

SFB 649 Discussion Paper 2010-043

Meteorological forecasts and the pricing of weather derivatives

Matthias Ritter*
Oliver Mußhoff**
Martin Odening***



*Humboldt-Universität zu Berlin and Georg-August-Universität
Göttingen, Germany

**Georg-August-Universität Göttingen, Germany

***Humboldt-Universität zu Berlin, Germany

This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".

<http://sfb649.wiwi.hu-berlin.de>
ISSN 1860-5664

SFB 649, Humboldt-Universität zu Berlin
Spandauer Straße 1, D-10178 Berlin



SFB 649 ECONOMIC RISK BERLIN

Meteorological forecasts and the pricing of weather derivatives*

Matthias Ritter^{ab**} Oliver Mußhoff^a Martin Odening^b

September 1, 2010

In usual pricing approaches for weather derivatives, forward-looking information such as meteorological weather forecasts is not considered. Thus, important knowledge used by market participants is ignored in theory. By extending a standard model for the daily temperature, this paper allows the incorporation of meteorological forecasts in the framework of weather derivative pricing and is able to estimate the information gain compared to a benchmark model without meteorological forecasts. This approach is applied for temperature futures referring to New York, Minneapolis and Cincinnati with forecast data 13 days in advance. Despite this relatively short forecast horizon, the models using meteorological forecasts outperform the classical approach and more accurately forecast the market prices of the temperature futures traded at the Chicago Mercantile Exchange (CME). Moreover, a concentration on the last two months or on days with actual trading improves the results.

Keywords: Weather forecasting, weather risk, price forecasting, financial markets, temperature futures, CME

JEL classification: C53, G13, G17, N23

1. Introduction

Weather derivatives are relatively new financial instruments to insure against economic consequences of weather risk such as fluctuations of temperature or precipitation. Most weather derivatives are traded at the Chicago Mercantile Exchange (CME), where futures and options on temperature, hurricanes, snowfall and frost are offered.

Because the underlying, the weather, is not tradeable, the market is incomplete, making it difficult to determine prices theoretically. Most pricing approaches use statistical models based

*The financial support from the German Research Foundation via the CRC 649 “Economic Risk”, Humboldt-Universität zu Berlin, is gratefully acknowledged. Moreover, the authors would like to thank the participants of the CRC 649 Workshop “Weather Derivatives and Risk” and the CRC 649 Conference 2010 for their helpful comments and discussions on this topic.

^aDepartment of Agricultural Economics and Rural Development, Georg-August-Universität Göttingen

^bDepartment of Agricultural Economics, Humboldt-Universität zu Berlin

**Corresponding author: Matthias.Ritter@agrar.hu-berlin.de

on historical weather data to simulate prices. Meteorological weather forecasts, however, are usually not included into the model. Thus, important knowledge about weather in the future and the resulting payoff is ignored.

There are only a few papers dealing with the incorporation of meteorological forecasts into pricing models for weather derivatives. Alaton et al. [1] mention that meteorological forecasts should be used for short-term pricing. For long-term pricing, they suggest considering a general trend which increases or lowers the parameters for the simulation without mentioning a specific method. Jewson & Caballero [16] describe how meteorological forecasts can be used for the pricing of weather derivatives. Via single and ensemble forecasts up to 12 days in advance, they derive probabilistic weather forecasts to price derivatives which have already begun or will begin in the forecast period only. Yoo [18] incorporates seasonal meteorological forecasts into a temperature model. These forecasts predict one of three possible states for the temperature in a future period (up to one year): above-normal, near-normal or below-normal. This forecasted scenario determines the parameters of the temperature modelled by an Ornstein-Uhlenbeck process. For pricing in spot markets, Benth & Meyer-Brandis [2] enlarge the filtration by future information to obtain a filtration representing all available information. They introduce new parameters called “information yield” and “information premium” and model these parameters for electricity markets. Benth & Meyer-Brandis [2] argue that “significant parts of the supposedly irregular market price of risk in electricity markets is in reality due to information misspecification in the model” [2, pg. 1]. Through index modelling, Dorfleitner & Wimmer [11] calculate the index outcome depending on the years of historical data used, with and without detrending. They add meteorological forecasts and compare the results with the CME prices of monthly and seasonal contracts for US cities between 2002 and 2006. Their approach, however, works only if the forecasts reach into the accumulation period. They conclude that meteorological weather forecasts have to be included into the pricing process.

In this paper, we will introduce a daily modelling approach which also regards meteorological temperature forecasts. Flexibility is the key advantage of daily modelling; any temperature index can be easily derived, when a temperature model is determined. Our method also allows the pricing to include meteorological forecasts, even if they do not reach into the accumulation period. Therefore, our model is the first model that makes general pricing with meteorological forecasts possible for every temperature contract and every trading time. Furthermore, we apply

this method to price temperature futures for New York City, Minneapolis and Cincinnati as well as compare the prices simulated with and without meteorological forecasts with CME market prices to determine the influence of meteorological forecasts on the pricing of weather derivatives. The results show that prices which include meteorological forecasts reflect market prices much better than prices without the use of forecasts. As one would expect, the impact increases with increasing proximity to the accumulation period.

The paper is structured as follows; in Section 2, we propose a time-discrete autoregressive model with seasonal variance extended by meteorological forecast data to derive prices of weather derivatives. The temperature data, the meteorological forecast data and the market price data for weather derivatives in New York City, Minneapolis and Cincinnati are also presented in this section. In Section 3, we compare market prices with theoretical prices with and without using meteorological forecast data and quantify the differences. In Section 4, we conclude with some further discussion.

2. Methods and data

2.1. Pricing of weather derivatives

Financial theory asserts that the arbitrage free price $F_{(t;\tau_1,\tau_2)}$ at time $t \in \mathbb{N}$ of a derivative with contract period $[\tau_1, \tau_2]$ and weather-dependent payoff $Y_M(W_M)$ at time $M > t$ is calculated as the discounted, conditional expected payoff based on the filtration \mathcal{F}_t which includes all the information available at time t [e.g. 12]:

$$F_{(t;\tau_1,\tau_2)} = \frac{1}{(1+r)^{M-t}} \mathbb{E}_{\mathbb{Q}} [Y_M(W_M) | \mathcal{F}_t], \quad (1)$$

where W_M is the value of a certain weather index at time M and $r \geq 0$ is the interest rate. While the risk neutral measure \mathbb{Q} can be uniquely derived for derivatives written on tradeable assets, this is not the case for weather derivatives because weather is not tradeable. Thus, the market for weather derivatives is incomplete. Consequently, the no-arbitrage condition no longer implies a clear-cut risk neutral measure.

Two directions have been proposed in the literature to overcome the pricing problem for incomplete markets. First, one could arbitrarily assume a particular risk utility function for the market participants. This idea is pursued in the context of marginal pricing [10], indifference

pricing [17] or equilibrium pricing [8, 9]. Second, some researchers have attempted to determine the unknown risk preference and thus the market price of weather risk implicitly from observed market price quotations. For example, Alaton et al. [1] compare two contracts for Stockholm and infer that the market price of risk is not constant. To derive their pricing formulas, Benth & Šaltytė-Benth [4] use a time-varying function θ_t which is not specified. Härdle & López Cabrera [14] analyse data of futures for Berlin to test different shapes of the market price of risk. They conclude that it is a deterministic, time dependent function with a seasonal structure.

Benth & Meyer-Brandis [2], however, conjecture that the filtration \mathcal{F}_t causes the irregular pattern of the market price of risk. For the pricing of temperature derivatives, the filtration \mathcal{F}_t usually contains only the historical temperature values (up to day $t-1$), yet there is more information available on the market. In particular, meteorological temperature forecast models exploit additional information such as air pressure or wind speed, and process this information within complex physical models. Because forecasts from these models are available to market participants, they should be used and included in temperature models, which form the basis of theoretical pricing approaches.

From a statistical perspective, several alternatives exist for modelling the stochastic process W in Equation (1) [15]. One approach is to estimate the distribution of the underlying weather index (e.g. accumulated temperatures or rainfall) at maturity M directly, either parametrically or non-parametrically. The estimated distribution is then used for simulating the derivative's payoff $Y_M(W_M)$. This procedure is called “index value simulation”.

Alternatively, one can estimate a stochastic process for the generic weather variable (e.g. temperature) and derive any index of interest. The latter approach, which is called “daily simulation”, is more flexible and statistically more reliable compared to the index value simulation. In the next subsection, a daily temperature model is presented.

2.2. Temperature model

The underlying weather indices of the derivatives traded at the CME for US cities are directly derived from the daily average temperature T_t , which is defined as the mean of the minimum and the maximum temperature at a weather station on a day t , $t \in \mathbb{N}$. Our model is based on a time-discrete model for the daily average temperature and is similar to the model which Benth et al. [5] use for Stockholm temperature data and which Härdle & López Cabrera [14] use for

Berlin temperature data. The daily average temperature T_t on day t , $t \in \mathbb{N}$, is given by:

$$\begin{aligned}
 T_t = & \underbrace{a + bt}_{\text{trend}} + \underbrace{\sum_{p=1}^P \left[a_p \cos\left(\frac{2\pi pt}{365}\right) + b_p \sin\left(\frac{2\pi pt}{365}\right) \right]}_{\text{seasonality}} \\
 & + \underbrace{\sum_{l=1}^L \rho_{t-l} T_{t-l}}_{\text{autoregression}} + \underbrace{\sigma_t \epsilon_t}_{\text{stochastic}} \quad , \\
 \sigma_t^2 = & \sum_{q=1}^Q \left[c_q \cos\left(\frac{2\pi qt}{365}\right) + d_q \sin\left(\frac{2\pi qt}{365}\right) \right] , \epsilon_t \sim \mathcal{N}(0, 1),
 \end{aligned} \tag{2}$$

with $L, P, Q \in \mathbb{N}$ and $a, a_1, \dots, a_P, b, b_1, \dots, b_P, \rho_{t-L}, \dots, \rho_{t-1}, c_1, \dots, c_Q, d_1, \dots, d_Q \in \mathbb{R}$.

It consists of the following components:

- A linear trend captures the long-term development of the temperature.
- On average, the temperature in the summer is higher than in the winter. Thus, the temperature contains a seasonal component which is described by a truncated Fourier series of order P .
- The temperature of one day depends on the temperature on the days before. This autoregression is captured by an AR-process of lag L .
- Moreover, the temperature is random, so a stochastic component is needed. Empirical data shows that the variance has a seasonal component and is, on average, higher in the winter than in the summer. We model the stochastic part by a seasonal variance and a standard normally distributed random variable.

The choice of the temperature model is a crucial decision, which is intensively discussed in the literature. As a new approach, Dupuis [13] proposes to model the minimum and maximum temperatures separately with extreme value theory. For modelling the daily average temperature, Campbell & Diebold [7] suggest a GARCH-process as an alternative for the variance to approximate the cyclical volatility. Jewson & Brix [15] suggest time series models such as ARMA, ARFIMA and (S)AROMA models. For continuous time, models such as an Ornstein-Uhlenbeck process [4] are widely used and are sometimes driven by a fractional Brownian motion [6] or a Lévy process [3]. A widely accepted model is the CAR (continuous autoregressive) model [5].

This model is the continuous-time counterpart to the model we use. Because of the continuity, this model is helpful for the theoretical analysis of pricing and hedging decisions. We chose the AR model because it captures the main properties of the temperature (trend, seasonality, autoregression and seasonal variance) and is not too difficult to estimate; however, depending on the observed city, other models might outperform the AR model. Nevertheless, the AR model is widely accepted as a sufficient compromise between manageability and accuracy.

The two indices used in this study and traded at the CME for US cities, heating degree days and cooling degree days, are derived from the daily average temperature T_t as follows: the (cumulative) heating degree days (HDD) over a period $[\tau_1, \tau_2]$, $\tau_1, \tau_2 \in \mathbb{N}$, $\tau_1 \leq \tau_2$, with threshold K (usually $18^\circ\text{C}/65^\circ\text{F}$) are defined as the sum of the daily heating degree days in the period, i. e.

$$\text{HDD}(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} \text{HDD}_t = \sum_{t=\tau_1}^{\tau_2} \max(0, K - T_t). \quad (3)$$

The (cumulative) cooling degree days (CDD) over a period $[\tau_1, \tau_2]$, $\tau_1, \tau_2 \in \mathbb{N}$, $\tau_1 \leq \tau_2$, with threshold K (usually $18^\circ\text{C}/65^\circ\text{F}$) are defined as the sum of the daily cooling degree days in the period, i. e.

$$\text{CDD}(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} \text{CDD}_t = \sum_{t=\tau_1}^{\tau_2} \max(0, T_t - K). \quad (4)$$

For most of the contracts, the contract period $[\tau_1, \tau_2]$ is one calendar month. After this accumulation period, the index outcome W_M is known and the “tick size” (20 \$ for US cities) converts the outcome into a monetary amount, which is the payoff of the contract $Y_M(W_M)$.

2.3. Extended model

In this subsection, we explain how we include the forecast data into the temperature simulation and how we simulate the prices of weather derivatives.

In usual pricing approaches for weather derivatives, the filtration \mathcal{F}_t consists only of the historical temperatures. We will later refer to the model based on this filtration by “NMF” for “no meteorological forecast”. Benth & Meyer-Brandis [2] propose to enlarge this filtration and to use the filtration \mathcal{G}_t , $\mathcal{F}_t \subset \mathcal{G}_t$, which contains all available information that is relevant for the future

development of the underlying. This filtration \mathcal{G}_t , however, is rather theoretical as it is difficult in practise to incorporate all of the available information into the model in practice.

For our extended model, we enlarge the filtration \mathcal{F}_t by adding a limited number of meteorological forecast data. Instead of using only the historical temperatures T_0, \dots, T_{t-1} available at time t , we additionally use meteorological temperature forecasts k days in advance, calculated on day t . We call this enlarged filtration $\mathcal{G}_t^{\text{MF}k}$ with $k \in \mathbb{N}_0$ describing the number of days in advance that meteorological forecasts are considered. For example, $\mathcal{G}_t^{\text{MF}0}$ means that we add meteorological forecast data for the same day ($t+0$) and $\mathcal{G}_t^{\text{MF}13}$ means that we use meteorological forecast data up to 13 days in advance ($t+0, t+1, \dots, t+13$). As a result of limited data, the maximum for k is 13 in this study, yet this number could in general be greater than 13. In the following, we use the notation:

$$\mathcal{F}_t \subset \mathcal{G}_t^{\text{MF}0} \subset \mathcal{G}_t^{\text{MF}1} \subset \mathcal{G}_t^{\text{MF}2} \subset \dots \subset \mathcal{G}_t^{\text{MF}13} \subset \mathcal{G}_t. \quad (5)$$

Apart from the historical temperature data, the weather service uses additional information such as atmospheric processes to forecast temperatures. We assume that the meteorological forecast data up to 13 days in advance is the best information we have about the future temperature evolution – even better than forecasts made by our statistical temperature model. Thus, the meteorological forecasts are added to the historical temperature data as if they were actually observed values.

If we want to simulate the price of a temperature futures with accumulation period $[\tau_1, \tau_2]$ on a day t (see Fig. 1), we know the historical daily average temperatures from a starting day (t_0) to one day before ($t-1$). Moreover, we have meteorological temperature forecast data which are calculated on day t for the same day (t), the next day ($t+1$) and up to k days in advance. Based on this new, extended time series with data from t_0 to $t+k$, we adjust the time-discrete AR-model (Subsection 2.2) to the data for day t . The orders of the two truncated Fourier series (P, Q) and the lag of the autoregression (L) are set beforehand, as they do not change every day. The other parameters, however, are estimated newly for every day t depending on the temperature and forecast data available on that day.

With this model, we can simulate the temperatures for every day after $t+k$, especially for $[\tau_1, \tau_2]$ which is the accumulation period of the contract, and calculate the resulting payoff of the

HDD (Eq. (3)) or CDD contract (Eq. (4)). By using Monte-Carlo simulation, this procedure is repeated 10 000 times for every day t . The mean of these 10 000 simulated prices represents the expected payoff $Y_M(W_M)$. The theoretical prices of the contract are set equal to the expected payoff under the real-world probability measure \mathbb{P} , depending on the filtration:

$$\hat{F}_{(t;\tau_1,\tau_2)}^{\text{NMF}} = \mathbb{E}_{\mathbb{P}} [Y_M(W_M)|\mathcal{F}_t], \quad (6)$$

$$\hat{F}_{(t;\tau_1,\tau_2)}^{\text{MF}k} = \mathbb{E}_{\mathbb{P}} [Y_M(W_M)|\mathcal{G}_t^{\text{MF}k}]. \quad (7)$$

Compared to Equation (1), we use the simplifying assumptions that the market price of risk θ and the interest rate r are 0, so $\mathbb{Q} = \mathbb{P}$. One reason is that the interest level is quite low and we deal with short-term pricing. Moreover, we aim to compare the differences between market prices and theoretical prices with and without meteorological forecasts and not to describe market prices perfectly. Thus, further assumptions about the market price of risk or the interest rate are not required, as they would change the absolute values of the theoretical prices only and not the differences.

Based on the number of meteorological forecasts we use, we obtain prices simulated at time t for the models NMF (no meteorological forecasts, Eq. (6)), MF0 (meteorological forecast for the same day), \dots , and MF13 (meteorological forecasts 13 days in advance, Eq. (7)). We repeat this procedure for every day t in the trading period and obtain theoretical prices for the futures based on the different models.

Thus far, this procedure considered the general case only in which neither the historical data nor the meteorological forecast data overlap with the accumulation period of the contract, so that all the temperatures in this period have to be simulated. If these data last into the accumulation period, however, we use these values directly and simulate temperatures only for the rest of the period. As a result, no simulation is required, if we are close to the end of the accumulation period where historical data and meteorological forecast data cover the whole period.

There are alternative approaches for incorporating the forecast data. Jewson & Caballero [16] suggest evaluating past meteorological forecast data to obtain additional information. Based on past meteorological forecast data, one could derive a conditional mean forecast and, furthermore, compare it with reality to find systematic errors.




Figure 1: Model for payoff simulation on day t

There are two reasons why we do not pursue this approach here. On the one hand, we concentrate on the simulation with information accessible on one day; usually, historical temperatures and meteorological forecasts calculated on the same day for a few days in advance are available. It is unlikely that market participants have an archive of historical meteorological weather forecasts, which they can evaluate. On the other hand, we presume that the weather service itself evaluates its meteorological forecasts. If the forecasts were biased, the weather service would probably adjust them.

Jewson & Caballero [16] also suggest the pruning method to increase weight of the meteorological forecasts. There, two probability densities are calculated. The first is based on meteorological forecasts and the second is based on historical temperatures. By weighting these densities, paths for future temperatures can be simulated. The problem with this method is that calculating a density based on 14 meteorological forecast values is not useful. Thus, this method would again lead to an evaluation of the historical meteorological forecast data.

After having simulated theoretical prices with different models depending on the number of meteorological forecasts used, we compare them with the daily observed market prices at the CME. Thereby, we will examine if theoretical prices with meteorological forecasts predict market prices better than theoretical prices which do not use meteorological forecasts. The goodness of

the prediction is calculated through the root mean squared prediction error (RMSPE) defined as

$$\text{RMSPE}(\hat{F}, F) = \sqrt{\frac{1}{N} \cdot \sum_{i=1}^N (\hat{F}_{t_i} - F_{t_i})^2}. \quad (8)$$

This RMSPE measures the difference between the actual observed market price F and the theoretical price \hat{F} which is calculated by different models without meteorological forecasts (NMF) and with meteorological forecasts (MF0–MF13) as it accumulates the quadratic difference between the prices for every day in the considered period $[t_1, t_N]$.

In addition to the RMSPE, we display the correlation coefficient. If meteorological weather forecasts influence the pricing of the derivatives, this correlation should increase when the meteorological forecasts are included in the pricing. Later, we will calculate the correlation for every model and contract as well as have a closer look at the means of the correlations for different cities which is obtained via Fisher transformation.

We will apply these methods in Section 3 to compare the results based on different models.

2.4. Data

In this subsection, we present the data of the historical temperatures, the meteorological forecasts and the market prices used in this study to examine the differences between the prices with and without the use of meteorological forecasts.

For historical temperature data, we use temperature data from New York City – La Guardia Airport, Minneapolis – St. Paul International Airport and Cincinnati – Northern Kentucky Airport from 01/01/1997 to 28/02/2010 provided by the CME. To obtain years of equal length, leap days were removed from the data.

Furthermore, we have point forecast data from WeatherOnline¹ from 29/12/2008 to 12/02/2010 with a few days missing². The dataset consists of forecasted minimum and maximum temperatures from 0 to 13 days in advance, i. e. 14 days, for New York City, Minneapolis and Cincinnati.

¹The authors cordially thank H. Werner and U. Römer from WeatherOnline for providing meteorological forecast data.

²No forecast data are available for 12/01/2009, 21/01/2009, 14/08/2009, 29/08/2009, 19/09/2009, 08/10/2009, 18/10/2009, 23/11/2009, 30/11/2009. Incomplete forecast data are available for 07/01/2009, 08/01/2009, 25/01/2009 (forecast 13 days in advance are missing), 26/01/2009 (forecasts 12 and 13 days in advance are missing).

We compare our theoretical prices with the market prices reported at the CME as “last price” for every weekday³. We consider monthly futures contracts for New York, Minneapolis and Cincinnati from February, 2009 to January, 2010 as they overlap with the period of the meteorological forecast data. We chose these cities due to their high total trading volume in the period, New York (42.253) and Minneapolis (11.190), compared to other cities in our dataset. The volume in Cincinnati is lower (6.782), but this volume is distributed over a higher number of trading days. A detailed description of the contracts, the number of trading days, the traded volume, the days with a volume larger than zero⁴ and the payoffs are listed in Table A.1 (New York), Table A.2 (Minneapolis) and Table A.3 (Cincinnati).

3. Results

In this section, we first present the results for the temperature model (Subsection 3.1). Subsection 3.2 contains a short analysis of the forecast data. The main part of this section consists of the presentation of the HDD and CDD futures prices and the comparison with market prices in Subsection 3.3. Finally, we will check our results for hypothetically perfect meteorological forecasts (Subsection 3.4).

3.1. Historical temperatures

At first, we study the historical temperature data until 31/12/2008 (12 complete years) and test the temperature properties mentioned in Subsection 2.2 for our temperature data. Moreover, we determine the orders of the two truncated Fourier series (P, Q) and the lag of the autoregression (L). Later, as mentioned above, these numbers are fixed, opposite to the other parameters of the model, which are estimated every day.

According to the AIC and BIC, the order of the Fourier series of the seasonality term P is set to equal three for all three cities. After removing trend and seasonality, we check the residuals for autoregression. The autocorrelation function and the partial correlation function of the residuals show an autoregressive pattern and indicate that the lag of the autoregression is $L=3$ for all

³We obtained the data and daily trading volume from Bloomberg via the Research Data Center (RDC) of the Collaborative Research Center (CRC) 649 “Economic Risk”.

⁴The question arises how a market price for a day is reported if the traded volume on that day is zero. According to the guidelines of the CME, bid/ask information, trade information and market information from third party sources are considered when obtaining prices for every trading day. (CME Group: Weather Futures and Options Settlement Guidelines, www.cmegroup.com/trading/weather/files/Weather_Futures_and_Options_Settlement_Guidelines.pdf)




Figure 2: Quadratic difference of the meteorological forecasts and the temperature model compared to the observed temperature in New York in 2009

three cities. The squared residuals without trend, seasonality and autoregression still show a seasonality in the autocorrelation function, indicating seasonality in the variance. In fact, a higher variance in the winter than in the summer is observed in most studies about temperature models [e.g. 1, 5, 14]. According to the AIC and BIC, the order of the Fourier series of the seasonality in the variance is set to $Q=1$ for all cities. The histograms of the residuals without trend, seasonality, autoregression and seasonal variance indicate a standard normal distribution of these residuals. The Kolmogorov-Smirnoff test does not reject the normality hypothesis at the 5% significance level.

3.2. Meteorological forecasts

Using meteorological forecast data in our model is appropriate and reasonable as these data are better than data simulated by the statistical temperature model. A comparison of the quadratic deviation of meteorological forecasts from observed temperatures and the quadratic deviation of temperatures simulated by the statistical model in New York in 2009 shows that meteorological forecasts outperform the statistical model (Figure 2). This finding confirms our motivation for using meteorological forecast data; however, the quadratic deviation declines with increasing forecast periods. For more than 10 days, the statistical model even appears to be more accurate than meteorological forecasts, at least in 2009.




Figure 3: Observed and simulated prices of an HDD contract for January, 2010 in New York

3.3. Futures prices and information premium

Applying the methods described in the previous section, we obtain theoretical prices for the different models and contracts. As an example, the results for the HDD contract in January, 2010 with a reference station in New York are depicted in Figure 3, measured in index points. The CME price as well as the theoretical prices with and without meteorological forecasts are constant for a rather long trading period. No significant difference between the NMF and the MF13 model can be observed in the first months. Looking at the last two months, however, reveals stronger fluctuations of the prices and differences between the models. For this HDD contract, the MF13 simulated price is closer to the CME price than the one simulated without meteorological forecasts.

Benth & Meyer-Brandis [2] introduce the term “information premium” and define it as the difference between the theoretical prices calculated with and without using additional information such as weather forecasts. This premium describes how theoretical prices change over time if meteorological forecasts are considered. Figure 4 shows the information premium for the same HDD contract as in Figure 3, as well as the mean of the information premium for all of the twelve contracts for New York, plotted against the time to maturity. Obviously, the information premium is close to zero until about 60 days before the contract expires. Thereafter, the information premium has a considerable size, either positive or negative. For an HDD contract, it is negative (positive) if the meteorological forecasts are higher (lower) than the forecasts from the statistical model. The opposite is true for a CDD contract.




Figure 4: Information premium for the HDD contract for January, 2010 in New York and its mean for all the twelve New York contracts

For each city, the RMSPE is calculated for every contract and model separately. The means of the RMSPE for the different models for New York, Minneapolis and Cincinnati are depicted in Table 1. The means show that the RMSPE decreases from 24.1 to about 20 (up to 18%) for New York, from 32.4 to about 30 (up to 9%) for Minneapolis, and from 34.7 to about 30 (around 13%) for Cincinnati if meteorological forecast data are used. The RMSPE for the twelve observed contracts with a reference station New York is reported in Table A.4. It decreases for every contract for some models which include meteorological forecasts.

The absolute value of the RMSPE does not have an important role because we compare the improvements caused by the use of the meteorological forecast values only. The error for the New York CDD contract in April, for example, is lower than the error for other months. However, this does not imply that the prediction for April is better. The theoretical prices for the April CDD contract (< 50) and its outcome (32.5) are much lower than the ones for the other contracts in New York (see Table A.1 for the outcomes), which causes the lower error. Note that we do not perfectly predict prices and that in most cases an extra premium such as the market price of risk should be taken into account, which would further reduce the absolute value of the RMSPE. The constant gap between the simulated and the actual prices in the first months observable in the left panel of Figure 3 likely reflects this aspect.

RMSPE	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
New York	24.1	23.8	23.4	23.1	22.3	22.1	21.6	21.6	20.9	20.7	20.3	20.0	19.8	19.9	20.0
Minneapolis	32.4	31.5	30.9	30.3	30.0	29.7	29.6	29.5	29.6	29.6	29.6	29.5	29.6	29.8	30.3
Cincinnati	34.7	34.3	33.5	32.8	32.2	32.1	31.8	31.6	31.2	30.8	30.6	30.3	30.2	30.0	30.5

Table 1: Mean of RMSPE for different contracts and models (whole trading period)

RMSPE	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
2 months															
New York	30.8	30.0	28.9	27.7	26.1	25.2	24.2	23.7	22.5	22.0	21.6	21.1	20.9	21.3	21.4
Minneapolis	42.8	40.5	38.9	37.5	36.4	35.3	33.6	32.9	32.6	32.3	32.6	32.7	33.0	33.3	34.1
Cincinnati	44.9	43.9	42.0	40.3	38.7	37.9	36.5	35.3	34.0	32.6	31.8	31.2	31.0	30.3	30.6

Table 2: Mean of the RMSPE for different contracts and models (last 2 months only)

Figure 4 reveals that the information premium is significantly different from zero in the last 60 days until maturity, resulting in larger differences between the models in this period. Thus, we calculate the RMSPE for the last two months of the trading period (the accumulation period and the month before) only. Compared to the RMSPE for the whole period, the means for the three cities (see Table 2) show a higher decrease (between 25% and 33%) if meteorological forecast data is used. The results for all the contracts and all the models for New York for the last two months are reported in detail in Table A.5.

As mentioned above, the CME calculates market prices even if there is no trading on that day. One could conjecture that only prices which are the result of active trading capture all relevant information. To check this hypothesis, we confine the calculation of the RMSPE to those days on which the trading volume is larger than zero. For the means, we obtain an improvement of up to 35% for New York, up to 40% for Minneapolis and up to 54% for Cincinnati (see Table 3). Nevertheless, these results should be interpreted with care because the calculation of the RMSPE is for some contracts based on only one or two days (see Tables A.1-A.3). The detailed results for New York for every model and contract are reported in Table A.6.

All the results for the RMSPE based on the whole period, the last two months and the days with a volume larger than zero, as well as for the correlation are shown in Figure 5 (New

RMSPE	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
Vol>0															
New York	36.6	35.8	35.4	33.2	30.6	29.1	27.6	27.2	25.3	24.8	24.1	23.2	22.9	23.8	23.8
Minneapolis	54.3	48.8	45.0	43.8	40.4	36.3	32.5	31.9	32.8	32.3	34.1	32.8	32.8	34.2	35.1
Cincinnati	58.0	55.1	49.1	45.7	47.8	45.5	44.2	40.9	35.8	31.1	29.5	28.6	28.7	27.2	26.8

Table 3: Mean of the RMSPE for different contracts and models (for days with trading volume>0)




Figure 5: Average nRMSPE (whole period, last 2 months, days with volume>0) and the correlation coefficient for New York

York), Figure 6 (Minneapolis) and Figure 7 (Cincinnati) where we compare the means for the normalized⁵ RMSPE (nRMSPE) and the correlation coefficients for the different pricing models. All RMSPE graphs decrease and the correlations increase if meteorological data is included. It is noticeable that many graphs seem to turn at the end, indicating that there is an optimal number of meteorological forecasts that should be used. This finding is in line with the aforementioned result that the meteorological forecasts many days in advance in 2009 were worse than the forecasts from the statistical model.

The fact that meteorological weather forecasts influence the pricing of weather derivatives is also confirmed by calculations of the correlation coefficient. The mean of the correlation coefficients of the three cities is presented in Table 4. Compared to the model without forecasts, the correlation increases for most of the contracts when all of the forecast data is used. The mean of the correlations rises from 0.62 to 0.78 (New York), from 0.65 to up to 0.75 (Minneapolis), and from 0.64 to 0.80 (Cincinnati). More detailed results for every contract and model for New York are reported in Table A.7.

⁵“Normalized” here means that we set the RMSPE of the model without meteorological forecasts equal to 1 for every calculation method. This helps to compare the improvements caused by meteorological forecasts for the RMSPE which are calculated by different methods.




Figure 6: Average nRMSPE (whole period, last 2 months, days with volume>0) and the correlation coefficient for Minneapolis




Figure 7: Average nRMSPE (whole period, last 2 months, days with volume>0) and the correlation coefficient for Cincinnati

Correlation	Model															
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13	
New York	0.62	0.63	0.64	0.65	0.68	0.68	0.69	0.67	0.71	0.73	0.74	0.76	0.77	0.77	0.78	
Minneapolis	0.65	0.67	0.68	0.69	0.70	0.72	0.74	0.75	0.75	0.74	0.73	0.73	0.72	0.72	0.70	
Cincinnati	0.64	0.65	0.69	0.71	0.73	0.73	0.75	0.76	0.77	0.79	0.80	0.81	0.80	0.81	0.80	

Table 4: Mean of the correlation coefficients for different contracts and models

RMSPE	Model		
	NMF	MF13	MF13perfect
New York	24.1	20.0	21.7
Minneapolis	32.4	30.3	34.9
Cincinnati	34.7	30.5	31.6

Table 5: Mean RMSPE for different contracts and models

3.4. Perfect forecasts

Clearly, a variety of meteorological forecast models exist and one might object that the results presented in this study are biased because market participants at the CME did not use the same forecast as we did. Perhaps the price is more influenced by other meteorological forecast data and better results were obtained with different forecast data.

Unfortunately, we do not have access to other meteorological forecast data to check this hypothesis. Nevertheless, from an ex-post perspective one can generate better and even perfect forecasts by using actual temperatures. Instead of adding meteorological forecast data to the filtration, we add the actual temperatures 13 days in advance as this is what the best weather service could forecast for the next days. We call this filtration $\mathcal{G}_t^{\text{MF13perfect}}$. Note that $\mathcal{G}_t^{\text{MF13perfect}} = \mathcal{F}_{t+14}$.

The means of the RMSPE for the model including perfect forecasts 13 days in advance for New York, Minneapolis and Cincinnati are presented in Table 5. More detailed results for New York can be found in Table A.8. It turns out that the model with actual (imperfect) temperature forecasts is superior to the model using perfect temperature forecasts for all three cities. This result is not surprising because market participants have access to imperfect meteorological forecasts only and no other information can be incorporated in the observed futures prices. Moreover, we conclude that the results would not significantly change if market participants used meteorological weather forecast different from ours.

4. Discussion and conclusion

In this article we compared the outcome of different pricing models for temperature futures at the CME. The pricing models differ in the information set that is used to predict the expected weather index at maturity. Standard pricing models use a time series approach, i.e. predictions of the underlying weather variable are simply based on its own history. Thereby, other information which is available at the present time is ignored. Nevertheless, such information is incorporated

into meteorological forecasts and it is likely that market participants utilize this information when trading weather derivatives. Our results demonstrate that the inclusion of meteorological forecasts has a clear impact on the price predicted by theoretical models. Theoretical prices which incorporate meteorological forecast are much closer to actual futures prices observed at the CME compared with a benchmark model without such forecasts. Admittedly, this effect is significant for the last two months before expiration only because reliable forecasts are only available in the short term. We also found that the value of meteorological forecasts declines with a longer forecast horizon.

The purpose of our modelling effort was not to outperform the market, but to explain observed market prices. Nevertheless, our findings are also relevant from an ex-ante perspective. As mentioned above, many researchers have attempted to estimate the market price of weather risk implicitly from observed price quotations and have applied this parameter to obtain arbitrage-free prices for other locations or for any other type of index. When measuring the market price of risk as a residual component which relates theoretical prices to observed prices, it is important that theoretical prices capture all available information. Otherwise, the estimation of the market price of risk will be biased and this error will translate into errors in the price of the weather derivative. Against this background, many empirical results on the size and the structure of the market price of risk appear questionable. Thus, we recommend an inclusion of meteorological forecasts when analyzing the market price of weather risk empirically.

Our results provide the first evidence for the value of meteorological forecasts when pricing weather derivatives. However, a generalization of our results requires a broader empirical study which includes more cities and contracts as well as other weather indices. Moreover, the incorporation of probabilistic instead of deterministic forecasts would be interesting, where probabilities for future temperatures instead of single values are reported by the weather service.

References

- [1] Alaton, P., Djehiche, B., & Stillberger, D. (2002). On modelling and pricing weather derivatives. *Applied Mathematical Finance*, 9(1), 1–20.
- [2] Benth, F. E. & Meyer-Brandis, T. (2009). The information premium for non-storable commodities. *Journal of Energy Markets*, 2(3), 111–140.

- [3] Benth, F. E. & Šaltytė-Benth, J. (2005). Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance*, 12(1), 53–85.
- [4] Benth, F. E. & Šaltytė-Benth, J. (2007). The volatility of temperature and pricing of weather derivatives. *Quantitative Finance*, 7(5), 553–561.
- [5] Benth, F. E., Šaltytė-Benth, J., & Koekebakker, S. (2007). Putting a price on temperature. *Scandinavian Journal of Statistics*, 34(4), 746–767.
- [6] Brody, D. C., Syroka, J., & Zervos, M. (2002). Dynamical pricing of weather derivatives. *Quantitative Finance*, 2, 189–198.
- [7] Campbell, S. D. & Diebold, F. X. (2005). Weather forecasting for weather derivatives. *Journal of the American Statistical Association*, 100, 6–16.
- [8] Cao, M. & Wei, J. (2004). Weather derivatives valuation and market price of weather risk. *Journal of Futures Market*, 24(11), 1065–1089.
- [9] Chaumont, S., Imkeller, P., Müller, M., & Horst, U. (2005). A simple model for trading climate risk. *Vierteljahrshefte zur Wirtschaftsforschung / Quarterly Journal of Economic Research*, 74(2), 175–195.
- [10] Davis, M. (2001). Pricing weather derivatives by marginal value. *Quantitative Finance*, 1, 305–308.
- [11] Dorfleitner, G. & Wimmer, M. (2010). The pricing of temperature futures at the Chicago Mercantile Exchange. *Journal of Banking & Finance*, 34(6), 1360–1370.
- [12] Duffie, D. (2001). *Dynamic Asset Pricing Theory*. Princeton University Press, 3rd edn.
- [13] Dupuis, D. J. (2010). Forecasting temperature to price CME temperature derivatives. *International Journal of Forecasting*.
- [14] Härdle, W. & López Cabrera, B. (2009). Implied market price of weather risk. *SFB 649 Discussion Paper 2009-001*.
- [15] Jewson, S. & Brix, A. (2005). *Weather Derivative Valuation: The Meteorological, Statistical, Financial and Mathematical Foundations*. Cambridge University Press.

- [16] Jewson, S. & Caballero, R. (2003). The use of weather forecasts in the pricing of weather derivatives. *Meteorological Applications*, 10(04), 377–389.
- [17] Musiela, M. & Zariphopoulou, T. (2004). A valuation algorithm for indifference prices in incomplete markets. *Finance and Stochastics*, 8(3), 399–414.
- [18] Yoo, S. (2004). Weather derivatives and seasonal forecast. *Asia-Pacific Journal of Financial Studies*.

A. Appendix

New York		Trading days	Traded volume	Days with vol>0	Payoff
Feb09	HDD	38/247	4225/4809	22/26	787.0
Mar09	HDD	61/217	5496/6216	26/30	718.5
Apr09	CDD	82/143	0	0	32.5
May09	CDD	102/249	4550	12	48.5
Jun09	CDD	124/185	3635	23	126.5
Jul09	CDD	145/206	2100	20	293.5
Aug09	CDD	166/228	5878	20	394.5
Sep09	CDD	187/249	4575	15	114.5
Oct09	HDD	66/68	1895	14	262.0
Nov09	HDD	172/177	2555	14	392.5
Dec09	HDD	185/189	2646/3211	21/23	872.5
Jan10	HDD	205/209	2829	14	992.5

Table A.1: Futures contracts for New York used in this study overlapping with the period of the meteorological forecast data; the number of trading days, the traded volume, the number of days with volume>0 and the payoffs are shown. If two numbers are depicted, this indicates that less data than available were used because of missing meteorological forecast data.

Minneapolis		Trading days	Traded volume	Days with vol>0	Payoff
Feb09	HDD	38/247	668/752	8/12	1238.5
Mar09	HDD	61/217	547/651	7/10	1017.0
Apr09	CDD	82/143	0	0	1.5
May09	CDD	102/249	100	1	38.0
Jun09	CDD	124/185	1400	5	153.5
Jul09	CDD	145/206	950	2	164.5
Aug09	CDD	166/228	2400	12	158.0
Sep09	CDD	187/249	1500	2	104.0
Oct09	HDD	64/66	150	2	676.5
Nov09	HDD	163/167	825	7	669.0
Dec09	HDD	184/188	1958/2262	15/17	1479.5
Jan10	HDD	201/206	200	4	1610.5

Table A.2: Futures contracts for Minneapolis used in this study overlapping with the period of the meteorological forecast data; the number of trading days, the traded volume, the number of days with volume>0 and the payoffs are shown. If two numbers are depicted, this indicates that less data than available were used because of missing meteorological forecast data.

Cincinnati		Trading days	Traded volume	Days with vol>0	Payoff
Feb09	HDD	38/247	272/440	7/10	819.0
Mar09	HDD	61/217	931/1139	9/12	543.0
Apr09	CDD	82/143	0	0	36.5
May09	CDD	102/249	1450	6	38.0
Jun09	CDD	124/185	425	6	240.5
Jul09	CDD	145/206	515	7	164.5
Aug09	CDD	166/228	225/325	2/3	225.0
Sep09	CDD	187/249	470	4	117.0
Oct09	HDD	64/66	350	6	408.0
Nov09	HDD	163/167	375	5	526.0
Dec09	HDD	184/188	725/875	7/8	983.5
Jan10	HDD	204/208	250	3	1183.0

Table A.3: Futures contracts for Cincinnati used in this study overlapping with the period of the meteorological forecast data; the number of trading days, the traded volume, the number of days with volume>0 and the payoffs are shown. If two numbers are depicted, this indicates that less data than available were used because of missing meteorological forecast data.

RMSPE	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
Feb09	28.8	28.2	29.4	30.5	28.1	27.3	26.1	28.9	28.0	27.1	26.1	24.0	20.3	17.9	18.8
Mar09	22.9	22.6	22.3	21.8	19.8	20.5	20.1	21.0	20.2	19.7	20.9	20.6	22.3	25.0	25.2
Apr09	6.8	6.8	6.7	6.4	6.2	6.4	6.7	6.7	6.8	7.0	7.2	7.3	7.5	7.8	7.7
May09	21.2	21.3	21.2	21.0	21.0	20.7	20.5	20.5	20.6	20.3	20.4	20.5	20.5	20.7	20.8
Jun09	17.2	16.8	15.9	15.4	15.0	14.4	14.2	14.6	14.8	15.5	16.4	17.3	18.4	19.2	19.9
Jul09	27.0	26.0	25.1	24.0	23.1	22.4	21.4	21.1	20.2	20.2	19.7	19.5	18.8	18.8	18.7
Aug09	20.2	20.6	20.5	20.6	20.3	20.0	19.9	19.6	18.2	17.5	16.7	16.1	16.0	16.2	15.5
Sep09	18.5	17.8	17.5	17.3	16.9	16.7	16.6	16.6	16.6	16.6	16.2	16.3	16.2	16.3	16.4
Oct09	24.1	23.3	22.4	21.9	21.3	21.6	19.5	17.4	14.9	13.5	9.9	9.6	8.3	7.6	8.0
Nov09	26.8	26.8	26.7	26.4	25.9	25.5	25.3	24.9	24.8	25.0	25.6	25.7	26.2	26.7	27.1
Dec09	38.3	37.6	36.4	35.3	34.4	34.3	34.1	33.1	32.5	32.3	32.1	31.7	31.7	31.8	31.7
Jan10	38.0	37.7	37.4	36.8	35.9	35.2	34.8	34.4	33.6	33.1	32.2	31.7	31.3	30.8	30.5
Mean	24.1	23.8	23.4	23.1	22.3	22.1	21.6	21.6	20.9	20.7	20.3	20.0	19.8	19.9	20.0

Table A.4: RMSPE for different contracts and models, New York

RMSPE	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
2 months															
Feb09	29.2	28.6	30.0	31.0	28.6	27.6	26.4	29.5	28.5	27.6	26.5	24.3	20.4	17.9	18.7
Mar09	21.3	21.2	20.6	19.9	16.4	17.9	17.0	18.9	17.7	16.9	19.0	18.5	21.5	25.2	26.6
Apr09	9.3	9.3	9.1	8.6	8.4	8.7	9.2	9.2	9.3	9.5	9.8	10.0	10.3	10.5	10.3
May09	22.1	21.9	21.6	20.8	20.6	19.7	18.9	18.4	18.2	17.8	17.7	17.5	17.6	17.7	17.8
Jun09	17.2	16.3	13.4	12.0	11.0	9.1	8.1	10.3	11.4	13.8	16.8	19.4	22.1	24.1	25.7
Jul09	44.5	42.1	40.2	37.8	36.1	34.5	32.3	31.6	30.1	29.9	28.8	28.2	26.8	26.7	26.2
Aug09	36.5	37.4	37.4	37.6	37.0	36.3	36.2	35.5	32.7	31.1	29.4	28.2	28.0	28.5	26.8
Sep09	25.0	22.1	20.7	19.6	17.7	16.9	15.9	16.1	15.1	15.0	13.6	13.1	12.5	12.4	12.7
Oct09	29.2	28.2	27.1	26.5	25.8	26.2	23.5	21.0	18.0	16.2	11.9	11.5	9.9	9.1	9.5
Nov09	29.3	29.3	28.5	27.0	25.2	23.3	22.1	19.8	19.2	19.6	22.4	23.1	24.7	26.7	27.6
Dec09	47.7	45.6	41.0	37.0	33.7	32.7	32.0	26.8	24.4	23.4	22.8	20.8	20.5	21.6	21.4
Jan10	58.5	57.7	56.8	55.0	52.2	50.0	48.8	47.3	45.1	43.2	40.0	38.1	36.8	35.1	33.6
Mean	30.8	30.0	28.9	27.7	26.1	25.2	24.2	23.7	22.5	22.0	21.6	21.1	20.9	21.3	21.4

Table A.5: RMSPE for different contracts and models, calculated for the last two months, New York

RMSPE	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
Vol>0															
Feb09	31.7	32.1	34.4	34.7	30.3	28.4	28.9	32.8	32.3	31.0	29.3	25.7	22.4	20.4	22.4
Mar09	25.4	24.8	25.3	24.2	20.7	22.2	22.4	24.5	22.6	21.7	23.6	23.7	26.7	31.7	32.9
Apr09	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
May09	16.2	16.5	16.6	16.1	16.3	13.8	11.9	10.7	10.8	11.5	13.1	14.2	15.2	16.1	16.6
Jun09	22.6	21.2	17.4	14.9	13.5	10.0	8.5	10.7	12.8	15.2	18.7	21.8	25.1	27.6	30.1
Jul09	53.7	50.8	47.4	44.9	42.7	39.2	34.3	33.8	31.2	31.3	29.6	28.4	24.8	24.6	25.1
Aug09	29.1	30.5	31.7	30.3	28.1	28.8	27.0	27.0	23.1	22.1	21.4	20.6	19.8	21.2	20.3
Sep09	32.4	26.9	24.5	22.5	18.9	17.9	15.3	15.1	14.5	15.3	12.7	11.6	11.5	11.9	12.7
Oct09	30.7	29.5	32.8	35.2	36.1	37.5	36.2	33.8	27.7	23.3	18.0	17.1	14.2	13.5	8.0
Nov09	40.1	40.9	39.7	35.6	30.3	26.7	23.6	24.5	23.4	24.3	28.6	27.2	29.0	31.3	32.8
Dec09	45.4	46.4	45.9	38.0	34.4	31.5	34.4	28.4	26.3	26.3	25.6	23.2	23.3	25.4	25.7
Jan10	75.4	74.1	73.8	68.4	65.5	64.5	60.5	58.0	53.6	50.5	44.2	41.8	39.4	37.8	35.8
Mean	36.6	35.8	35.4	33.2	30.6	29.1	27.6	27.2	25.3	24.8	24.1	23.2	22.9	23.8	23.8

Table A.6: RMSPE for different contracts and models, calculated for days with volume>0, New York

Corr. coeff.	Model														
	NMF	MF0	MF1	MF2	MF3	MF4	MF5	MF6	MF7	MF8	MF9	MF10	MF11	MF12	MF13
Feb09	0.41	0.39	0.31	0.25	0.31	0.31	0.32	0.05	0.13	0.24	0.32	0.41	0.56	0.62	0.60
Mar09	0.53	0.58	0.60	0.62	0.68	0.56	0.52	0.48	0.54	0.62	0.56	0.58	0.57	0.55	0.62
Apr09	0.56	0.56	0.58	0.61	0.64	0.63	0.62	0.61	0.59	0.58	0.56	0.54	0.52	0.51	0.53
May09	0.65	0.63	0.63	0.67	0.65	0.65	0.68	0.69	0.67	0.69	0.68	0.68	0.68	0.67	0.67
Jun09	0.94	0.95	0.96	0.96	0.97	0.98	0.98	0.98	0.98	0.98	0.98	0.97	0.96	0.96	0.96
Jul09	0.88	0.90	0.91	0.92	0.93	0.94	0.95	0.95	0.95	0.95	0.96	0.96	0.97	0.97	0.97
Aug09	0.26	0.20	0.22	0.20	0.20	0.24	0.25	0.29	0.47	0.52	0.55	0.58	0.57	0.56	0.61
Sep09	0.71	0.75	0.77	0.78	0.79	0.80	0.81	0.80	0.80	0.80	0.82	0.83	0.83	0.84	0.84
Oct09	0.62	0.64	0.67	0.69	0.72	0.72	0.77	0.81	0.86	0.88	0.94	0.94	0.96	0.96	0.96
Nov09	0.85	0.84	0.85	0.87	0.88	0.89	0.90	0.92	0.92	0.92	0.90	0.91	0.90	0.89	0.89
Dec09	0.63	0.68	0.74	0.79	0.82	0.81	0.82	0.85	0.85	0.86	0.84	0.84	0.83	0.81	0.81
Jan10	0.40	0.41	0.42	0.47	0.54	0.59	0.63	0.67	0.70	0.74	0.80	0.84	0.87	0.89	0.90
Mean	0.62	0.63	0.64	0.65	0.68	0.68	0.69	0.67	0.71	0.73	0.74	0.76	0.77	0.77	0.78

Table A.7: Correlation coefficient for different contracts and models, New York

RMSPE	Model		
	NMF	MF13	MF13perfect
Feb09	28.8	18.8	33.3
Mar09	22.9	25.2	24.8
Apr09	6.8	7.7	5.8
May09	21.2	20.8	21.0
Jun09	17.2	19.9	21.5
Jul09	27.0	18.7	17.0
Aug09	20.2	15.5	19.4
Sep09	18.5	16.4	17.3
Oct09	24.1	8.0	11.3
Nov09	26.8	27.1	25.3
Dec09	38.3	31.7	32.4
Jan10	38.0	30.5	31.4
Mean	24.1	20.0	21.7

Table A.8: RMSPE for different contracts and models, New York

SFB 649 Discussion Paper Series 2010

For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 001 "Volatility Investing with Variance Swaps" by Wolfgang Karl Härdle and Elena Silyakova, January 2010.
- 002 "Partial Linear Quantile Regression and Bootstrap Confidence Bands" by Wolfgang Karl Härdle, Ya'acov Ritov and Song Song, January 2010.
- 003 "Uniform confidence bands for pricing kernels" by Wolfgang Karl Härdle, Yarema Okhrin and Weining Wang, January 2010.
- 004 "Bayesian Inference in a Stochastic Volatility Nelson-Siegel Model" by Nikolaus Hautsch and Fuyu Yang, January 2010.
- 005 "The Impact of Macroeconomic News on Quote Adjustments, Noise, and Informational Volatility" by Nikolaus Hautsch, Dieter Hess and David Veredas, January 2010.
- 006 "Bayesian Estimation and Model Selection in the Generalised Stochastic Unit Root Model" by Fuyu Yang and Roberto Leon-Gonzalez, January 2010.
- 007 "Two-sided Certification: The market for Rating Agencies" by Erik R. Fasten and Dirk Hofmann, January 2010.
- 008 "Characterising Equilibrium Selection in Global Games with Strategic Complementarities" by Christian Basteck, Tijmen R. Daniels and Frank Heinemann, January 2010.
- 009 "Predicting extreme VaR: Nonparametric quantile regression with refinements from extreme value theory" by Julia Schaumburg, February 2010.
- 010 "On Securitization, Market Completion and Equilibrium Risk Transfer" by Ulrich Horst, Traian A. Pirvu and Gonçalo Dos Reis, February 2010.
- 011 "Illiquidity and Derivative Valuation" by Ulrich Horst and Felix Naujokat, February 2010.
- 012 "Dynamic Systems of Social Interactions" by Ulrich Horst, February 2010.
- 013 "The dynamics of hourly electricity prices" by Wolfgang Karl Härdle and Stefan Trück, February 2010.
- 014 "Crisis? What Crisis? Currency vs. Banking in the Financial Crisis of 1931" by Albrecht Ritschl and Samad Sarferaz, February 2010.
- 015 "Estimation of the characteristics of a Lévy process observed at arbitrary frequency" by Johanna Kappusl and Markus Reiß, February 2010.
- 016 "Honey, I'll Be Working Late Tonight. The Effect of Individual Work Routines on Leisure Time Synchronization of Couples" by Juliane Scheffel, February 2010.
- 017 "The Impact of ICT Investments on the Relative Demand for High-Medium-, and Low-Skilled Workers: Industry versus Country Analysis" by Dorothee Schneider, February 2010.
- 018 "Time varying Hierarchical Archimedean Copulae" by Wolfgang Karl Härdle, Ostap Okhrin and Yarema Okhrin, February 2010.
- 019 "Monetary Transmission Right from the Start: The (Dis)Connection Between the Money Market and the ECB's Main Refinancing Rates" by Puriya Abbassi and Dieter Nautz, March 2010.
- 020 "Aggregate Hazard Function in Price-Setting: A Bayesian Analysis Using Macro Data" by Fang Yao, March 2010.
- 021 "Nonparametric Estimation of Risk-Neutral Densities" by Maria Grith, Wolfgang Karl Härdle and Melanie Schienle, March 2010.

SFB 649 Discussion Paper Series 2010

For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 022 "Fitting high-dimensional Copulae to Data" by Ostap Okhrin, April 2010.
- 023 "The (In)stability of Money Demand in the Euro Area: Lessons from a Cross-Country Analysis" by Dieter Nautz and Ulrike Rondorf, April 2010.
- 024 "The optimal industry structure in a vertically related market" by Raffaele Fiocco, April 2010.
- 025 "Herding of Institutional Traders" by Stephanie Kremer, April 2010.
- 026 "Non-Gaussian Component Analysis: New Ideas, New Proofs, New Applications" by Vladimir Panov, May 2010.
- 027 "Liquidity and Capital Requirements and the Probability of Bank Failure" by Philipp Johann König, May 2010.
- 028 "Social Relationships and Trust" by Christine Binzel and Dietmar Fehr, May 2010.
- 029 "Adaptive Interest Rate Modelling" by Mengmeng Guo and Wolfgang Karl Härdle, May 2010.
- 030 "Can the New Keynesian Phillips Curve Explain Inflation Gap Persistence?" by Fang Yao, June 2010.
- 031 "Modeling Asset Prices" by James E. Gentle and Wolfgang Karl Härdle, June 2010.
- 032 "Learning Machines Supporting Bankruptcy Prediction" by Wolfgang Karl Härdle, Rouslan Moro and Linda Hoffmann, June 2010.
- 033 "Sensitivity of risk measures with respect to the normal approximation of total claim distributions" by Volker Krätschmer and Henryk Zähle, June 2010.
- 034 "Sociodemographic, Economic, and Psychological Drivers of the Demand for Life Insurance: Evidence from the German Retirement Income Act" by Carolin Hecht and Katja Hanewald, July 2010.
- 035 "Efficiency and Equilibria in Games of Optimal Derivative Design" by Ulrich Horst and Santiago Moreno-Bromberg, July 2010.
- 036 "Why Do Financial Market Experts Misperceive Future Monetary Policy Decisions?" by Sandra Schmidt and Dieter Nautz, July 2010.
- 037 "Dynamical systems forced by shot noise as a new paradigm in the interest rate modeling" by Alexander L. Baranovski, July 2010.
- 038 "Pre-Averaging Based Estimation of Quadratic Variation in the Presence of Noise and Jumps: Theory, Implementation, and Empirical Evidence" by Nikolaus Hautsch and Mark Podolskij, July 2010.
- 039 "High Dimensional Nonstationary Time Series Modelling with Generalized Dynamic Semiparametric Factor Model" by Song Song, Wolfgang K. Härdle, and Ya'acov Ritov, July 2010.
- 040 "Stochastic Mortality, Subjective Survival Expectations, and Individual Saving Behavior" by Thomas Post and Katja Hanewald, July 2010.
- 041 "Prognose mit nichtparametrischen Verfahren" by Wolfgang Karl Härdle, Rainer Schulz, and Weining Wang, August 2010.
- 042 "Payroll Taxes, Social Insurance and Business Cycles" by Michael C. Burda and Mark Weder, August 2010.
- 043 "Meteorological forecasts and the pricing of weather derivatives" by Matthias Ritter, Oliver Mußhoff, and Martin Odening, September 2010.