

Modeling Asymmetric Volatility in Weekly Dutch Temperature Data*

Philip Hans Franses[†]

Econometric Institute, Erasmus University Rotterdam

Jack Neele[‡]

Econometric Institute, Erasmus University Rotterdam

Dick van Dijk[§]

Tinbergen Institute, Erasmus University Rotterdam

September 21, 1998

Abstract

In addition to clear-cut seasonality in mean and variance, weekly Dutch temperature data appear to have a strong asymmetry in the impact of unexpectedly high or low temperatures on conditional volatility. Furthermore, this asymmetry also shows fairly pronounced seasonal variation. To describe these features, we propose a univariate seasonal time series model with asymmetric

*This paper was prepared for a special issue of *Environmental Modeling and Software*. We thank the editor Michael McAleer for several helpful comments and Richard Tol for making the Dutch temperature data available to us.

[†]Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR, Rotterdam, The Netherlands, email: franses@few.eur.nl (corresponding author).

[‡]Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR, Rotterdam, The Netherlands.

[§]Tinbergen Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, email: djvandijk@few.eur.nl

conditionally heteroskedastic errors. We fit this (and other, nested) model(s) to 25 years of weekly data. We evaluate its forecasting performance for 5 years of hold-out data and find that the imposed asymmetry leads to better out-of-sample forecasts of temperature volatility.

1 Introduction

High frequency temperature data, like daily or weekly data, have several characteristic features. First and most obvious, the mean temperature shows substantial seasonal variation. In The Netherlands for example, daytime temperatures range between 0 and 4 degrees Celsius in winter, while daytime temperatures vary from 16 to 20 degrees Celsius in summer. Second, the volatility of temperatures is not constant within the year but appears to follow a fairly regular seasonal pattern as well. At the beginning of Dutch winters, the standard deviation of weekly temperatures is almost twice as large as at the end of summer. This implies that temperatures are less predictable in winter than in summer.

A third feature of Dutch temperature data, documented in Tol (1996), is that large (small) absolute deviations from the mean tend to cluster. As a consequence, the conditional forecastability of temperatures also varies within summer and winter. Interestingly, the same feature holds for many high frequency financial time series (such as daily stock market returns and interest rates). To describe this volatility clustering in empirical finance one often uses the so-called Autoregressive Conditionally Heteroskedastic [ARCH] model put forward by Engle (1982). Over the last fifteen years, this model has been the subject of intensive research, see the surveys of Bollerslev, Chou and Kroner (1992), Bera and Higgins (1993), and Bollerslev, Engle and Nelson (1994), among others. A popular extension of the basic ARCH model is the Generalized ARCH [GARCH] model, see Bollerslev (1986). To capture time-varying predictability, Tol (1996) fits a GARCH model to daily Dutch temperature data in winter and summer periods, and demonstrates its usefulness for describing the volatility clustering feature of the data.

In the present paper we show that Dutch temperature data have yet another feature. This fourth property is that the impact of temperatures lower than expected on conditional volatility is different from the impact of temperatures higher than expected. Furthermore, this impact is changing over the year as well. In particular,

the correlation between conditional volatility and the ‘surprise’ in the temperature is negative in winter and positive in summer. Hence, in winter (summer) temperatures lower than expected lead to larger (smaller) conditional variance than temperatures higher than expected.

The aim of this paper is to develop a time series model which is capable of describing the above-mentioned four features for weekly temperatures in The Netherlands, observed over a period of 30 years. The plan of the rest of this paper is as follows. First, in Section 2 we discuss the stylized facts of the weekly temperature data in more detail. In Section 3, we introduce a variant of the Quadratic GARCH [QGARCH] model of Engle and Ng (1993) and Sentana (1995) that can capture all observed features. Section 4 presents the in-sample estimation and out-of-sample forecasting results. We estimate the proposed model and two nested versions using the first 25 years of data, while we hold out the last 5 years to evaluate their forecasting performance. Both the in- and out-of-sample evidence suggests that a model with asymmetric volatility is to be preferred. In Section 5, we conclude this paper with some remarks.

2 Weekly Dutch Temperature Data

In this section we document the four characteristic features of Dutch temperature data mentioned in the Introduction. The time series under scrutiny, denoted y_t , is the mean weekly temperature in The Netherlands, which is constructed from the daily series analyzed in Tol (1996) by simple averaging (over 7 days). The first observation in every year is taken to be the week which starts on the first day of February. This month is halfway through the Dutch winter, and usually has the lowest temperatures. Given this choice we can describe seasonal patterns in the data quite easily, see below. Due to the fact that there are 365 or 366 daily observations per year, the 52-nd weekly observation concerns an average over 8 or 9 days, respectively. The sample ranges from 1961 until 1990. We use the first 25

years (1300 observations) for estimation of the models to be presented in Section 3 and leave out the last 5 years for forecast evaluation. It should be noted that our time series ends in December 1990, so that our hold-out sample contains 255 and not 260 weekly observations.

- insert Figure 1 about here -

In Figure 1 the weekly temperatures for the estimation sample 1961-1985 are plotted against the week of observation. Obviously, there is a clear-cut seasonal pattern in the data, with the lowest values being attained in January and February, and the highest values in July and August.

- insert Figure 2 about here -

Closer inspection of Figure 1 also suggests that the variation in temperatures is larger in winter than in summer. In Figure 2 we give the weekly standard deviation, which is calculated using only the observations for the particular week. It is seen that the standard deviation at the onset of winter is almost twice as large as at the end of summer. This suggests that it may be less easy to predict temperatures during the winter. Also note that the standard deviation is declining from January until the end of August, approximately, and increasing from September until the end of December. Hence, the increase in volatility occurs much faster than the decrease.

To obtain some more insight into the seasonal variation of the mean and standard deviation of the temperature data, we estimate the following nonparametric regression model with heteroskedastic errors,

$$y_t = m(T_t) + s(T_t)\eta_t, \quad t = 1, \dots, n, \quad (1)$$

where $m(T_t)$ and $s(T_t)$ are the conditional mean and standard deviation of y_t , respectively, T_t denotes the number of the week, i.e., $T_t = t \bmod 52$, η_t are i.i.d. random variables, $E(\eta_t) = 0$, $E(\eta_t^2) = 1$, and n denotes sample size, i.e., $n = 1300$.

The functions $m(T_t)$ and $s(T_t)$ are unknown and are estimated using kernel regression techniques. To be more precise, we estimate the conditional mean m by the Nadaraya-Watson estimator

$$\hat{m}(x) = \frac{\sum_{t=1}^n K((T_t - x)/h)y_t}{\sum_{t=1}^n K((T_t - x)/h)}, \quad (2)$$

while the conditional variance s^2 is estimated by

$$\hat{s}^2(x) = \frac{\sum_{t=1}^n K((T_t - x)/h)y_t^2}{\sum_{t=1}^n K((T_t - x)/h)} - \hat{m}^2(x). \quad (3)$$

For the kernel function $K(\cdot)$ we use the following variant of the Epanechnikov kernel

$$K(u) = \frac{3}{4} \left[1 - \{\min(|u|, 52 - |u|)\}^2 \right] 1_{\{|u| < 1\}}, \quad (4)$$

since this kernel takes into account the observed seasonality, see also Hyndman and Wand (1997). The bandwidth h is taken to be four weeks. For an elaborate discussion of the model given in (1) and estimators of the functions $m(T_t)$ and $s(T_t)$ we refer to Härdle and Tsybakov (1997). More general introductions to kernel regression and other nonparametric techniques can be found in Wand and Jones (1995) and Fan and Gijbels (1996), among others. The estimates $\hat{m}(T_t)$ and $\hat{s}(T_t)$ are shown as solid lines in Figures 1 and 2, respectively, and confirm the observations made above.

- insert Figure 3 about here -

From Figure 3 we observe that the skewness of the temperature series also varies throughout the year. Interestingly, there is positive skewness in summer while skewness is negative in winter weeks. This means that one may expect more weeks which are warmer than average in summer, and more weeks which are colder than average in winter.

The observations on the variance and skewness made above lead us to consider the relationship between the level of this week's temperature and the volatility of next week's temperature. In particular, we try to address the questions whether this

relationship is symmetric, in the sense that an unexpectedly low temperature has the same effect on volatility than an unexpectedly high temperature, and whether it is constant throughout the year. To investigate this we hypothesize the following model for the residuals $\varepsilon_t \equiv s(T_t)\eta_t$ from (1),

$$\varepsilon_t^2 = \beta_0(T_{t-1}) + \beta_1(T_{t-1})\varepsilon_{t-1} + \beta_2(T_{t-1})\varepsilon_{t-1}^2 + \nu_t, \quad (5)$$

where $\beta_i(T_t), i = 0, 1, 2$ are unknown functions of T_t and ν_t are i.i.d. random variables. If the $\beta_i(T_t), i = 0, 1, 2$, are constant, the model reduces to the Quadratic GARCH [QGARCH] model of Sentana (1995). The purpose of estimating (5) is to obtain some preliminary idea as to whether it is worthwhile to specify a QGARCH model with time-varying coefficients. We will return to this model in the next section. It is useful to note here that, assuming that the distribution of ε_t is symmetric, the regressors in (5) are orthogonal. It follows that $\beta_2(T_t) = \text{corr}_{\varepsilon_t^2, \varepsilon_{t-1}^2}(T_t)$ and $\beta_1(T_t) = \text{cov}_{\varepsilon_t^2, \varepsilon_{t-1}}(T_t)/\text{var}_{\varepsilon_t}(T_t)$. Hence β_2 equals the ‘local’ first-order autocorrelation of the squared residuals, while β_2 measures the (local) asymmetry in the impact of positive and negative shocks to the temperature on volatility.

The unknown functions $\beta_i(T_t), i = 0, 1, 2$, are estimated using a variant of the locally weighted regression [LWR] technique first introduced by Cleveland (1979), see also Cleveland and Devlin (1988) and Cleveland, Devlin and Grosse (1988). In the original LWR technique, the variables which determine the parameters are assumed to be the same as the regressors in the model, while here they are different. We employ the local linear estimator developed by Chen (1998) and Cai, Fan and Yao (1998) in the context of autoregressive models to estimate $\beta(x) = (\beta_0(x), \beta_1(x), \beta_2(x))'$, which is given by

$$\hat{\beta}(x) = (X'W_xX)^{-1} X'W_xY, \quad (6)$$

where $X = (X'_1, \dots, X'_{n-1})'$ with $X_t = (1, \varepsilon_t, \varepsilon_t^2)'$, $Y = (\varepsilon_2^2, \dots, \varepsilon_n^2)'$ and W_x is a diagonal matrix with t -th diagonal element equal to $K((T_{t-1} - x)/h)$, which is computed according to (4). Again a bandwidth h of four weeks is used. To see the

intuition behind this estimator, notice that (6) is the solution to the weighted least squares problem

$$\hat{\beta}(x) = \arg \min_{\beta} \sum_{t=2}^n (\varepsilon_t^2 - X'_{t-1} \beta)^2 K((T_{t-1} - x)/h). \quad (7)$$

- insert Figure 4 about here -

The estimates of $\beta_i(x)$, $i = 0, 1, 2$ together with 90% pointwise confidence bands are displayed in Figure 4. The upper panel shows the by now familiar seasonal pattern in volatility. The middle panel demonstrates that the relation between this week's volatility and last week's 'surprise' in temperature, is nonzero, suggesting that the impact of temperatures lower and higher than expected is asymmetric indeed. Furthermore, it appears that this asymmetry shows seasonal variation as well: in winter it appears to be negative, while in summer it is positive, that is, in winter (summer) temperatures lower than expected lead to larger (smaller) conditional variance than temperatures higher than expected. Finally, the lower panel of this Figure suggests that the correlation in the squared series ε_t^2 is significant (ly positive) only in the second half of the year.

- insert Figure 5 about here -

As a final characteristic of the temperature series we give the weekly kurtosis, again measured over the first 25 years, in Figure 5. In contrast to the previous measures, there is not much seasonal variation in the kurtosis, apart from a few weeks in summer. Notice also the kurtosis exceeds 3 (that is, the kurtosis corresponding to the normal distribution) in summer weeks, indicating that there have been some very warm summers in the Netherlands.

3 The Models

In this section we introduce the model which is used to describe the various features of our weekly temperature series. Summarizing the evidence presented in Section

2 and Tol (1996), the model should allow for i) seasonal variation in the mean, ii) seasonal variation in the variance, iii) volatility clustering, and iv) changing asymmetry in the relation between last week's temperature and the volatility of this week's temperature.

Given the visual evidence in Figures 1 and 2, we decide to consider simple parabolic functions to describe the seasonal variation in the mean and variance. For this purpose we use the variable T_t as defined in Section 2, which runs from 1 to 52 in every year. We include T_t as well as its square in both the conditional mean and conditional variance equation. Furthermore, some preliminary experimentation shows that we need to include y_{t-1} in the conditional mean equation to accommodate serial correlation. Hyndman and Wand (1997) demonstrate that the correlation for a daily Australian temperature series is also changing over the year. It appears that this is not a critical issue for our weekly temperature series – hence, we assume this correlation to be constant.

The possible presence of volatility clustering (or: time-varying predictability) in temperature data can be accommodated by means of a GARCH model, see also Tol (1996). A modification of this model that also allows for asymmetry in the impact of innovations on the conditional variance is the so-called Quadratic GARCH [QGARCH] model, proposed in Engle and Ng (1993) and analyzed in detail in Sentana (1995). There are several other extensions of the basic GARCH model that are able to describe this correspondence, see Hentschel (1995) for a concise review. However, in, for example, Franses and Dijk (1996), it is found that the QGARCH model frequently outperforms its rivals in terms of out-of-sample forecasting. We therefore limit our attention in this paper to the QGARCH model.

In sum, our proposed model for the weekly temperature data is

$$y_t = \mu_0 + \mu_1 T_t + \mu_2 T_t^2 + \phi y_{t-1} + \varepsilon_t, \quad t = 1, 2, \dots, 1300, \quad (8)$$

$$\varepsilon_t = \sigma_t \eta_t, \quad (9)$$

$$\sigma_t^2 = \omega_0 + \omega_1 T + \omega_2 T^2 + \alpha \left(\varepsilon_{t-1} - \gamma_0 - \gamma_1 T - \gamma_2 T^2 \right)^2 + \beta \sigma_{t-1}^2, \quad (10)$$

where y_0 is the last week of January 1961 and η_t again are i.i.d. random variables with mean zero and variance one. In addition, we assume the η_t are normally distributed. The third equation describes the conditional variance of ε_t . We follow Tol (1996) and only include ε_{t-1} and σ_{t-1}^2 to describe the volatility dynamics. The γ_0, γ_1 and γ_2 parameters deal with the asymmetric impact of shocks to the temperature on conditional volatility and with the possible change in this impact over the year. Notice that (10) reduces to a standard GARCH(1,1) model (with deterministic seasonal variation in the - unconditional - variance) when all three γ parameters are equal to zero. Finally, (10) can be rewritten in terms of ε_t^2 as

$$\begin{aligned} \varepsilon_t^2 = & \omega_0 + \omega_1 T + \omega_2 T^2 + \left(\gamma_0 + \gamma_1 T + \gamma_2 T^2 \right)^2 \\ & - 2\alpha \left(\gamma_0 + \gamma_1 T + \gamma_2 T^2 \right) \varepsilon_{t-1} + (\alpha + \beta) \varepsilon_{t-1}^2 + \nu_t - \beta \nu_{t-1}, \end{aligned} \quad (11)$$

where $\nu_t = \varepsilon_t^2 - \sigma_t^2$. Comparing this with (5), the similarities of our QGARCH model with the nonparametric model (5) are easily seen.

The unrestricted model (8)-(10) should be capable of describing all four features of the temperature series. In addition we estimate two simpler, nested models by imposing restrictions on some of the parameters. First, we restrict γ_1 and γ_2 equal to zero to investigate whether the asymmetric impact of shocks in the temperature on the conditional volatility is changing over the year. Second, γ_0 is also restricted to zero to check whether there is any asymmetry in this impact at all. Summarizing, in the next section we compare the following models:

GARCH	$\gamma_0 = \gamma_1 = \gamma_2 = 0$;
QGARCH-I	no restrictions;
QGARCH-II	$\gamma_1 = \gamma_2 = 0$;

Given the visual evidence in Figure 4, we hypothesize that the QGARCH models outperform the linear GARCH model.

4 Empirical Results

In this section we evaluate the in-sample estimation results and the out-of-sample forecasting performance of the three competing GARCH models.

4.1 In-sample estimation

The parameters in the three models are estimated using maximum likelihood methods. The relevant Gauss-code is available upon request.

In addition to the parameter estimates, we calculate the Box-Pierce test statistic for residual autocorrelation at lags 1 to 10 for the scaled residuals $\hat{\eta}_t \equiv \hat{\varepsilon}_t/\hat{\sigma}_t$ and their squares $\hat{\eta}_t^2$ to verify if (8) and (10) include enough dynamics to render i.i.d. residuals. We also calculate the Jarque-Bera test for normality of the scaled residuals, as assumed in (9). Clear rejection of this assumption may lead to consideration of alternative distributions for η_t . Finally, we compute the log likelihood and, the Akaike Information Criterion [AIC] and Schwarz's BIC in order to compare the three models.

- insert Table 1 about here -

The in-sample estimation results are given in Table 1. From its bottom panel we see that the three models pass the residual correlation tests (the 5% critical value is 16.92), albeit the GARCH model passes the test for the scaled residuals only at the 2.5% level. Furthermore, all three models display non-normality of the scaled residuals at the 5% level, but not at the 1% level. This last result is perhaps not unexpected given the large number of observations and the rather small number of model parameters. To us, assuming alternative distributions for η_t therefore does not seem necessary. Finally, the values of the AIC indicate that the QGARCH-II model is to be preferred, whereas the values of the BIC suggest to select the GARCH model. It is well-known that Schwarz's criterion penalizes the inclusion of additional parameters rather severely, such that the improvement in fit has to be substantial in order to be justified.

The top panel of Table 1 contains the parameter estimates and associated t -ratios. The parameter α is not significant in the estimated GARCH model, which may be due to neglected asymmetry. The significant and positive estimates of γ_0 in the

QGARCH-I and -II models shows that a negative value of the temperature (after correction for seasonality and dynamics) implies a larger value of the conditional variance than does an above average value of the same size. In the QGARCH-I model the γ_1 and γ_2 parameters are not significant. Hence, the seasonality in the asymmetry suggested by Figure 4 appears not to be significant after all.

4.2 Out-of-sample Forecasting of Volatility

As an alternative way to evaluate the three volatility models and to compare their ability to describe the features in the temperature series we investigate their out-of-sample forecasting performance. We calculate one-step ahead forecasts of the conditional variance σ_t for the 255 observations in the hold-out sample. For every forecast, we re-estimate the parameters in the three models, using all observations prior to the forecast origin. Given the possible presence of aberrant observations in this hold-out sample, we evaluate the models using the Median Squared Error (MedSE) criterion. As a measure of the true variance, we use the squared residuals $\hat{\varepsilon}_t^2$ from (8) (obtained in each estimation round) when no GARCH model is fitted to the data.

- insert Table 2 about here -

The forecasting results are summarized in Table 2. The linear GARCH model appears to be the best for 1986. The QGARCH-I model beats the other models in 1987 and 1989, while the QGARCH-II model is best for 1988 and 1990. For the whole period, the QGARCH-II model clearly outperforms its rivals.

- insert Figure 6 about here -

Figure 6 shows the median of the difference of the squared forecast errors, which enables a pair-wise comparison of the three models. The middle panel for example reveals that the QGARCH-II models achieves the largest gains in forecast accuracy

relative to the GARCH model in the months July, September and December. The bottom panel suggests that the main difference in the QGARCH-I and QGARCH-II models occurs in April. These results provides us with additional confidence in the usefulness of the nonlinear QGARCH model.

5 Concluding Remarks

In this paper we have proposed and evaluated a nonlinear GARCH model for weekly temperatures in The Netherlands. Both the in-sample estimation results and out-of-sample forecasting performance suggest that our nonlinear GARCH model is superior to a linear GARCH model, thereby confirming visual evidence on an asymmetric relation between this week's surprise in temperature and the volatility of next week's temperature. Our model implies that temperatures lower than expected lead to larger conditional forecasting intervals (and hence less predictability) than do temperatures higher than expected. Whether this empirical observation can be attributed to certain meteorological phenomena is a topic for further research. Additionally, it may be of interest to examine if this model for Dutch data also fits temperature data for other countries.

References

- Bera, A.K. and M.L. Higgins, 1993, A survey of ARCH models: properties, estimation and testing, *Journal of Economic Surveys* **7**, 305–366.
- Bollerslev, T., 1986, Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics* **31**, 307–327.
- Bollerslev, T., R.F. Engle and D.B. Nelson, 1994, ARCH models, in R.F. Engle and D.L. McFadden (editors), *Handbook of Econometrics IV*, Elsevier Science, Amsterdam, pp. 2961–3038.

- Bollerslev, T., R.Y. Chou and K.F. Kroner, 1992, ARCH modeling in finance: a review of the theory and empirical evidence, *Journal of Econometrics* **52**, 5–59.
- Cai, Z., J. Fan and Q. Yao, 1998, Functional-coefficient regression models for non-linear time series, working paper, University of South Carolina.
- Chen, R., 1998, Functional coefficient autoregressive models: estimation and tests of hypotheses, Working paper 98-10, Humboldt Universität Berlin.
- Cleveland, W.S., 1979, Robust locally weighted regression and smoothing scatterplots, *Journal of the American Statistical Association* **74**, 829–836.
- Cleveland, W.S. and S.J. Devlin, 1988, Locally weighted regression: an approach to regression analysis by local fitting, *Journal of the American Statistical Association* **83**, 596–610.
- Cleveland, W.S., S.J. Devlin and E. Grosse, 1988, Regression by local fitting - methods, properties, and computational algorithms, *Journal of Econometrics* **37**, 87–114.
- Engle, R.F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* **50**, 987–1007.
- Engle, R.F. and V.K. Ng, 1993, Measuring and testing the impact of news on volatility, *Journal of Finance* **48**.
- Fan, J. and I. Gijbels, 1996, *Local Polynomial Modeling and its Applications*, Chapman and Hall, London.
- Franses, Ph.H. and D. Van Dijk, 1996, Forecasting stock market volatility using nonlinear GARCH models, *Journal of Forecasting* **15**, 229–235.
- Härdle, W. and A. Tsybakov, 1997, Local polynomial estimators of the volatility function in nonparametric autoregression, *Journal of Econometrics* **81**, 223–242.

- Hentschel, L.F., 1995, All in the family: Nesting linear and nonlinear GARCH models, *Journal of Financial Economics* **39**, 139–164.
- Hyndman, R.J. and M.P. Wand, 1997, Nonparametric autocovariance function estimation, *Australian and New Zealand Journal of Statistics* **39**, 313–324.
- Sentana, E., 1995, Quadratic ARCH models, *Review of Economic Studies* **62**, 639–661.
- Tol, R.J.S., 1996, Autoregressive conditional heteroscedasticity in daily temperature measurements, *Environmetrics* **7**, 67–75.
- Wand, M.P. and M.C. Jones, 1995, *Kernel Smoothing*, Chapman and Hall, London.

Table 1: In-sample parameter estimates¹

Parameters	Model		
	GARCH	QGARCH-I	QGARCH-II
	<u>Autoregressive part</u>		
μ_0	-0.42 (-1.60)	-0.58 (-2.22)	-0.45 (-1.89)
μ_1	6.53 (15.28)	6.46 (15.31)	6.29 (15.75)
μ_2	-1.30 (-16.40)	-1.33 (-16.32)	-1.30 (-16.76)
ϕ	0.54 (22.44)	0.53 (22.23)	0.54 (22.75)
	<u>GARCH part</u>		
ω_0	0.35 (1.09)	0.15 (0.64)	0.42 (1.44)
ω_1	-0.37 (-2.01)	-0.37 (-2.83)	-0.42 (-2.61)
ω_2	0.11 (3.21)	0.12 (4.57)	0.12 (3.85)
α	0.01 (0.49)	0.03 (2.57)	0.03 (2.63)
β	0.94 (26.12)	0.90 (38.92)	0.89 (31.99)
γ_0		3.02 (2.38)	1.91 (2.77)
γ_1		0.17 (0.42)	
γ_2		-0.15 (-1.27)	
Diagnostics			
BP(10) for $\hat{\eta}_t$	17.74	16.09	15.86
BP(10) for $\hat{\eta}_t^2$	11.74	9.20	11.46
JB for $\hat{\eta}_t$	7.93	9.42	7.79
LL	-2993.2	-2990.5	-2990.7
AIC	6004.4	6005.1	6001.4
BIC	6050.9	6067.0	6053.1


¹ The various models are estimated using 1300 weekly observations from 1961-1985. t -statistics are given in parentheses below the parameter estimates. BP(10) denotes the Box-Pierce statistic for autocorrelation, JB denotes the Jarque-Bera test for normality and LL denotes log likelihood. The AIC is computed as $-2LL + 2k$, with k the number of parameters. The BIC is computed as $-2LL + k \ln(n)n$, with n the number of observations.

Table 2: Out-of-sample volatility forecasting performance based on the Median Squared Error¹

Year	Model		
	GARCH	QGARCH-I	QGARCH-II
1986	16.81	19.41	18.95
1987	12.55	10.70	12.65
1988	17.45	16.01	12.11
1989	18.55	14.57	15.94
1990	16.61	14.32	12.16
1986-1990	16.60	15.86	14.63


¹ Entries denote the Median Squared Forecasting Error for the conditional variance of weekly Dutch temperature. The squared residuals from estimating (8) by least squares are used as measure of true volatility.

Figure 1: Weekly temperatures




Note: Weekly temperatures in the Netherlands, 1961-1985, plotted against the week of observation. The solid line is the estimate of the conditional mean function $m(T_t)$ in (1).

Figure 2: Standard deviation of weekly temperatures




Note: Estimates of the (conditional) standard deviation for weekly temperatures in the Netherlands. The estimates shown by circles are obtained using only the observations from each individual week. The solid line is the estimate of the conditional standard deviation $s(T_t)$ in (1).

Figure 3: Skewness of weekly temperatures



Note: Estimated skewness of weekly temperatures in the Netherlands, 1961-1985. The estimates are obtained by computing the skewness using only the observations from a particular week.

Figure 4: Correlation and asymmetry in the volatility of weekly temperatures



Note: Nonparametric estimates of the unknown functions $\beta_0(x)$, $\beta_1(x)$ and $\beta_2(x)$ (from top to bottom) in the model (5) (solid lines), together with upper and lower limits of 90% confidence intervals (dashed lines). The estimates are obtained from (6) using (4) as kernel with the bandwidth h equal to four weeks.

Figure 5: Kurtosis of weekly temperatures