An Artificial Neural Network Application

Predicting the Nordic Electric Spot Market

Abstract

This paper studies the Nordic electric-market relationship between one-week and one-month forward market prices versus realized average next week and next month spot prices. The forward products are traded at the Nordpool Financial Market and the spot-market prices are settled at Nordpool Spot. The microstructures of the forward and spot markets are used to evaluate a formal model based on actual reaction from new market information from the limited number of market participants. The objective is directional prediction ability applying an Artifical Neural Network (ANN) to predict next week and month spot market prices. From ANN directional accuracy short and long positions are traded at the Nordpool financial market for delivery in the spot market inducing risk. The forecasts are evaluated with both statistical and economic criteria. In terms of statistical criteria, the ANN model shows directional accuracy at 5% for most of the years and at the 1% level for the whole period. In terms of economic criteria, the paper shows a directional accuracy for short (majority) and long forward market position that is surprisingly high. For the whole period the weekly contracts show a success-ratio of 2.61 and for monthly contracts the success-ratio are as high as 7. The ANN prediction model therefore shows very promising results, inducing first that the spot energy market may inefficiently price electricity and second, application of the ANN-prediction model can be very profitable at the Nordpool Spot electricity market bearing some risk.

Classification: C45; C53; Q41; C13

Keywords: Financial and Spot Markets, Artificial Neural networks, Prediction Model, Risk

1 Introduction

The main objective of this paper is to investigate whether it is possible to exploit the existence of nonlinearities in the Nordic energy market to improve forecasting over especially short time horizons. For markets in general, many empirical studies have shown evidence rejecting the Efficient Market Hypothesis (EMH). At the same time, the interest in non-linear financial time series models has greatly increased. The main reason is that these models are capable of explaining observed aspects such as asymmetric information, differences in target and negotiations times and that agents with complex algorithms might be able to make better us of available information. Once a linear model is rejected in favour of some form for nonlinear model, a wealth of possible nonlinear structures can be used to describe and forecast financial time series. In this paper the out-of-sample forecasts are analyzed using artificial neural networks. Artificial Neural Network models (ANNs) are nowadays used in a large variety of modelling and forecasting problems. In line with the increased interest in non-linear models, the last two decades have shown that neural networks can be used in financial applications. Conferences such as “Neural Networks in the Capital Markets”, a large number of books (i.e. Gately, 1996) and articles in scientific journals dealing with financial applications of neural networks (for an overview see Qi, 1996), characterizes the increased popularity[1].

The main reason for the use of ANN is that these models seem to be able to approximate any nonlinear function arbitrarily close. Hence, as shown by Solibakke (2002, 2006) the hydro-electric time series is characterized by truly nonlinear dynamic relationships, the ANN will detect any nonlinear function and will therefore provide a superior fit compared to linear time series models. The ANN model’s parameters are difficult, if not impossible, to interpret. Hence, an estimated ANN model does not necessary provides information on which type of parametric model might be suitable to describe the nonlinear patterns. This paper is no exception; the ANN

[1] Reviews of ANN from statistical and econometric perspectives can be found in Cheng and Titterington (1994) and Kuan and White (1994), respectively.

model is constructed mainly for the purpose of pattern recognition and forecasting. Moreover, the author is well aware of the overfitting problem; superior in-sample fitting is no guarantee for that an ANN performs well in out-of-sample forecasting. The number of hidden units will therefore always be kept to a minimum.

The ANN procedure that is adopted for establishing a single predicted output for the spot prices, are therefore a combination of fundamental and technical details. The formulation of the problem is established using a formal model for the market participants (microstructure) and all factors affecting the spot prices are evaluated and weighted using sigmoid functions in ANN hidden units. The results are encouraging. Since the models first prediction in 1999/2000, its extensions and modifications, the model has in majority of cases made very good predictions, making a significant contribution to overall profits. The overall results show that the model has a very significant $N(0,1)$ directional accuracy test statistic of about 6.7 for the whole period from 1999 to 2007. Moreover, the profit distribution is skewed to the right suggesting that holding other variables equal between mistakes to successes, the results show higher positive than negative profits.

The rest of this article is now organised as follows. Section 2 gives the ANN representation. Section 3 conducts an evaluation of the spot and one-week forward data series and Section 4 builds the necessary model for establishing the ANN model. Section 5 show actual processing and spot predictions with evaluation of forecasting fit and finally, section 5 contains summarises and conclusions.

2 Artificial Neural Networks Representation

Starting with a STAR (Smooth Transition AutoRegressive) model for a univariate time series y_t, which may be price, return or absolute return on a financial asset

$y_t = \phi_0 + \beta_1 \cdot G\left(\gamma\left[y_{t-1} - c\right]\right) + \varepsilon_t$, where G(-) is the logistic (sigmoid) function $G(z) = \dfrac{1}{1+\exp(-z)}$. This model describes the situation where the conditional mean of y_t depends on the value of y_{t-1} relative to the threshold c. For $y_{t-1} << c$, the conditional mean of y_t is equal to ϕ_0, while it changes gradually to $\phi_0 + \beta_1$ as y_{t-1} increases. From this STAR representation, an Artificial Neural Network (ANN) can now be obtained by assuming that the conditional mean of y_t depends on the value of a linear combination of p lagged values of y_{t-1}, ..., y_{t-p} relative to threshold c. The SETAR model above then becomes $y_t = \phi_0 + \beta_1 \cdot G\left(\gamma\left[\tilde{x}_t^{'} \cdot \delta - c\right]\right) + \varepsilon_t$, where $\tilde{x}_t = \left(y_{t-1}, ..., y_{t-p}\right)^{'}$, or after some manipulation $y_t = \phi_0 + \beta_1 \cdot G\left(\tilde{x}_t^{'}\gamma_1\right) + \varepsilon_t$, where $\tilde{x}_t = \left(1, \tilde{x}_t^{'}\right)^{'}$ and the individual elements of the parameter vector $\gamma_1 = \left(\gamma_{0,1}, \gamma_{1,1}, ..., \gamma_{p,1}\right)^{'}$, are easily obtained from γ, δ and c. The ANN can be interpreted as a switching-regression model, where the switching is determined by a particular linear combination of the p lagged variables in the vector x_t. In our application of neural networks we will not focus on this regime switching. Instead, the aim is to model the, possibly nonlinear, relationship between y_t and x_t. The usual way to do this is to include additional logistic components in the model, which gives $y_t = \phi_0 + \sum_{j=1}^{q} \beta_j \cdot G\left(\tilde{x}_t^{'}\gamma_j\right) + \varepsilon_t$. Suppose that an appropriate relationship between y_t and x_t is given by $y_t = g\left(x_t ; \xi\right) + \eta_t$, where $g\left(x_t ; \xi\right)$ is a continuous function. It can be shown that ANNs of this form can approximate any such function $g\left(x_t ; \xi\right)$ to any desired degree of accuracy, provided that the number of nonlinear components q is sufficiently large. The description ANN(k,q,1) is normally used to identify that the network has k input variables $y_{t-1}, \ldots, y_{t-p}$, $z_{1,t}, \ldots, z_{m,t}$, q logistics components $G\left(x_t^{'}\gamma_j\right)$ and one (1) output variable y_t. Because we consider only the case with a single output, we abbreviate to ANN(k,q).

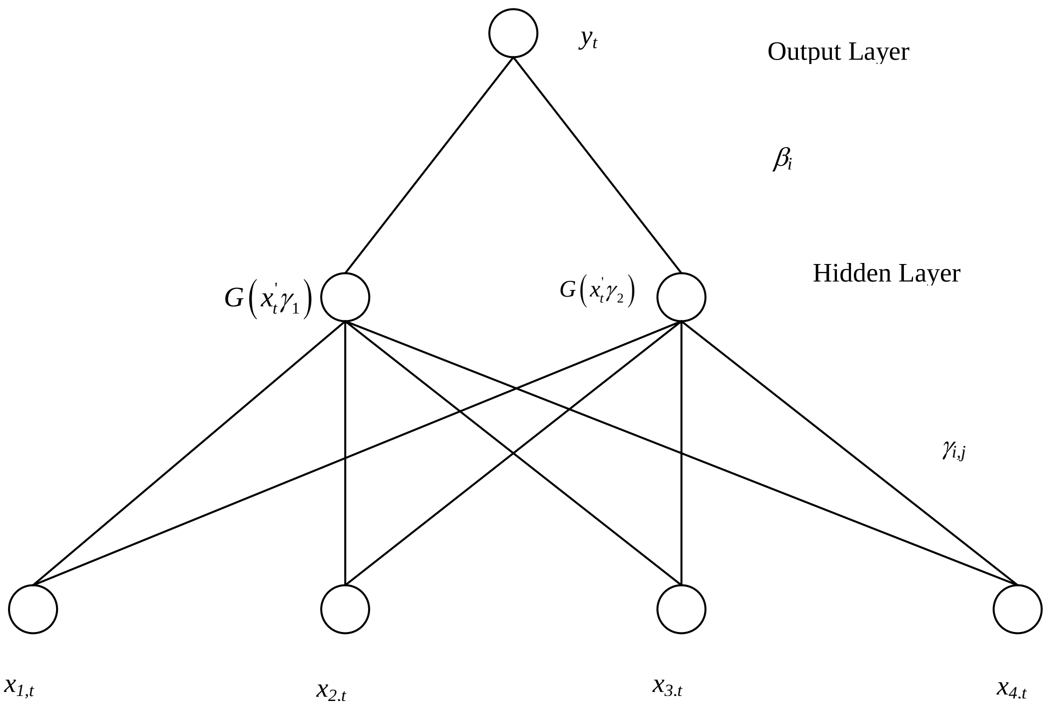

Figure 1. Graphical Representation of single hidden layer feed forward neural network ANN(k,q); k = 4 and q = 2.

The terminology which is commonly used in discussions of neural networks is rather different from usual econometric practice (see Kuan and White, 1994; Warner and Misra, 1996, for more extensive descriptions). Consider the graphical representation of the ANN(k,q) model with k = 4 input variables and q = 2 hidden layers shown in Figure 1. The network is seen to consist of three different layers. At the basis is the input layer, consisting of the explanatory variables in x_t, which usually are called input variables. These inputs are multiplied by so-called connection strengths $\gamma_{i,j}$ as they enter the hidden layer, which consists of q hidden units – that is, the logistic (sigmoid) function G(-). The name "hidden layer" arises from the fact that it is not directly observed. In the hidden layer, the linear combination $x_t^{'}\gamma_j$ are formed and transformed into a value between 0 and 1 by the activation functions G(-). Finally, these are multiplied by weights β_j to produce the output y_t. This type of ANN is usually referred to as single hidden layer feed forward network model, because it contains only one hidden layer and information flows only in one direction, from inputs to outputs. Extensions of the model to allow for multiple hidden layers or some form of feedback are possible. Importantly, to gain some insight in the properties of the nonlinear part of the model, it is appropriate to inspect the joint distribution of the nonlinear components; that is, the hidden layers together or the individual contributions of each hidden layer unit.

3 Nordpool Spot Market and the Nordpool Financial Market

On the Nordpool Sport market (Elspot), hourly power contracts are traded daily for physical delivery in the next day's 24-hour period. The price calculation is based on the balance between bids and offers from all market participants – finding the intersection point between the market's supply curve and demand curve. This trading method is referred to as equilibrium point trading, auction trading, or simultaneous price setting. The price mechanism in Elspot adjusts the flow of power across the interconnectors - and also on certain connections within the Norwegian grid - to the available trading capacity given by the Nordic transmission system operators. Elspot is therefore a common power market for the Nordic countries with an implicit capacity auction on the interconnectors between the bidding areas. The Elspot market's system price is denoted "the unconstrained market clearing price", because the system price is the price that balances sale and purchase in the exchange area while not considering any transmission constraints. Finally, the use of the concept implicit auctions is important for the market. That is, the spot concept is based on bids for purchase and sale of hourly contracts using three different bidding types: hourly bids, block bids and flexible hourly bids that cover some or all of the 24 hours of the next day.

The financial market is a commercial centre where price securing contracts are traded. At present the contract types traded on Nord Pool's Financial Market comprise of power derivates, electricity certificates and EUA's. The derivates are base load futures, forwards, options, and Contracts for Difference. The reference price for these contracts is the System Price of the unconstrained total Nordic power market (Elspot). The maximum trading time horizon is currently four/five years. There is no physical delivery of financial market electricity contracts. Cash settlement is made throughout the delivery period, starting at the due date of each contact.

The initial contract launched for the electricity certificates market is a spot contract with physical delivery.The EUA’s are forward contracts with physical delivery.

4 Models for Spot Market Predictions and Data

The interest in primary commodities are centuries old, dating back, e.g. to Smith (1776) who give an elaborate analysis of the price quotation of wheat in England from 1202 to 1764[2]. Out standing features of this publication (and many others in the literature) is the abruptness of changes in both level and volatility of commodity prices. Large fluctuations in the prices of primary commodities regularly invoke political actions, both domestically and internationally. Crop failures in North Korea and Central Africa lead to international assistance. Falling US grain prices build momentum for Congress to rush emergency financial assistance for US farmers. The larger than anticipated declines in oil prices forces oil-producing countries to take macroeconomic measures in support of government spending and domestic currencies. The energy crisis in California 2001/2002 may also indicate the need for public regulations and ownership of the electric power industry. The experience from Scandinavia during December 2002 and the autumn of 2006, inducing very high spot- and forward-prices, may also suggest a need for close public regulations of market behaviour.

The formal model focuses on producers and consumers of electricity. Regulatory authorities are not explicitly modelled but are implicitly present through a high degree of ownership of the Scandinavian utilities. This high concentration and influence of public ownership in production companies together with an overall limited number of producers/suppliers in the total Nordic energy market may spur market power and information asymmetries. The energy market, in contrast to more liquid and symmetric markets, may therefore induce inefficiencies and therefore the ability of price predictability. Consider a competitive risk-averse producer of electric power, denoted by Q, that is deterministic because of diversifiable production. The assumption is that all the daily production is sold applying the one-day forward quoted Euro price (ElSpot). The quoted price at the one-day (week) future date, denoted by $\tilde{q}_{t+1}$, is random. The producer's revenues denominated in domestic currencies are given by $(\tilde{q} \cdot \tilde{e})Q$ where $\tilde{e}$ is the random foreign-exchange rate at a future date and are defined as the domestic price of foreign currency. Revenues are therefore the inextricable product of two random variables, which complicates solving the maximization problem.

The producer's output Q is assumed to be a linear function of inputs y_t and is deterministic. In this analysis the stochastic element (+/-20% energy) is ignored for convenience. However, it should not infer any changes to our conclusion regarding the spot price/exchange rate relationship. The output next day per producer is therefore based on the observed optimal demand/supply solution reported to Nordpool Spot at time *t*. The *Q* is therefore influenced by storage and price/currency considerations. Hence, $\tilde{q}_{t+1}$ defines the $Q_{t+1} = y_{t+1}$ production. Costs in domestic currency, is assumed to consist of a linear input cost and a quadratic production cost. The notation is now changed from *t+1* to *t* to enhance readability. Hence, the producer's profit function reads: $\tilde{V}_p = [\tilde{q}_t \cdot \tilde{e}_t]Q - \omega_t Q - \frac{1}{2}Q^2$, where ω_t is the proportional cost of input. A typical assumption is to assume that the producer (and other representative agents) maximizes a mean-variance utility function with an absolute risk aversion parameter α_p, a typical expression of the form: $EU_p = E\tilde{V}_p - \frac{\alpha_p}{2} Var(\tilde{V}_p)$. It is assumed that the first two moments of the exchange rate and the commodity price exists. Hence, $E(\tilde{e}) = \overline{e}, E(\tilde{q}) = \overline{q}, Var(\tilde{e}) = \sigma_e^2, Var(\tilde{q}) = \sigma_q^2$ are the respective price levels and conditional variances? Furthermore the covariance of $\tilde{e}$ and $\tilde{q}$ is

[2] See also The Economist, which, at regular intervals, publishes its own commodity price index.

given by $\rho(\tilde{q}_t,\tilde{e}_t)\sigma_e\sigma_q$ where ρ is a symbol denoting the correlation coefficient. The maximization of the mean-variance utility function above with respect to the production variable y_t gives: $E[V_p]=(E[q_t\cdot e_t]-\omega_t)Q-\frac{1}{2}Q^2 = AQ-\frac{1}{2}Q^2, \quad A = (E[q_t\cdot e_t]-\omega_t)$ and constant. Moreover, the variance becomes $Var(V_p)=Q^2Var(q_t\cdot e_t)=BQ^2$, $B = Var(q_t\cdot e_t)$ and is constant. Expected utility from these expressions therefore become

$EU_p = E\tilde{V}_p-\frac{\alpha_p}{2}Var(\tilde{V}_p)=AQ-\frac{1}{2}Q^2-\frac{\alpha_p}{2}BQ^2 = AQ-\frac{1}{2}(1+\alpha_p B)Q^2$, which has the maximum at $Q_{\max}=\frac{A}{1+\alpha_p B}$. Now applying Taylor series expansion for the expectation and variance, we can write the production decision using the following expression[3]

$$Q_{\max}=\frac{E[q_t\cdot e_t]-\omega_t}{1+\alpha_p\cdot Var(q_t\cdot e_t)}=\frac{\mu_{q_t}\cdot\mu_{e_t}-\omega_t}{1+\alpha_p\cdot\left[(\frac{\partial(q\cdot e)}{\partial q})^2\sigma_q^2+(\frac{\partial(q\cdot e)}{\partial e})^2\sigma_e^2+2(\frac{\partial(q\cdot e)}{\partial q})(\frac{\partial(q\cdot e)}{\partial e})Cov(q,e)\right]}$$

which is positively related to the expected market based commodity price μ_q and μ_e, but negatively related to the input cost ω_t. An increase in total variance $Var(q\cdot e)$, will therefore lead to lower production volumes. A higher stockpiling absorbs the lower volume of production. Hence, an increase in the variance for the spot price and the exchange rate would results in lower production by increasing stockpiling inducing higher production in periods with lower volatility. We assume that a typical producer maximizes the expected utility of profits, where its von-Neumann Morgenstern utility function $U(\cdot)$ satisfies $U'>0 \textit{ and } U''<0$. Producer's profits are assumed to be denominated in foreign currency and to follow from three activities. Producers perform stockholding by transferring commodities from periods with a low price to a period with a high price and thus evening out price fluctuations. Stockholding, I_t (reservoir), involves zero costs and is positively related to the difference of the future spot price $\tilde{q}_F$ and the current

[3] This is usually an assumption. However as the function is $X\cdot Y$, higher then the first derivative is zero implying close to perfect fit.

known spot price[4] q_{t-1}, the latter compounded by r_{t-1} (one plus the foreign interest rate). Next, the commodity price and exchange risk, the producer's optimal behaviour consists in the reduction or even removal of these risks by using risk-sharing tools. As the producer is long in the commodity, the producer can sell forward a quantity K_t of the commodity of the futures price $f_t^{\tilde{q}_F}$, the transaction adding $(f_t^{\tilde{q}_F} - q_{t-1})K_t$ to its foreign currency receipts. Hence, the producer's profits in foreign currency are expressed as:

$$\tilde{V} = \tilde{e}_t \left[\tilde{q}_t Q + (\tilde{q}_F^{Spot} - r \cdot q_{t-1}) I_t + (f_t^{q_F} - q_{t-1}) K_t \right]$$

The producer selects at each period *t* the variables (Q_t, I_t, K_t) so as to maximize $EU(\tilde{V}_m)$. The necessary and sufficient conditions for optimum are therefore:

$$\begin{aligned} EU'(\tilde{V}_m)\left[\tilde{q}_F^{Spot} - r \cdot q_{t-1} \right] &= 0 \\ EU'(\tilde{V}_m)\left[f_t^{\tilde{q}_F} - q_{t-1} \right] &= 0 \end{aligned}$$

where a prime indicates partial derivatives. These two last equations also suggest that $EU'(\tilde{V}_m)\left[f_t^{\tilde{q}_F} - r \cdot q_{t-1} \right] = 0$. That is, the storage level is decided by the fact that the forward price is the expected forward spot price. The optimal storage level of the producer at any date *t* is chosen so as to equate the storage return to the marginal cost of storage, which is close to zero. Maximizing the expected utility function with respect to Q_t, I_t *and* K_t gives the following optimal solutions for the stochastic income function $V = c(qQ + (q_F - rq)I + (f - q)K)$ and expected utility $EU = E\left[V\right] - \frac{\alpha_p}{2} Var\left[V\right]$. We use the notation $\mu_{cq} = E[cq], \quad \mu_{c^2q^2} = E[c^2q^2]$ etc.

The implications are the following expression

$$\begin{aligned} EU &= \mu_{cq}(Q - rI - K) + \mu_{cq_F} I + \mu_c fK \\ &+ \frac{\alpha}{2}(\mu_{cq}(Q - rI - K) + \mu_{cq_F} I + \mu_c fK)^2 \\ &- \frac{\alpha}{2}[\mu_{c^2q^2}(Q - rI - K)^2 + 2\mu_{c^2qq_F}(QI - rI^2 - KI) \\ &+ 2\mu_{c^2q} f(QK - rIK - K^2) + 2\mu_{c^2q_F} fIK + 2\mu_{c^2q_F^2} I^2 + \mu_{c^2} f^2K^2] \end{aligned}$$

where $\sigma_{cq} = Cov(c,q) = \mu_{cq} - \mu_c\mu_q = E[cq] - E[c]E[q]$ or we can write it as $\mu_{cq} = \sigma_{cq} + \mu_c\mu_q$.

[4] The notation is *t*-1 due to earlier change from *t*+1 to *t*. *t*-1 is therefore available information for decision making.

From these equations, we observe that the producer's optimal behaviour establishes a link between both the currency and the spot commodity-market prices through the correlation coefficient. Expected utility has origin from two components: the expected utility that derives from its main activity; that is the price determination, production volume and stockpiling, and that from risk management activity. If a producer's expectations do not deviate from the forward/future prices i.e. $\overline{q} = f_t^{q_F}$, the producer will hedge all transactions completely, the gains from speculation vanishes and only gains from stockpiling remain. Otherwise, the correlation between exchange rate and commodity price affects the utility of the producer, either directly by entering the *EU* above or indirectly through $\overline{q} - f_t^{q_F}$.

5 Data and In of Sample Forecasts

For spot price predictions the paper applies the one-week and the one month forward time series together with the above defined production and consumption variables. Weekly forward prices are available on daily basis from October 1995, while monthly prices are available from October 2004. Nordpool spot prices (Elspot) are available from the end of 1992 and are used to calculate existing and future volatility together with a calculation of the forward-spot difference. This difference is the speculative profit per kwh available to market participant. The actual profit/loss is dependent on actual size of positions in the market. Moreover, the spot-price changes series are used to calculate the BIC-optimal[5] ARMA-GARCH model with a constant, serial correlated mean adjusting for seasonal effects (holidays) and clustered conditional volatility (Engle, 1982). The AR-GARCH model is used for one-period-ahead volatility prediction and used as input to the ANN-model. Moreover, demand is accounted for applying weather forecast data available for minimum 16 days ahead. Hydro balance, demand (temperature) and supply incidents are all accounted for using available data from well established information sources already available for all market participants. Table 1 report the average spot and forward prices,

[5] Bayes Information Criterion (Schwarz, 1978). See also Ljung & Box (1978) for dimensions of a model.

Table 1. Market Information Summary

Period	Average prices contracts			Weekly Volatility	Fundamental Market Information					Market participants	
					Reservoar (+/- Normal)	Consumption		Production			
	Spot	Weekly Futures	Monthly Futures	GARCH(1,1)	+/- Relative to Normal	Gwh	% Change	Gwh	% Change	Spot	Financial
1997	14.75430	14.40648	-----	15.65923	11.32426	114648	----	111002	----	40	245
1998	13.45192	13.23748	-----	13.87851	7.56432	119692	4.4	115997	4.5	42	250
1999	13.88288	13.55553	-----	14.96064	-2.73077	119931	0.2	121935	5.1	43	300
2000	13.19792	13.11663	-----	15.54446	8.93588	122763	2.4	141834	16.3	45	300
2001	23.32753	23.34364	-----	12.86969	-5.83216	124610	1.5	121017	-14.7	43	325
2002	28.06961	29.95285	-----	13.47642	-10.13843	119986	-3.7	129668	7.1	45	350
2003	36.48690	36.71045	34.83027	14.39749	-53.35216	114224	-4.8	106550	-17.8	40	365
2004	29.42186	29.20244	29.64576	11.23252	-16.34115	120122	5.3	108652	2.4	38	380
2005	30.10677	30.20433	31.09280	11.91180	0.53647	124731	3.8	136904	26.0	43	400
2006	49.83553	49.73654	50.80250	12.89956	-17.44255	121167	-2.9	120371	-12.1	45	405
2007	27.23161	26.33625	26.02000	14.18863	13.46267	118619	-6.3	115347	-16.0	45	500

predicted GARCH volatility using all available historic information, consumption (actual versus normal), percent magazines and temperature information for the period from January 1999 to April 2007. Importantly, all information is known at prediction time; that is – no look-ahead biases.

An econometric formulation of a feed-forward (non-recursive)[6] single hidden layer artificial neural network process (ANN) is $y_h = F\left(\beta_{h0} + \sum_{j=1}^{q} G\left(\tilde{x}'\gamma_j\right)\cdot\beta_{hj}\right) \quad h = 1,\dots g$. The y_h is a gx1 vector of endogenous variables, $\tilde{x}_j = \left(1, x_1, \dots, x_k\right)'$ is a kx1 vector of explanatory variables, $\gamma_j = (\gamma_{j0}, \gamma_{j1}, \dots, \gamma_{jk})'$ is a k+1x1 vector of hidden weights, q is the number of hidden units, G is the transformation (sigmoid: $\frac{1}{(1+e^{-\tilde{x}'\cdot\gamma})}$) applied in the output layer. Hence, this is a system of g non-linear equations, with some common coefficients (γ) across equations. As for the AR-GARCH model above the BIC-optimal ANN is the preferred in-sample prediction model.

6 Results and out of sample forecasts

The out-of-sample forecasting ability of the BIC-optimal ANN models in terms of statistical accuracy and economic criteria is evaluated next. The analysis is based on one-step-ahead forecasts (which can of course be extended to several steps)[7]. The periods are suitable to describe the robustness of the results achieved. The main characteristics of the in-sample nonlinear ANN models are one hidden layer (q), several explanatory variables in the single hidden layer (x), and g is vectors of hidden weights. G(-) is the transformation applied in the hidden layer and β is a vector of output weights. The ANN models are “trained” using ordinary nonlinear least squares methods over a given period using a minimum of 500 iterations (or convergence). Once the model is estimated, we construct its one-period a head out-of-sample

[6] Learning by means of steepest descent is shown to be inefficient compared to ordinary least squares methods.
[7] Multiple steps forecasts are not used in actual applications neither for weekly nor monthly spot forecasts.

forecasts. The forecasts are used to buy or sell future/forward contracts speculating on the average next week/month Nordpool spot price. The predictions are performed using the latest market information and market prices /volatilities.

Based on success/failure classification from the short/long automated model decisions, parameters are either retained or discarded. The prediction model will then be ready for estimation for another week and month.

6 Evaluation of the out-of-sample forecasts

Forecasting with neural networks is analogous to forecasting with other parametric nonlinear models, such as STAR models. The main features can be summarized as follows. A 1-step of y_{n+1} can be computed directly from an ANN(p,q) model as: $\hat{y}_{t+1|t} = x_t^{'} \cdot \phi + \sum_{j=1}^{q} \beta_j \cdot G(x_t^{'} \cdot \gamma_j)$, where $x_t = (1, y_t, ..., y_{t-p+1})^{'}$. Under the additional assumption that the shocks ε_t are normally distribute, the 1-step-ahead forecast error $e_{t+1|t} = y_{t+1} - y_{t+1|t}$ is normally distributed (since $e_{t+1|t} = \varepsilon_{t+1}$ by definition) and forecast confidence intervals can be constructed in the usual way. For multiple step-ahead forecasts things become much more complicated. No closed-form expressions exist for $\hat{y}_{t+h|t}$ where h > 1 and one has to rely on simulation techniques to obtain such forecasts.

Table 2 shows statistics related to the out-of sample forecast performance of ANN(p,q) models for weekly spot prices from the Nordpool Spot market.

Table 2 Weekly and Monthly Forecast for Nordpool Spot Market prices from January 1999 to April 2007

Panel A Weekly Contracts from the first week of 1999

Average In-the-sample parameter values (only β reported)					Out-of-the-sample Prediction		Directional Positions		Ratios	Dicrectional		Assuming normally distributed	
Year	β_1	β_2	β_3	β_4	y_t	Realised	Short	Long	Success/Failure	Accurancy		P(S) (>=0)	P(F) (<0)
1999	-1580.541	1979.072	405.3493	-406.3685	13.6006	13.8829	706.41	-655.41	3.25000	3.71007	**	70.49 %	29.51 %
2000	-1307.999	1548.807	339.4873	-339.5650	14.8217	13.1979	686.29167	-634.2917	2.46667	1.80280	*	69.63 %	30.37 %
2001	-1187.093	1381.331	307.1549	-307.1991	22.8912	23.3275	1213.0317	-1161.032	3.72727	1.82837	*	70.29 %	29.71 %
2002	-1093.766	1263.113	318.7079	-318.8235	27.1498	28.0696	1487.6896	-1434.69	2.78571	1.71232	*	61.29 %	38.71 %
2003	-1021.552	1146.600	278.0715	-277.3614	36.7239	36.4869	1897.3188	-1845.319	2.46667	3.53607	**	63.81 %	36.19 %
2004	-1009.052	1130.456	270.7059	-270.5640	29.1154	29.4219	1529.9367	-1477.937	3.00000	1.58872		75.07 %	24.93 %
2005	-1001.730	1119.212	270.7135	-271.1618	29.8356	30.1068	1565.5521	-1513.552	1.73684	1.56778		57.80 %	42.20 %
2006	-974.9948	1109.342	280.4850	-280.4131	46.9632	49.8355	2591.4475	-2539.448	2.05882	1.95989	*	58.30 %	41.70 %
2007 (15)	-888.021	1073.496	279.5026	-279.0421	27.4655	27.2316	458.20583	-441.2058	2.00000	1.81334	*	57.89 %	42.11 %
1999-2007	-1118.306	1305.714	305.5753	-305.6109	27.6185	27.9512	12135.884	-11702.88	2.61022	6.47157	**	67.10 %	32.90 %

Panel B. Montly Contracts from the month of October 2003

Average In-the-sample parameter values					Out-of-the-sample Prediction		Directional Positions		Ratios	Dicrectional		Assuming normally distributed	
Year	β_1	β_2	β_3	β_4	y_t	Realised	Short	Long	Success/Failure	Accurancy		P(S) (>=0)	P(F) (<0)
2006	-27.6301	19.02253	45.9414	58.33726	45.87240	49.91624	9	3	4.00000	2.56234	**	84.33 %	15.67 %
2007 (04)	-106.3474	15.126817	130.2341	165.6481	31.29205	20.05299	2	2	3.00000	1.89346	*	67.28 %	32.72 %
2006-2007	-48.62138	17.983673	68.41945	86.9535	47.10687	47.36197	11	5	7.00000	4.25362	**	77.52 %	22.48 %

The predictions are mainly lower than the unbiased future contract price, inducing therefore short positions which are much larger in numbers than long future positions. The trading strategies are simple. A predicted price at or lower than the simultaneously quoted future price contracts induce a short position; that is – selling forward contracts that go to physical delivery (no closing out). With forecasted prices strictly above the simultaneously quoted future price contracts induce a long position; that is – buying forward contract that go to physical delivery (no closing out).

From Table 2 Panel A the ANNs clearly show predictive power. First, compared with linear models (OLS) (not reported[8]) the nonlinear components have predictive power. Secondly, focusing on Table 2, predicted values are on average lower than realized average values in 6 out of the 9 years. The market therefore seems to show a too high risk premium in future contracts. The number of short positions in the respective years confirms this average numbers. However, the relative low number of long positions relative to average numbers suggests a price distribution clearly skewed to the right. The predictive power of the neural network is clearly manifested in the success-failure (10) and the directional accuracy (11) columns (Pesaran & Timmermann, 1992). For all the years from 1999 to 2007 the number of successes relative to failures is greater than 2.5. That is, on average for 1 failure (wrong short/long position) you will find a minimum of 2.7 successes. For 2004 alone on average, the same ratio was the impressive number 7.67. For the whole period from 1999 to 2005 (week 17), the success-failure ratio was 4.4. Assuming symmetric price changes around the future contract prices relative to the time of trading the model should be very profitable. The numbers for directional accuracy shows for almost all years statistical significance at 5%. Moreover, for the whole period the significance is at 1%. Figure 1 plots the predicted prices versus realized values. The impressive 7.67 year 2004 suggest that the model perform best in low volatile periods. The more volatile years 2005 and

[8] Linear model results (OLS) are available from author upon request.

2006 show a clearly lower model fit. The overall probability for success is as high as 67.1% for weekly data series.

The same figures for monthly data are lower in number and therefore less reliable. However, the model predicts the next month's short/long positions with an impressive success – failure ratio greater or equal to 3 and with a success – total ratio greater or equal to 75%. On average the probability for success for monthly data series from 2004 to 2007 is around 77.5%. The predictions, forward prices and realized spot prices are reported in Figure 2.

Finally, a word of caution is justified. Even though the results are impressive, each month will be a new challenge. You never know whether the prediction is true or not. For weekly results there is on average a 32.5% probability for a wrong long/short position. Moreover, there are no escapes. When the contract goes to delivery, there are no "closing out" opportunities except trading day contracts mainly following the every day spot level. For monthly contracts the escape is weekly contracts, which however also follow closely the average level of the spot contracts. However, smaller contracts working in the long run will eventually perform and produce significant surpluses for market participants applying the prediction model.

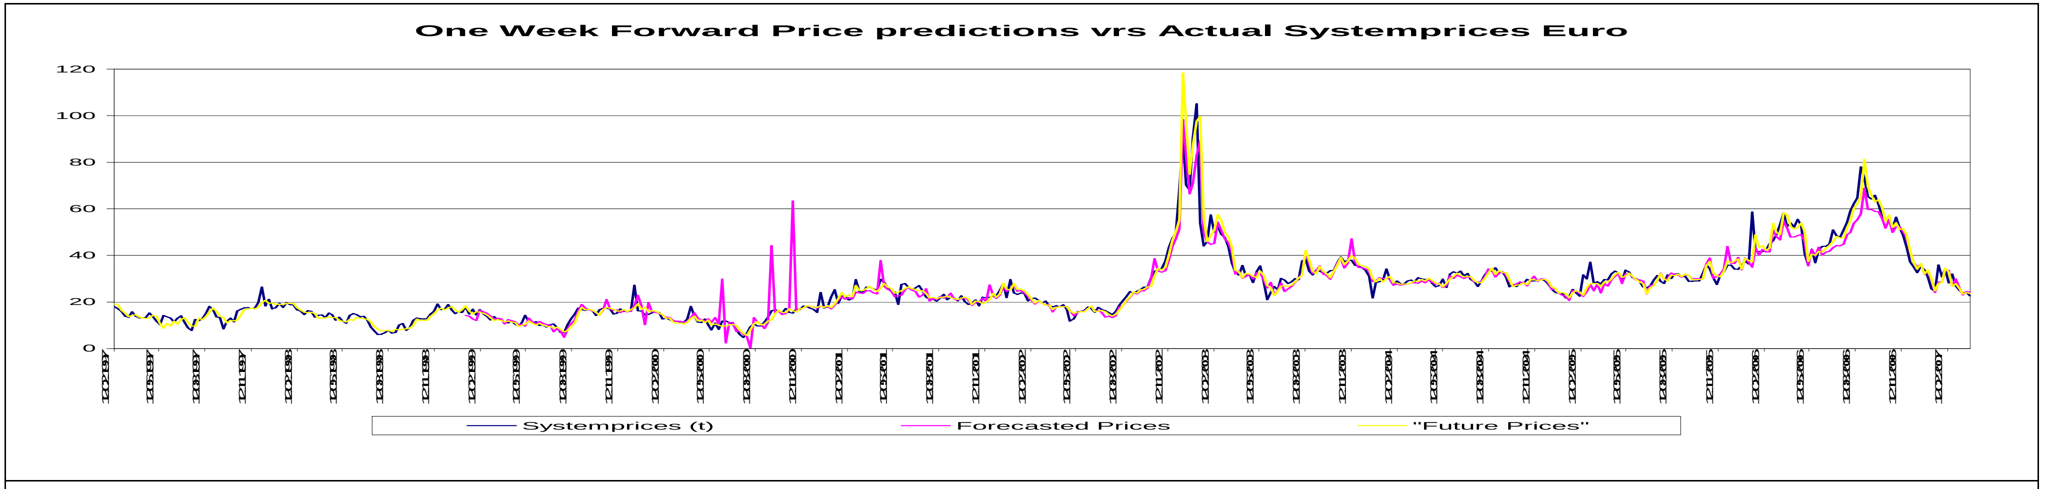

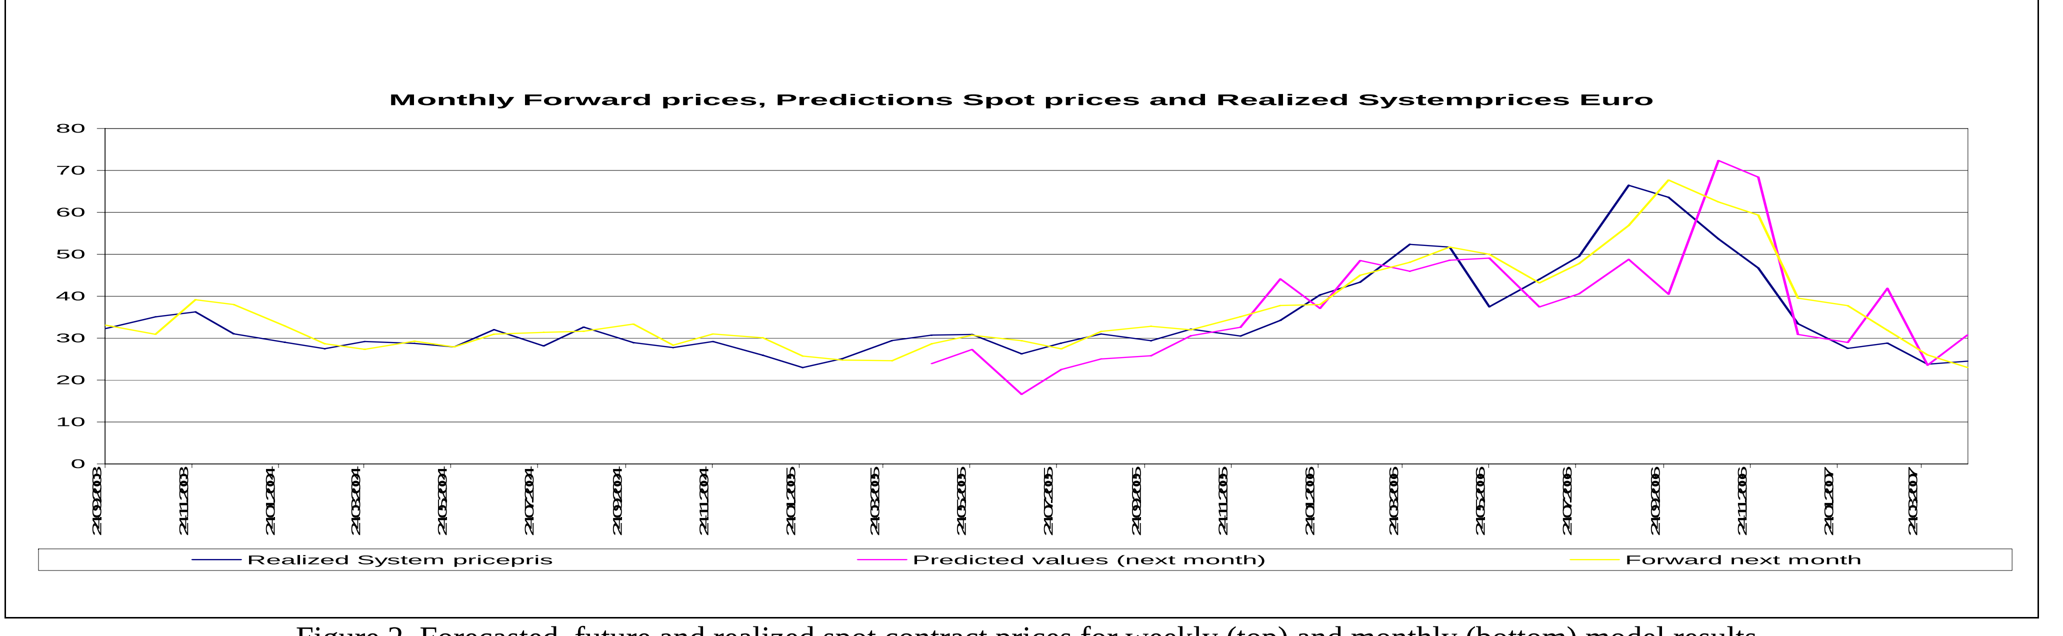

Figure 2. Forecasted, future and realized spot contract prices for weekly (top) and monthly (bottom) model results

7 Summary and Conclusions

The ANN model is evaluated based on achieved returns over the last 7.5 years for weekly contracts and for the last 2.5 years for monthly contracts data. The ANN is estimated using a feed-forward non-linear least square estimation (non-recursive) and one week and month forward spot predictions are obtained from the network. Based on the predictions directional positions are traded on the Nordpool Financial market; that is – either a short or a long forward/future contract. Success – Failure ratios and directional accuracy calculations suggest price predictability indicating an inefficient spot market relative to the forward/future market. The commodity spot pricing at Nordpool Spot are therefore not efficient and most likely attributable to the low number of participants following the formal spot pricing model outlined in this paper.

References

Cheng, B. and D.M.Titterington, 1994, Neural Networks: a review from a statistical perspective, Statistical Science 9, pp. 2-54.

Engle, R.F., 1982, Auto-regressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation, Econometrica, 50, 987-1008.

Gately, E., 1996, Neural Networks for Financial Forecasting, New York: John Wiley.

Kuan, C.-M. and H.White, 1994, Artificial Neural Networks: an econometric perspective, Econometric Reviews 13, pp. 1-143 (with discussions).

Ljung, G.M. and G.E.P. Box, 1978, On a Measure of Lack of Fit in Time Series Models, Biometrika, 66, pp. 67-72.

Qi, M., 1996, Financial applications of artificial neural networks, in G.S. Maddala and C.R. Rao (eds.), Handbooks of Statistics, 14, Amsterdam: Elsevier Science, pp. 529-552.

Pesaran, M.H. and A. Timmermann, 1992, A simple nonparametric test for predictive performance, Journal of Time Series Analysis 10, pp. 461-465.

Schwarz, G., 1978, Estimating the Dimension of a Model, Annals of Statistics, 6, pp. 461-464.

Smith, A., 1776, An Inquiry into the Nature and Causes of the Wealth of Nations. Strahan and Cadell, London

Solibakke, Per Bjarte, 2002, Efficiently Estimated Mean and Volatility Characteristics for the Nordic Spot Electric Power Market, International Journal of Business, 7(2), 17-35.

Solibakke, Per Bjarte, 2006, Describing the Nordic Forward Electric Power Market: A Stochastic Model Approach, International Journal of Business, 7(2), 17-35.

Warner, B. and M. Misra, 1996, Understanding Neural Networks as statistical tools, American Statistican 50, pp. 284-93.