultaneously. The latter could potentially lower the volatility and increase the information ratio.

The paper would benefit from the inclusion of other statistical prediction techniques, to be used as benchmarks against which the ANNs could be evaluated. E.g. the observed autocorrelation in the AUD/USD exchange rate time series could probably be exploited using tools other than ANNs.

To conclude, the results presented in this paper cannot be used to discard completely the potential gains from applying ANNs in the FX market. Some models do perform well, but the overall impression is that of a modest performance. However, given the tremendous daily turnover of the FX spot market, the fact that it is reasonably efficient is not entirely surprising.

References

[1] Chen, A., Leung, M., and Daouk, H. Application of neural networks to an emerging financial market: Forecasting and trading the taiwan stock index. *Computers & Operations Research 30*, 6 (2003), 901 – 923.

[2] Dunis, C., and Williams, M. Applications of advanced regression analysis for trading and investment. In *Applied Quantitative Methods for Trading and Investment (ed. Dunis, C.L. and Laws, J. and Naïm, P.)*. John Wiley & Sons, Inc., Chichester, United Kingdom, 2003, pp. 1–40.

[3] Fine, T. *Feedforward Neural Network Methodology.* Springer, New York, NY, USA, 1999.

[4] Hagan, M., and Menhaj, M. Training feedforward networks with the Marquardt algorithm. *IEEE Transactions on Neural Networks 5* (November 1994), 989–993.

[5] Hornik, K., Stinchcombe, M., and White, H. Multilayer feedforward networks are universal approximators. *Neural Networks 2* (1989), 359–366.

[6] Hudson Beale, M., Hagan, M., and Demuth, H. *Neural Network Toolbox – User's Guide*, 7.0.1 ed. The MathWorks, Inc., Natick, MA, USA, April 2011.

[7] Hult, H., Lindskog, F., Hammarlid, O., and Rehn, C. Risk and portfolio analysis. Draft avaialble at the Department of Mathematics, Royal Institute of Technology, Stockholm, 2011.

[8] Kaastra, I., and Boyd, M. Designing a neural network for forecasting financial and economic time series. *Neurocomputing 10* (April 1996), 215–236.

[9] Marsland, S. *Machine Learning.* Chapman & Hall/CRC, Boca Raton, FL, USA, 2009.

[10] Mehta, M. Foreign exchange markets. In *Neural Networks in the Capital Markets (ed. Refenes, A.P.)*. John Wiley & Sons, Inc., Chichester, United Kingdom, 1995, pp. 176–198.

[11] Parzen, E. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics 33* (1962), 1065–1076.

[12] Shmilovici, A., Kahiri, Y., Ben-Gal, I., and Hauser, S. Measuring the efficiency of the intraday forex market with a universal data compression algorithm. *Computational Economics 33*, 2 (March 2009), 131–154.

[13] Specht, D. Vectorcardiographic diagnosis using the polynomial discriminant method of pattern recognition. *IEEE Transactions on Biomedical Engineering BME-14* (1967), 90–95.

[14] von Kleist, K., Mallo, C., Mesny, P., and Grouchko, S. Triennial central bank survey – report on global foreign exchange market activity in 2010. Tech. rep., Bank for International Settlements–Monetary and Economic Department, Basel, Switzerland, December 2010.

[15] Wasserman, P. *Advanced Methods in Neural Computing.* John Wiley & Sons, Inc., New York, NY, USA, 1993.

[16] Zhang, G., and Hu, M. Neural network forecasting of the british pound/us dollar exchange rate. *Omega 26*, 4 (1998), 495 – 506.

[17] Zhang, G., Patuwo, E., and Hu, M. Forecasting with artificial neural networks: The state of the art. *International Journal of Forecasting 14* (1998), 35 – 62.

A Tables

Variable	EcoWin Ticker	Time (GMT)	Description
AUD/USD	ew:aus1900110	21:15	Spot price of AUD in USD.
USD/CAD	ew:can19001	21:15	Spot price of USD in CAD.
EUR/USD	ew:emu19101	21:15	Spot price of EUR in USD.
USD/JPY	ew:jpn19001	21:15	Spot price of USD in JPY.
NZD/USD	ew:nzl1900110	21:15	Spot price of NZD in USD.
USD/NOK	ew:nor19001	21:15	Spot price of USD in NOK.
USD/SEK	ew:swe19001	21:15	Spot price of USD in SEK.
USD/CHF	ew:che19001	21:15	Spot price of USD in CHF.
GBP/USD	ew:gbr1900110	21:15	Spot price of GBP in USD.
Swap rate USD	ew:usa14912	21:15	2 year swap rate. Ask price.
Swap rate AUD	ew:aus14886	10:30	2 year swap rate. Ask price.
Swap rate CAD	ew:can14886	21:14	2 year swap rate. Ask price.
Swap rate EUR	ew:emu14882	21:15	2 year swap rate. Ask price.
Swap rate JPY	ew:jpn14886	21:14	2 year swap rate. Ask price.
Swap rate NZD	ew:nzl14886	03:35	2 year swap rate. Ask price.
Swap rate NOK	ew:nor14886	21:15	2 year swap rate. Ask price.
Swap rate SEK	ew:swe14800	21:15	2 year swap rate. Ask price.
Swap rate CHF	ew:che14886	21:15	2 year swap rate. Ask price.
Swap rate GBP	ew:gbr14886	21:15	2 year swap rate. Ask price.
S&P500	ew:usa15510	22:00	S&P 500 Composite Index.
ASX	ew:aus15500	05:30 or 08:00	ASX All Ordinaries Index.
S&P/TSX	ew:can15585	23:45	S&P/TSX Composite 60 Index.
STOXX	ew:emu15550	19:45/20:45*	STOXX Blue Chip 50 Index.
Nikkei	ew:jpn15500	07:00	Nikkei 225 Index.
NZX	ew:nzl15520	04:45	NZX 50 Index.
OSEBX	ew:nor15595	17:15	OSE Benchmark Index.
OMX	ew:swe15580	22:15	OMXS30 Index.
SMI	ew:che15415	Unknown	Swiss Market Index.
FTSE	ew:gbr15500	16:30	FTSE 100 Index.
WTI Oil	ew:com20460	00:45	WTI Light Crude Oil. Spot price in USD.

Table 12: Overview of the variables retrieved from EcoWin. Column 3 depicts the time of the day at which the observations are made. For the respective equity indices, the time series retrieved consist of close prices. Where spot prices are retrieved, the price quoted is the average between the bid price and the ask price.
*During summer time/winter time

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.7559	-0.1939	0.7160	-0.0973
Annualised return	0.0502	-0.0222	0.0473	-0.0106
Annualised volatility	0.0986	0.1047	0.0981	0.1048
Information ratio	0.5091	-0.2116	0.4820	-0.1010
Annualised transaction costs	0.0188	0.0040	0.0111	0.0021
Number of trades	1455	255	867	137
Winning trades	51.27%	55.69%	49.94%	41.61%
Losing trades	48.73%	44.31%	50.06%	58.39%
Average gain	0.0047	0.0101	0.0047	0.0107
Average loss	0.0044	0.0123	0.0044	0.0116
Average gain/Average loss	1.0735	0.8207	1.0590	0.9261
Winning up periods	48.31%	49.12%	32.57%	17.11%
Winning down periods	52.19%	56.25%	66.71%	79.41%

Table 13: USD/CAD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.3730	0.2267	0.0177	0.1150
Annualised return	-0.0398	0.0215	0.0015	0.0114
Annualised volatility	0.1053	0.0985	0.1049	0.0985
Information ratio	-0.3779	0.2180	0.0143	0.1156
Annualised transaction costs	0.0184	0.0037	0.0104	0.0022
Number of trades	1422	237	817	140
Winning trades	48.10%	49.79%	48.47%	55.71%
Losing trades	51.90%	50.21%	51.53%	44.29%
Average gain	0.0048	0.0108	0.0048	0.0104
Average loss	0.0048	0.0098	0.0048	0.0102
Average gain/Average loss	0.9849	1.1001	1.0059	1.0267
Winning up periods	49.23%	46.80%	38.04%	48.00%
Winning down periods	48.20%	52.40%	61.43%	52.80%

Table 14: USD/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.2210	0.6180	-0.0898	-0.0920
Annualised return	-0.0215	0.0513	-0.0080	-0.0100
Annualised volatility	0.1431	0.1433	0.1431	0.1434
Information ratio	-0.1502	0.3581	-0.0560	-0.0696
Annualised transaction costs	0.0185	0.0039	0.0098	0.0014
Number of trades	1433	251	771	92
Winning trades	47.87%	49.40%	48.77%	53.26%
Losing trades	52.13%	50.60%	51.23%	46.74%
Average gain	0.0067	0.0152	0.0066	0.0142
Average loss	0.0066	0.0144	0.0067	0.0155
Average gain/Average loss	1.0065	1.0597	0.9791	0.9139
Winning up periods	52.20%	54.74%	57.09%	69.12%
Winning down periods	45.74%	47.91%	42.50%	28.37%

Table 15: NZD/USD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	1.0995	0.5122	0.0612	1.3805
Annualised return	0.0666	0.0439	0.0051	0.0944
Annualised volatility	0.1300	0.1345	0.1300	0.1340
Information ratio	0.5126	0.3268	0.0392	0.7043
Annualised transaction costs	0.0186	0.0040	0.0104	0.0019
Number of trades	1438	257	818	124
Winning trades	51.74%	52.53%	50.98%	57.26%
Losing trades	48.26%	47.47%	49.02%	42.74%
Average gain	0.0062	0.0144	0.0061	0.0142
Average loss	0.0060	0.0135	0.0062	0.0136
Average gain/Average loss	1.0367	1.0732	0.9776	1.0420
Winning up periods	47.28%	49.34%	29.98%	33.19%
Winning down periods	54.84%	52.77%	70.05%	73.80%

Table 16: USD/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.0575	0.6677	-0.4864	-0.0955
Annualised return	0.0049	0.0546	-0.0554	-0.0104
Annualised volatility	0.1335	0.1266	0.1335	0.1268
Information ratio	0.0365	0.4317	-0.4149	-0.0819
Annualised transaction costs	0.0185	0.0038	0.0104	0.0015
Number of trades	1432	245	812	99
Winning trades	48.60%	50.61%	46.55%	47.47%
Losing trades	51.40%	49.39%	53.45%	52.53%
Average gain	0.0064	0.0139	0.0061	0.0122
Average loss	0.0061	0.0123	0.0064	0.0142
Average gain/Average loss	1.0480	1.1325	0.9489	0.8600
Winning up periods	46.37%	43.42%	29.35%	15.35%
Winning down periods	51.40%	56.99%	68.28%	84.56%

Table 17: USD/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.3408	0.1983	-0.0980	-0.0009
Annualised return	-0.0356	0.0190	-0.0088	-0.0001
Annualised volatility	0.1161	0.1145	0.1162	0.1146
Information ratio	-0.3066	0.1659	-0.0755	-0.0008
Annualised transaction costs	0.0180	0.0041	0.0112	0.0019
Number of trades	1394	261	875	125
Winning trades	47.35%	48.28%	48.57%	48.00%
Losing trades	52.65%	51.72%	51.43%	52.00%
Average gain	0.0053	0.0123	0.0053	0.0115
Average loss	0.0055	0.0113	0.0054	0.0121
Average gain/Average loss	0.9602	1.0890	0.9850	0.9528
Winning up periods	44.97%	46.98%	33.19%	28.88%
Winning down periods	54.14%	51.49%	66.03%	70.52%

Table 18: USD/CHF trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.0719	0.5127	-0.0699	-0.2590
Annualised return	0.0061	0.0440	-0.0062	-0.0307
Annualised volatility	0.0952	0.1001	0.0953	0.1002
Information ratio	0.0636	0.4394	-0.0648	-0.3062
Annualised transaction costs	0.0187	0.0037	0.0116	0.0023
Number of trades	1450	236	909	148
Winning trades	48.76%	52.97%	48.62%	43.92%
Losing trades	51.24%	47.03%	51.38%	56.08%
Average gain	0.0044	0.0106	0.0045	0.0102
Average loss	0.0045	0.0102	0.0045	0.0107
Average gain/Average loss	0.9935	1.0399	1.0074	0.9502
Winning up periods	53.49%	53.99%	48.10%	51.71%
Winning down periods	47.26%	51.90%	51.04%	44.73%

Table 19: GBP/USD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	1.0316	-0.6213	-0.2549	-0.0737
Annualised return	0.0636	-0.0960	-0.0248	-0.0079
Annualised volatility	0.1117	0.1036	0.1123	0.1045
Information ratio	0.5692	-0.9269	-0.2212	-0.0759
Annualised transaction costs	0.0187	0.0039	0.0121	0.0019
Number of trades	1444	250	946	124
Winning trades	50.28%	46.40%	49.58%	52.42%
Losing trades	49.72%	53.60%	50.42%	47.58%
Average gain	0.0051	0.0095	0.0051	0.0105
Average loss	0.0050	0.0113	0.0051	0.0104
Average gain/Average loss	1.0332	0.8401	0.9913	1.0108
Winning up periods	54.18%	44.53%	48.65%	47.92%
Winning down periods	48.86%	45.53%	49.86%	50.21%

Table 20: AUD/CAD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	1.1650	-0.1813	0.2393	0.4047
Annualised return	0.0695	-0.0206	0.0185	0.0360
Annualised volatility	0.1180	0.1111	0.1183	0.1110
Information ratio	0.5888	-0.1854	0.1565	0.3242
Annualised transaction costs	0.0185	0.0039	0.0101	0.0028
Number of trades	1431	248	792	179
Winning trades	51.15%	48.39%	50.51%	51.96%
Losing trades	48.85%	51.61%	49.49%	48.04%
Average gain	0.0052	0.0111	0.0049	0.0107
Average loss	0.0049	0.0106	0.0053	0.0110
Average gain/Average loss	1.0746	1.0390	0.9276	0.9658
Winning up periods	50.19%	46.12%	25.00%	32.76%
Winning down periods	51.41%	48.13%	76.50%	72.39%

Table 21: EUR/AUD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.6131	0.4534	0.0523	0.1462
Annualised return	-0.0793	0.0397	0.0044	0.0143
Annualised volatility	0.1861	0.1767	0.1854	0.1768
Information ratio	-0.4261	0.2245	0.0236	0.0808
Annualised transaction costs	0.0183	0.0038	0.0091	0.0026
Number of trades	1418	247	709	165
Winning trades	48.38%	55.47%	50.21%	52.12%
Losing trades	51.62%	44.53%	49.79%	47.88%
Average gain	0.0070	0.0158	0.0071	0.0163
Average loss	0.0076	0.0169	0.0075	0.0163
Average gain/Average loss	0.9145	0.9339	0.9510	1.0015
Winning up periods	52.44%	56.03%	64.66%	58.16%
Winning down periods	46.96%	51.38%	35.00%	41.28%

Table 22: AUD/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.1822	0.8715	0.0368	-0.2502
Annualised return	0.0147	0.0674	0.0031	-0.0295
Annualised volatility	0.0801	0.0771	0.0802	0.0776
Information ratio	0.1832	0.8739	0.0386	-0.3804
Annualised transaction costs	0.0190	0.0041	0.0108	0.0018
Number of trades	1472	262	850	115
Winning trades	50.41%	59.16%	49.41%	54.78%
Losing trades	49.59%	40.84%	50.59%	45.22%
Average gain	0.0037	0.0086	0.0037	0.0082
Average loss	0.0037	0.0080	0.0037	0.0085
Average gain/Average loss	0.9961	1.0718	1.0115	0.9667
Winning up periods	48.65%	53.01%	41.90%	41.73%
Winning down periods	53.09%	58.97%	57.75%	53.85%

Table 23: AUD/NZD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.2789	-0.2434	0.0536	0.0552
Annualised return	-0.0280	-0.0286	0.0045	0.0056
Annualised volatility	0.1219	0.1232	0.1219	0.1232
Information ratio	-0.2300	-0.2320	0.0367	0.0454
Annualised transaction costs	0.0184	0.0039	0.0121	0.0031
Number of trades	1427	253	948	196
Winning trades	50.60%	47.43%	48.63%	51.53%
Losing trades	49.40%	52.57%	51.37%	48.47%
Average gain	0.0052	0.0113	0.0054	0.0117
Average loss	0.0055	0.0118	0.0054	0.0113
Average gain/Average loss	0.9516	0.9602	0.9991	1.0352
Winning up periods	48.95%	51.53%	49.30%	43.13%
Winning down periods	51.49%	45.38%	51.11%	56.72%

Table 24: AUD/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.1839	0.8764	-0.3319	0.8351
Annualised return	0.0148	0.0676	-0.0339	0.0652
Annualised volatility	0.1191	0.1120	0.1193	0.1121
Information ratio	0.1242	0.6038	-0.2841	0.5816
Annualised transaction costs	0.0181	0.0041	0.0115	0.0025
Number of trades	1400	261	899	159
Winning trades	50.07%	52.49%	49.17%	56.60%
Losing trades	49.93%	47.51%	50.83%	43.40%
Average gain	0.0054	0.0113	0.0054	0.0110
Average loss	0.0053	0.0108	0.0054	0.0111
Average gain/Average loss	1.0355	1.0534	0.9958	0.9954
Winning up periods	51.06%	55.56%	52.04%	53.26%
Winning down periods	48.16%	53.14%	45.30%	58.16%

Table 25: AUD/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.2341	-0.5008	-0.3794	-0.0885
Annualised return	0.0185	-0.0697	-0.0400	-0.0096
Annualised volatility	0.1473	0.1447	0.1470	0.1450
Information ratio	0.1253	-0.4818	-0.2720	-0.0662
Annualised transaction costs	0.0181	0.0039	0.0108	0.0024
Number of trades	1402	248	849	157
Winning trades	48.43%	46.77%	49.82%	50.32%
Losing trades	51.57%	53.23%	50.18%	49.68%
Average gain	0.0062	0.0130	0.0059	0.0134
Average loss	0.0060	0.0142	0.0063	0.0139
Average gain/Average loss	1.0306	0.9118	0.9398	0.9654
Winning up periods	49.12%	46.91%	55.19%	59.64%
Winning down periods	50.66%	47.56%	44.70%	38.67%

Table 26: AUD/CHF trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.1522	-0.2347	-0.4591	0.1154
Annualised return	0.0124	-0.0274	-0.0512	0.0114
Annualised volatility	0.1225	0.1175	0.1227	0.1175
Information ratio	0.1012	-0.2336	-0.4172	0.0972
Annualised transaction costs	0.0188	0.0037	0.0112	0.0023
Number of trades	1456	236	877	146
Winning trades	49.04%	48.73%	46.52%	47.95%
Losing trades	50.96%	51.27%	53.48%	52.05%
Average gain	0.0054	0.0106	0.0054	0.0108
Average loss	0.0054	0.0119	0.0055	0.0118
Average gain/Average loss	1.0148	0.8875	0.9756	0.9146
Winning up periods	47.99%	46.48%	26.96%	22.07%
Winning down periods	52.03%	53.66%	68.82%	76.31%

Table 27: GBP/AUD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.1054	-0.0996	-0.4158	-0.0497
Annualised return	-0.0096	-0.0108	-0.0449	-0.0053
Annualised volatility	0.1018	0.1001	0.1025	0.1000
Information ratio	-0.0947	-0.1084	-0.4383	-0.0529
Annualised transaction costs	0.0184	0.0040	0.0125	0.0020
Number of trades	1420	254	983	126
Winning trades	48.31%	43.70%	49.44%	42.86%
Losing trades	51.69%	56.30%	50.56%	57.14%
Average gain	0.0049	0.0114	0.0048	0.0106
Average loss	0.0049	0.0103	0.0051	0.0109
Average gain/Average loss	0.9988	1.1022	0.9460	0.9734
Winning up periods	49.44%	48.12%	37.63%	26.36%
Winning down periods	49.83%	45.21%	61.10%	72.03%

Table 28: EUR/CAD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.2385	0.1960	0.1809	0.5029
Annualised return	0.0188	0.0188	0.0143	0.0433
Annualised volatility	0.1486	0.1483	0.1479	0.1482
Information ratio	0.1264	0.1267	0.0968	0.2921
Annualised transaction costs	0.0180	0.0035	0.0094	0.0022
Number of trades	1389	225	739	141
Winning trades	49.53%	52.00%	47.90%	53.90%
Losing trades	50.47%	48.00%	52.10%	46.10%
Average gain	0.0066	0.0148	0.0065	0.0154
Average loss	0.0065	0.0151	0.0065	0.0144
Average gain/Average loss	1.0123	0.9840	0.9950	1.0716
Winning up periods	52.17%	61.65%	47.52%	56.77%
Winning down periods	48.21%	40.17%	53.85%	44.44%

Table 29: CAD/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.9256	-0.0979	-0.0293	-0.4930
Annualised return	0.0587	-0.0107	-0.0025	-0.0682
Annualised volatility	0.1162	0.1174	0.1168	0.1170
Information ratio	0.5049	-0.0908	-0.0218	-0.5831
Annualised transaction costs	0.0185	0.0040	0.0126	0.0021
Number of trades	1428	258	988	135
Winning trades	50.98%	47.29%	47.98%	49.63%
Losing trades	49.02%	52.71%	52.02%	50.37%
Average gain	0.0058	0.0132	0.0058	0.0116
Average loss	0.0055	0.0121	0.0055	0.0136
Average gain/Average loss	1.0567	1.0914	1.0527	0.8517
Winning up periods	50.87%	46.72%	46.13%	39.00%
Winning down periods	50.39%	47.30%	51.30%	58.92%

Table 30: NZD/CAD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.3435	0.2254	-0.3329	-0.2310
Annualised return	0.0260	0.0214	-0.0340	-0.0270
Annualised volatility	0.1129	0.1098	0.1132	0.1098
Information ratio	0.2305	0.1945	-0.3006	-0.2455
Annualised transaction costs	0.0179	0.0040	0.0096	0.0022
Number of trades	1383	254	754	144
Winning trades	49.10%	48.82%	48.28%	45.14%
Losing trades	50.90%	51.18%	51.72%	54.86%
Average gain	0.0055	0.0116	0.0053	0.0121
Average loss	0.0054	0.0115	0.0057	0.0111
Average gain/Average loss	1.0118	1.0146	0.9390	1.0891
Winning up periods	49.10%	52.55%	41.23%	39.61%
Winning down periods	52.16%	50.20%	59.25%	51.84%

Table 31: CAD/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.2007	0.2226	-0.1298	-0.0054
Annualised return	-0.0193	0.0211	-0.0118	-0.0006
Annualised volatility	0.1141	0.1066	0.1146	0.1067
Information ratio	-0.1692	0.1981	-0.1032	-0.0053
Annualised transaction costs	0.0179	0.0039	0.0102	0.0021
Number of trades	1386	252	797	137
Winning trades	49.71%	46.43%	50.44%	47.45%
Losing trades	50.29%	53.57%	49.56%	52.55%
Average gain	0.0055	0.0120	0.0055	0.0115
Average loss	0.0056	0.0110	0.0056	0.0114
Average gain/Average loss	0.9785	1.0899	0.9785	1.0061
Winning up periods	50.30%	47.06%	44.36%	37.25%
Winning down periods	49.36%	52.24%	56.14%	62.86%

Table 32: CAD/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.0100	0.3349	-0.2026	-0.3240
Annualised return	-0.0009	0.0305	-0.0192	-0.0399
Annualised volatility	0.1252	0.1263	0.1253	0.1264
Information ratio	-0.007	0.2414	-0.153	-0.3157
Annualised transaction costs	0.0184	0.0038	0.0116	0.0020
Number of trades	1423	246	911	130
Winning trades	48.63%	61.38%	49.51%	49.23%
Losing trades	51.37%	38.62%	50.49%	50.77%
Average gain	0.0057	0.0120	0.0056	0.0127
Average loss	0.0058	0.0140	0.0058	0.0130
Average gain/Average loss	0.9790	0.8596	0.9624	0.9747
Winning up periods	49.42%	57.20%	44.52%	39.77%
Winning down periods	51.62%	54.66%	56.24%	56.36%

Table 33: CAD/CHF trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.0081	-0.0742	-0.2757	0.6272
Annualised return	-0.0007	-0.0080	-0.0272	0.0519
Annualised volatility	0.099	0.1006	0.0993	0.1002
Information ratio	-0.0071	-0.0794	-0.2740	0.5181
Annualised transaction costs	0.0185	0.0036	0.0121	0.0025
Number of trades	1431	233	953	162
Winning trades	50.38%	47.64%	48.90%	53.09%
Losing trades	49.62%	52.36%	51.10%	46.91%
Average gain	0.0048	0.0103	0.0047	0.0108
Average loss	0.0048	0.0113	0.0049	0.0108
Average gain/Average loss	0.9968	0.9043	0.9744	0.9967
Winning up periods	46.57%	43.04%	39.32%	33.91%
Winning down periods	53.51%	59.26%	59.65%	72.22%

Table 34: GBP/CAD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.4394	-0.3359	0.3540	0.4635
Annualised return	0.0322	-0.0417	0.0263	0.0404
Annualised volatility	0.1286	0.1219	0.1288	0.1219
Information ratio	0.2502	-0.3419	0.2038	0.3314
Annualised transaction costs	0.0187	0.0036	0.0098	0.0014
Number of trades	1446	232	771	89
Winning trades	50.55%	52.59%	52.53%	55.06%
Losing trades	49.45%	47.41%	47.47%	44.94%
Average gain	0.0057	0.0114	0.0056	0.0122
Average loss	0.0055	0.0132	0.0056	0.0124
Average gain/Average loss	1.0382	0.8681	0.9980	0.9879
Winning up periods	54.49%	56.93%	57.29%	59.55%
Winning down periods	45.40%	42.49%	43.88%	46.35%

Table 35: EUR/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.0224	-0.2647	-0.1525	0.1239
Annualised return	-0.0020	-0.0315	-0.0140	0.0122
Annualised volatility	0.1179	0.1125	0.1181	0.1125
Information ratio	-0.0167	-0.2798	-0.1189	0.1086
Annualised transaction costs	0.0190	0.0041	0.0101	0.0015
Number of trades	1469	266	794	95
Winning trades	50.03%	47.37%	49.12%	53.68%
Losing trades	49.97%	52.63%	50.88%	46.32%
Average gain	0.0055	0.0121	0.0053	0.0116
Average loss	0.0056	0.0117	0.0058	0.0122
Average gain/Average loss	0.9734	1.0333	0.9263	0.9524
Winning up periods	49.08%	47.21%	25.85%	20.17%
Winning down periods	51.95%	46.07%	73.82%	80.15%

Table 36: EUR/NZD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.2398	0.3328	0.5956	-0.0690
Annualised return	0.0189	0.0303	0.0408	-0.0074
Annualised volatility	0.0746	0.0788	0.0741	0.0790
Information ratio	0.2531	0.3848	0.5499	-0.0938
Annualised transaction costs	0.0184	0.0041	0.0112	0.0021
Number of trades	1422	263	878	135
Winning trades	47.61%	52.09%	50.68%	49.63%
Losing trades	52.39%	47.91%	49.32%	50.37%
Average gain	0.0033	0.0079	0.0034	0.0070
Average loss	0.0033	0.0074	0.0032	0.0085
Average gain/Average loss	0.9970	1.0722	1.0361	0.8274
Winning up periods	49.61%	50.87%	34.49%	23.91%
Winning down periods	52.69%	52.96%	67.73%	79.26%

Table 37: EUR/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.0021	-0.2396	-0.1145	-0.1416
Annualised return	0.0002	-0.0281	-0.0103	-0.0158
Annualised volatility	0.0732	0.0669	0.0733	0.0670
Information ratio	0.0026	-0.4199	-0.1411	-0.2351
Annualised transaction costs	0.0187	0.0039	0.0116	0.0018
Number of trades	1445	252	912	116
Winning trades	48.93%	48.81%	46.60%	35.34%
Losing trades	51.07%	51.19%	53.40%	64.66%
Average gain	0.0032	0.0061	0.0032	0.0065
Average loss	0.0032	0.0070	0.0032	0.0067
Average gain/Average loss	0.9861	0.8735	1.0171	0.9715
Winning up periods	50.96%	48.81%	45.26%	51.98%
Winning down periods	49.76%	49.60%	52.67%	44.76%

Table 38: EUR/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.1289	-0.1109	-0.0768	0.2613
Annualised return	-0.0119	-0.0121	-0.0068	0.0244
Annualised volatility	0.0710	0.0743	0.0707	0.0742
Information ratio	-0.168	-0.1635	-0.0963	0.3295
Annualised transaction costs	0.0182	0.0037	0.0113	0.0017
Number of trades	1404	236	885	111
Winning trades	46.72%	46.19%	49.94%	54.95%
Losing trades	53.28%	53.81%	50.06%	45.05%
Average gain	0.0026	0.0061	0.0026	0.0067
Average loss	0.0026	0.0063	0.0026	0.0058
Average gain/Average loss	1.0239	0.9730	0.9814	1.1518
Winning up periods	46.75%	51.16%	46.64%	39.53%
Winning down periods	50.28%	46.28%	53.33%	61.57%

Table 39: EUR/CHF trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.1958	0.1409	-0.5138	0.4719
Annualised return	0.0157	0.0138	-0.0598	0.0410
Annualised volatility	0.0819	0.0785	0.0823	0.0783
Information ratio	0.1915	0.1759	-0.7264	0.5241
Annualised transaction costs	0.0185	0.0039	0.0121	0.0026
Number of trades	1428	250	947	164
Winning trades	51.68%	51.60%	46.88%	58.54%
Losing trades	48.32%	48.40%	53.12%	41.46%
Average gain	0.0037	0.0084	0.0037	0.0084
Average loss	0.0039	0.0078	0.004	0.0077
Average gain/Average loss	0.958	1.0762	0.9389	1.0838
Winning up periods	52.85%	53.23%	48.49%	59.32%
Winning down periods	50.90%	45.99%	48.29%	45.57%

Table 40: EUR/GBP trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.0165	-0.0597	-0.2001	0.0798
Annualised return	-0.0014	-0.0064	-0.0189	0.0080
Annualised volatility	0.1798	0.1736	0.1793	0.1736
Information ratio	-0.0080	-0.0367	-0.1055	0.0462
Annualised transaction costs	0.0181	0.0038	0.0092	0.0019
Number of trades	1403	241	724	123
Winning trades	49.25%	50.62%	47.38%	52.03%
Losing trades	50.75%	49.38%	52.62%	47.97%
Average gain	0.0076	0.0164	0.0073	0.0165
Average loss	0.0076	0.0169	0.0079	0.0168
Average gain/Average loss	0.9923	0.9697	0.9316	0.9787
Winning up periods	50.44%	51.43%	70.14%	62.86%
Winning down periods	49.81%	49.09%	28.21%	35.91%

Table 41: NZD/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.6255	0.0120	-0.3972	-0.1097
Annualised return	0.0432	0.0012	-0.0424	-0.0120
Annualised volatility	0.1563	0.1560	0.1559	0.1560
Information ratio	0.2762	0.0079	-0.2717	-0.0770
Annualised transaction costs	0.0192	0.0039	0.0082	0.0019
Number of trades	1488	249	643	122
Winning trades	50.87%	49.80%	48.37%	50.00%
Losing trades	49.13%	50.20%	51.63%	50.00%
Average gain	0.0067	0.0149	0.0065	0.0147
Average loss	0.0068	0.0161	0.0070	0.0162
Average gain/Average loss	0.9841	0.9259	0.9332	0.9102
Winning up periods	55.04%	49.46%	66.58%	66.79%
Winning down periods	47.63%	55.16%	31.43%	32.74%

Table 42: NOK/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.1900	0.5703	-0.3117	-0.2892
Annualised return	-0.0182	0.0481	-0.0314	-0.0349
Annualised volatility	0.1602	0.1483	0.1597	0.1483
Information ratio	-0.1134	0.3241	-0.1969	-0.2351
Annualised transaction costs	0.0184	0.0039	0.0107	0.0023
Number of trades	1425	249	838	149
Winning trades	50.74%	49.40%	47.61%	41.61%
Losing trades	49.26%	50.60%	52.39%	58.39%
Average gain	0.0069	0.0155	0.0068	0.0152
Average loss	0.0069	0.0142	0.0070	0.0145
Average gain/Average loss	1.0031	1.094	0.9751	1.0451
Winning up periods	51.15%	56.49%	58.47%	38.93%
Winning down periods	47.45%	44.54%	39.96%	55.04%

Table 43: SEK/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.1803	0.0058	3.2307	0.3365
Annualised return	0.0145	0.0006	0.1313	0.0306
Annualised volatility	0.1325	0.1320	0.1320	0.1319
Information ratio	0.1096	0.0046	0.9946	0.2322
Annualised transaction costs	0.0186	0.0039	0.0102	0.0020
Number of trades	1439	248	803	127
Winning trades	47.81%	50.00%	52.30%	48.03%
Losing trades	52.19%	50.00%	47.70%	51.97%
Average gain	0.0057	0.0136	0.0058	0.0136
Average loss	0.0055	0.0124	0.0054	0.0123
Average gain/Average loss	1.0470	1.0940	1.0806	1.1025
Winning up periods	51.01%	46.96%	52.17%	34.41%
Winning down periods	47.62%	48.62%	52.69%	64.82%

Table 44: GBP/JPY trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.2135	0.1312	-0.1449	-0.4771
Annualised return	0.0170	0.0129	-0.0133	-0.0652
Annualised volatility	0.1244	0.1294	0.1245	0.1291
Information ratio	0.1364	0.0998	-0.1068	-0.5052
Annualised transaction costs	0.0186	0.0038	0.0117	0.0021
Number of trades	1439	241	918	133
Winning trades	49.55%	46.06%	48.04%	45.11%
Losing trades	50.45%	53.94%	51.96%	54.89%
Average gain	0.0059	0.0138	0.0059	0.0125
Average loss	0.0058	0.0122	0.0059	0.0134
Average gain/Average loss	1.0181	1.1253	1.0003	0.9300
Winning up periods	50.00%	48.08%	43.78%	41.15%
Winning down periods	50.25%	47.92%	55.54%	52.92%

Table 45: NZD/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.1507	0.1587	-0.0892	-0.4162
Annualised return	-0.0141	0.0154	-0.008	-0.0544
Annualised volatility	0.1232	0.1175	0.1235	0.1173
Information ratio	-0.1146	0.1314	-0.0645	-0.4642
Annualised transaction costs	0.0186	0.0040	0.0112	0.0022
Number of trades	1440	259	878	138
Winning trades	49.31%	47.88%	52.39%	54.35%
Losing trades	50.69%	52.12%	47.61%	45.65%
Average gain	0.0058	0.0133	0.0058	0.0115
Average loss	0.0060	0.0117	0.0060	0.0134
Average gain/Average loss	0.9709	1.1357	0.9766	0.8547
Winning up periods	50.63%	49.43%	52.73%	52.47%
Winning down periods	49.86%	46.41%	47.69%	46.41%

Table 46: NZD/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.5811	-0.4382	-0.5885	0.0348
Annualised return	-0.0729	-0.0582	-0.0731	0.0036
Annualised volatility	0.1450	0.1420	0.1447	0.1422
Information ratio	-0.5028	-0.4100	-0.5053	0.0251
Annualised transaction costs	0.0186	0.0039	0.0118	0.0024
Number of trades	1438	252	928	155
Winning trades	47.57%	50.79%	49.25%	52.90%
Losing trades	52.43%	49.21%	50.75%	47.10%
Average gain	0.0063	0.0136	0.0063	0.0145
Average loss	0.0066	0.0147	0.0066	0.0138
Average gain/Average loss	0.9464	0.9204	0.9676	1.0509
Winning up periods	49.38%	50.75%	54.11%	55.64%
Winning down periods	48.68%	44.87%	42.19%	41.45%

Table 47: NZD/CHF trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.0551	-0.3516	-0.0005	0.0786
Annualised return	-0.0049	-0.0441	0.0000	0.0079
Annualised volatility	0.1197	0.1152	0.1200	0.1153
Information ratio	-0.0411	-0.3824	-0.0004	0.0685
Annualised transaction costs	0.0183	0.0037	0.0106	0.0019
Number of trades	1416	238	828	121
Winning trades	50.99%	43.70%	49.76%	50.41%
Losing trades	49.01%	56.30%	50.24%	49.59%
Average gain	0.0056	0.0117	0.0056	0.0121
Average loss	0.0057	0.0123	0.0058	0.0119
Average gain/Average loss	0.9707	0.9530	0.9699	1.0173
Winning up periods	47.10%	47.68%	28.31%	21.94%
Winning down periods	53.65%	47.53%	70.75%	75.67%

Table 48: GBP/NZD trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.3191	-0.0770	-0.2327	-0.0525
Annualised return	0.0244	-0.0083	-0.0224	-0.0056
Annualised volatility	0.0786	0.0744	0.0787	0.0744
Information ratio	0.3103	-0.1116	-0.2847	-0.0751
Annualised transaction costs	0.0181	0.0039	0.0121	0.0023
Number of trades	1399	249	948	145
Winning trades	47.53%	51.81%	47.68%	56.55%
Losing trades	52.47%	48.19%	52.32%	43.45%
Average gain	0.0038	0.0074	0.0036	0.0078
Average loss	0.0035	0.0075	0.0037	0.0071
Average gain/Average loss	1.0695	0.9810	0.9693	1.0867
Winning up periods	50.28%	45.85%	45.44%	44.27%
Winning down periods	49.00%	53.04%	53.68%	50.20%

Table 49: NOK/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.4219	0.0336	0.1089	-0.4487
Annualised return	-0.0466	0.0034	0.0089	-0.0600
Annualised volatility	0.1060	0.1148	0.1055	0.1144
Information ratio	-0.439	0.0300	0.0842	-0.5247
Annualised transaction costs	0.0187	0.0040	0.0120	0.0022
Number of trades	1443	258	944	142
Winning trades	45.95%	48.45%	50.11%	47.18%
Losing trades	54.05%	51.55%	49.89%	52.82%
Average gain	0.0043	0.0104	0.0043	0.0096
Average loss	0.0044	0.0105	0.0044	0.0113
Average gain/Average loss	0.9767	0.9887	0.9790	0.8482
Winning up periods	47.47%	53.50%	43.11%	20.99%
Winning down periods	49.38%	47.86%	58.68%	74.32%

Table 50: CHF/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.6350	-0.2899	-0.0562	-0.2324
Annualised return	0.0437	-0.0350	-0.0049	-0.0271
Annualised volatility	0.1026	0.1035	0.1028	0.1035
Information ratio	0.4257	-0.338	-0.0480	-0.2621
Annualised transaction costs	0.0178	0.0041	0.0110	0.0019
Number of trades	1380	266	861	121
Winning trades	49.93%	49.25%	51.22%	52.07%
Losing trades	50.07%	50.75%	48.78%	47.93%
Average gain	0.0048	0.0101	0.0047	0.0101
Average loss	0.0047	0.0109	0.0049	0.0110
Average gain/Average loss	1.0398	0.9210	0.9594	0.9188
Winning up periods	49.24%	51.28%	30.40%	23.08%
Winning down periods	52.24%	46.62%	69.27%	72.93%

Table 51: GBP/NOK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.1406	0.1391	-0.5970	-0.5083
Annualised return	0.0115	0.0136	-0.0748	-0.0712
Annualised volatility	0.1090	0.1090	0.1085	0.1086
Information ratio	0.1056	0.1250	-0.6892	-0.6553
Annualised transaction costs	0.0177	0.0040	0.0112	0.0020
Number of trades	1365	259	880	131
Winning trades	49.82%	56.37%	46.36%	45.80%
Losing trades	50.18%	43.63%	53.64%	54.20%
Average gain	0.0045	0.0093	0.0043	0.0096
Average loss	0.0044	0.0106	0.0046	0.0102
Average gain/Average loss	1.0355	0.8763	0.9479	0.9467
Winning up periods	50.59%	56.15%	42.07%	25.41%
Winning down periods	48.69%	53.13%	53.70%	62.11%

Table 52: CHF/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	0.4632	-0.4382	-0.4973	-0.0899
Annualised return	0.0337	-0.0582	-0.0571	-0.0097
Annualised volatility	0.1057	0.0972	0.1056	0.0976
Information ratio	0.3184	-0.5987	-0.5406	-0.0999
Annualised transaction costs	0.0185	0.0037	0.0107	0.0029
Number of trades	1431	240	838	183
Winning trades	50.17%	49.17%	49.28%	46.45%
Losing trades	49.83%	50.83%	50.72%	53.55%
Average gain	0.0049	0.0095	0.0047	0.0097
Average loss	0.0048	0.0107	0.0050	0.0106
Average gain/Average loss	1.0176	0.8826	0.9379	0.9183
Winning up periods	49.40%	48.07%	33.98%	33.05%
Winning down periods	52.35%	46.82%	63.61%	67.04%

Table 53: GBP/SEK trading results.

	FFNN daily	FFNN weekly	PNN daily	PNN weekly
Cumulative return	-0.0453	-0.0569	0.0801	0.0571
Annualised return	-0.0040	-0.0061	0.0066	0.0058
Annualised volatility	0.1038	0.1029	0.1037	0.1029
Information ratio	-0.0388	-0.0590	0.0638	0.0562
Annualised transaction costs	0.0187	0.0040	0.0098	0.0026
Number of trades	1444	254	767	166
Winning trades	49.65%	50.79%	45.76%	48.80%
Losing trades	50.35%	49.21%	54.24%	51.20%
Average gain	0.0045	0.0099	0.0045	0.0096
Average loss	0.0045	0.0100	0.0046	0.0103
Average gain/Average loss	1.0075	0.9924	0.9838	0.9364
Winning up periods	48.97%	47.76%	44.32%	40.41%
Winning down periods	50.31%	51.37%	57.11%	63.53%

Table 54: GBP/CHF trading results.