Regression Models for Ordinal Valued Time Series: Applications in High Frequency Finance and Medicine

Gernot Müller, Claudia Czado, Stefan Antes, Martin Rottenwallner

Zentrum Mathematik
Technische Universität München
Boltzmannstraße 3
85747 Garching bei München
Germany

May 2003

Abstract

In this paper we investigate intraday data of the IBM stock and a time series representing the sleep states of a newborn child. In both cases we are interested in the influence of several covariates observed together with the response series. For this purpose we use on the one hand the regression model proposed in Müller and Czado (2002), on the other hand the ordered probit model. The parameters are estimated with the GM-MGMC algorithm described in Müller and Czado (2002). Predictions are computed to test the results.

Keywords: Discrete-valued time series; High-frequency finance; Markov Chain Monte Carlo; Ordered Probit; Regression.

1 Introduction

In this paper we investigate two discrete valued time series and search for significant covariates in both cases. The first time series represents the absolute price differences of the IBM stock at the New York Stock Exchange (NYSE) on November 29, 2000. The data is taken from the TAQ2 data base of the NYSE. In recent years many propositions were made to model such high frequency financial data. For a global overview about high frequency finance see Bauwens and Giot (2001) or Dacorogna, Gençay, Müller, Olsen, and Pictet (2001). In our time series only price differences that are integer multiples of one sixteenth of a dollar are observed together with several covariates. This time series is investigated in Section 2.
The second time series, which is investigated in Section 3, represents the sleep states of a newborn child measured on an integer scale from one to four. Here we have information about the heart rate and the temperature of the child.
Therefore in both cases we have an observation of a discrete response time series $\{Y_t, t = 1, ..., T\}$, where Y_t takes on only K different values and a $(p+1)$-dimensional vector $\mathbf{X}_t = (1, X_{t1}, ..., X_{tp})'$ of real-valued covariates for each $t \in \{1, ..., T\}$. We assume that there exists an underlying unobserved real-valued time series $\{Y_t^*, t = 1, ..., T\}$ which produces Y_t by thresholding. More precisely:

$$
\begin{aligned}
Y_t &= k \iff Y_t^* \in [\alpha_{k-1}, \alpha_k)\,, \qquad k \in \{1, \ldots, K\} \qquad (1.1)\\
Y_t^* &= \mathbf{X}_t'\boldsymbol{\beta} + \phi Y_{t-1}^* + \varepsilon_t^*, \qquad t \in \{1, \ldots, T\} \qquad (1.2)
\end{aligned}
$$

where $\boldsymbol{\beta}$ is the vector of the regression coefficients and ϕ the autoregressive parameter. The cutpoints α_j have to fulfill the order condition $-\infty = \alpha_0 < \alpha_1 < ... < \alpha_{K-1} < \alpha_K = \infty$. All latent variables are marked with an asterisk. Further we assume $\varepsilon_t^* \sim N(0, \delta^2)$ i.i.d.. Since the vector of covariates contains an intercept we have to fix α_1 for reasons of identifiability. In particular we set $\alpha_1 = 0$. For the same reasons we have to fix the variance δ^2. Therefore we set $\delta^2 = 1$.
For more details concerning this model and a MCMC estimation algorithm for the latent variables and parameters see Müller and Czado (2002).

2 IBM data

In this section we want to apply the model defined above to intraday data from the IBM stock on 29th of November 2000. The response variable we are interested in is the absolute value of the transaction price change between two subsequent transactions of the IBM stock. This variable is abbreviated by PC (absolute *P*rice *C*hange). As the IBM stock is frequently traded we have considerable data available for estimation (usually more than 2000 observations per day), even if we only use the observations from 9:55am until 3:25pm. This way we omit observations from the opening and the closing period which might exhibit a different behaviour. The price changes take on only values which are integer multiples of one sixteenth of a dollar. The categories associated with the absolute price changes are given in Table 1, together with the observed frequencies.

absolute value of price change	PC	frequency
0$	1	863
1/16$	2	919
2/16$	3	252
$\geq$ 3/16$	4	59

Table 1: Absolute price changes: associated categories and observed frequencies.

For our analysis we have the following covariate information available:

- TD, the logarithm of the time elapsed between two subsequent transactions in seconds,
- SIZE, the transaction volume in numbers of shares,
- BOS, the last available bid-offer-spread,
- DIR, the sign of the price change (DIR=0 if PC=0).

We will further examine the influence of lagged covariates on PC. For this purpose we introduce the notation SIZE$_{-1}$,SIZE$_{-2}$,... for the transaction volume one, two, ... transactions before. An analogous notation will be used for the other covariates.
Since the response PC, which corresponds to Y_t in the model definition, is discrete and takes on only four values, ordinary scatter plots will not show us if there is a linear dependency between the latent response variable Y_t^* and the covariates (assuming nearly equidistant cutpints). Therefore we categorize the covariates as shown in Table 2. Now for each covariate we can compute the average of the absolute price changes (PC) per category and then search for linear, quadratic, logarithmic, or other nonlinear relationships between the covariate and PC. We emphasize that for the statistical analysis presented in Section 2.2 we use scored values $0.2, 0.6, \ldots, 2.6$ of TD instead of categories $1, \ldots, 7$.

logarithm of elapsed time	TD	frequency
≤ 0.4	1	89
> 0.4 and ≤ 0.8	2	161
> 0.8 and ≤ 1.2	3	224
> 1.2 and ≤ 1.6	4	194
> 1.6 and ≤ 2.0	5	440
> 2.0 and ≤ 2.4	6	391
> 2.4	7	594

transaction volume	SIZE	frequency
≤ 500	1	669
> 500 and ≤ 1000	2	351
> 1000	3	1073

last bid-offer-spread	BOS	frequency
0.0625	0.0625	261
0.1250	0.1250	679
0.1875	0.1875	562
0.2500	0.2500	545
0.3125	0.3125	46

price change	DIR	frequency
< 0	-1	603
0	0	863
> 0	1	627

Table 2: Associated categories and observed frequencies. Note that the covariate BOS still contains the original values.

2.1 Exploratory analysis of the IBM data

We present now an exploratory analysis investigating the impact of the different covariates on the absolute price changes PC using the procedure outlined above.

2.1.1 Impact of the logarithmic time differences (TD) on the absolute price changes

From the first row in Figure 1 we see that the average response PC clearly increases from 1.39 for transactions with TD=1 to 1.80 for transactions with TD=7. We also see that at least for TD between 1 and 5 the dependency seems to be approximately linear. However, between TD=5 to TD=7 we even see that the response decreases slightly from

1.86 to 1.8. Nevertheless we will continue to use TD as covariate.
There is no hint for a significant influence of TD_{-1} on PC: The average response fluctuates without any trend between 1.68 and 1.86 (cf. first row in Figure 1). A similar result is obtained from the examination of TD_{-2} and TD_{-3} (cf. Figure 1).

2.1.2 Impact of the last bid-offer-spread (BOS) on the absolute price change

There seems to be a significant dependency between BOS and PC (cf. second row in Figure 1). From the plot we see that the dependency is almost linear, even if the slope seems to increase slightly with increasing BOS. The average response increases from 1.48 for BOS=0.0625 to 2.39 for BOS=0.3125. We have a similar result for BOS_{-1}. Here the differnce between maximal and minimal average response is $2.23 - 1.52 = 0.71$ (cf. Figure 1).
Even BOS_{-2} still seems to have a strong positive relationship with PC (cf. Figure 1). The average PC increases depending on the value of BOS_{-2} from 1.6 to 2.28. But the dependency is not linear, which is evident if we consider the average response for BOS_{-2}=0.3125 that is higher than we would expect in the case of linearity. For small values of BOS_{-3} the average of PC is 1.63, for big ones it is 2.06.
For all covariates BOS, BOS_{-1}, ... the curve in Figure 1 is convex. Thus, we now consider squared covariates BOS^2, BOS^2_{-1}, ... Then we have in all four cases a curve indicating a linear dependency (cf. Figure 2). So from now on, we will focus on this transformation of BOS as a covariate.

2.1.3 Impact of the transaction volume (SIZE) on the absolute price change

The corresponding row in Figure 1 indicates a positive dependency of PC and SIZE. The curve is approximately linear. With increasing SIZE the average of the corresponding price changes also increases from 1.647 to 1.846. However, compared to the results we presented for the covariates TD and BOS, this increase is not as high. For $SIZE_{-1}$ and $SIZE_{-3}$ we see again a positive linear dependency. But the differences between the averages of the response are relatively small (for $SIZE_{-1}$: $1.72 \rightarrow 1.79$, $SIZE_{-3}$: $1.69 \rightarrow 1.81$). For $SIZE_{-2}$ there is no linear dependency observable.

2.1.4 Impact of sign of the price change (DIR) on the absolute price change

We do not take DIR itself into consideration as a covariate, because the price direction contains deterministic information about the absolute price change, since DIR=0 $\Longrightarrow$

PC=1, which also would cause an inconsistency with the model assumption of a normally distributed underlying variable Y_t^*. From the last row in Figure 1 we see that the average values of PC are nearly equal, regardless whether DIR = -1 or DIR = 1 and therefore no influence of DIR on the absolute price difference can be expected.
From the plots corresponding to DIR_{-1} and DIR_{-2} one can see that there is no linear relationship between PC and these lagged covariates. However, one could consider for example the absolute value of DIR_{-1} as a covariate to model the fact that PC $\neq 1$ in $t-1$ makes it more likely that the price also changes in t. But this effect is already incorporated in the model through the AR(1)-component, and so we do not examine models with this covariate.

2.1.5 Impact of interactions of the covariates on the absolute price change

The plots illustrating interaction effects can be found in Figures 3 and 4. For the examination of the interactions we proceeded in the following way: First we multiplied the categorized covariates, which we described in Table 2, with each other. Then we have again grouped the created covariates into categories such that we can now conduct the analysis just as we have done it in the sections above. The interactions of TD with BOS, TD with SIZE and BOS with SIZE all show an approximately linear relationship in the plot. As TD$*$BOS increases the average value of PC increases from 1.44 to 2.44, while it changes from 1.55 to 1.97 as TD$*$SIZE increases. Finally the average value of PC grows from 1.56 to 2.4 as BOS$*$SIZE increases.
The relationships between PC and interactions with DIR are again quite symmetric and are no used in the following for the same reasons as outlined in Section 2.1.4.
Figure 4 shows the observed relationships for the twoway interacting terms involving the transformed variable BOS^2. Here the curves seem to be even closer to a straight line than for BOS, showing again the usefulness of this transformation.

2.2 Estimations with the GM-MGMC Gibbs-Sampler

For the estimation of the parameters we use the Grouped Move Multigrid Monte Carlo (GM-MGMC) algorithm proposed by Müller and Czado (2002) for models of the form (1.1) and (1.2). Here and in the following sections the regression part $\mathbf{X}_t'\boldsymbol{\beta}$ always contains the intercept. The corresponding regression coefficient is denoted by β_0. Finally, ϕ denotes the autoregressive parameter. We applied the GM-MGMC Gibbs-sampler for 10000 iterations and considered a burn-in of 2000 iterations as sufficient after examining the time sequence plots.

2.2.1 Model with covariates TD, SIZE, BOS, and BOS^2

As a first model, we investigate the model with TD, SIZE and BOS as covariates. Posterior mean estimates, 95%- and 90%- credibility intervals for the parameters are given in Table 3. For comparison we also give the values where BOS^2 is used in addition

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)	ϕ
mean	1.3592	2.3368	-0.9615	0.2061	0.0803	4.2005	0.0629
2.5%	1.2857	2.2100	-1.1550	0.1348	0.0241	3.4238	0.0099
5.0%	1.2971	2.2285	-1.1255	0.1470	0.0332	3.5459	0.0186
95.0%	1.4241	2.4501	-0.7940	0.2642	0.1286	4.8654	0.1070
97.5%	1.4377	2.4723	-0.7627	0.2762	0.1366	4.9802	0.1147

Table 3: Posterior mean estimates and estimated posterior quantiles for the model with TD, SIZE and BOS.

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)	β_4(BOS^2)	ϕ
mean	1.3631	2.3351	-0.6728	0.2068	0.0778	0.3526	11.1189	0.0612
2.5%	1.2871	2.2059	-1.0182	0.1369	0.0208	-3.5720	0.0446	0.0088
5.0%	1.2986	2.2253	-0.9644	0.1481	0.0296	-2.9715	1.5378	0.0176
95.0%	1.4271	2.4471	-0.3825	0.2657	0.1255	3.6562	20.4890	0.1059
97.5%	1.4394	2.4666	-0.3273	0.2768	0.1338	4.2742	22.3493	0.1153

Table 4: Posterior mean estimates and estimated posterior quantiles for the model with TD, SIZE, BOS, and BOS^2.

(cf. Table 4). We see that in both cases the covariates and the autoregressive parameter are significantly different from zero. In the second model BOS remains still in the model since its square is significant. In the following we refer to the second model as the basic model.

To demonstrate that the GM-MGMC algorithm has converged we present in Figure 5 the sampled MCMC iterations together with the sample autocorrelations in Figure 6. They show a satisfactory behavior of the algorithm. We checked that all other models presented below behave similarly satisfactory and omit therefore the corresponding time sequence and autocorrelation plots.

We will now investigate whether interactions and lagged covariates will improve the fit.

2.2.2 Models with interactions

We consider three models that allow for interactions. In each of these models we use the covariates TD, SIZE, BOS, and BOS^2. In the first model we add the interaction of BOS resp. BOS^2 and SIZE, in the second model the interaction of BOS resp. BOS^2

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)
mean	1.3650	2.3596	-0.3386	0.0184	-0.2277	0.4281
2.5%	1.2911	2.2324	-0.7175	-0.0955	-0.3847	-3.5600
5.0%	1.3018	2.2516	-0.6581	-0.0773	-0.3612	-2.9027
95.0%	1.4287	2.4691	-0.0171	0.1123	-0.0963	3.7343
97.5%	1.4391	2.4893	0.0396	0.1326	-0.0712	4.3302

	β_4(BOS^2)	β_5(TD$*$SIZE)	ϕ
mean	10.8791	0.1658	0.0643
2.5%	-0.4034	0.0874	0.0100
5.0%	1.4382	0.0996	0.0196
95.0%	20.2412	0.2322	0.1092
97.5%	22.1794	0.2456	0.1172

Table 5: Posterior mean estimates and estimated posterior quantiles for the model with interaction of TD and SIZE.

and TD, and in the third model the interaction of TD and SIZE. The quantile estimates show that in the first two models the interactions are not significant on the 90% level (not shown here). The parameter and quantile estimates for the third model are given in Table 5. Here the interaction of TD and SIZE is significant. The fact, that the credibility intervals for the coefficients corresponding to TD and BOS contain zero, does not play any role, since TD is part of the interaction and BOS^2 is significant on the 90% level, too.

Therefore we conclude that if we want to use a model with interactions we use the third one, i.e. the model with covariates TD, SIZE, BOS, BOS^2, and with the interaction of TD and SIZE.

2.2.3 Models with lagged covariates

The explorative analysis in Section 2.1 indicates that the lagged variable TD has no significant influence on PC. This explorative result is confirmed by examining the quantile estimates for the model with the covariates TD, SIZE, BOS, BOS^2, and TD_{-1}: The 90% credibility interval for TD_{-1}, $[-0.1007, 0.0178]$, contains zero. So we concentrate on the

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)
mean	1.3675	2.3373	-0.7191	0.2096	0.0743	0.3390
2.5%	1.2923	2.2098	-1.0733	0.1377	0.0162	-3.5473
5.0%	1.3034	2.2293	-1.0189	0.1485	0.0260	-3.0336
95.0%	1.4340	2.4481	-0.4185	0.2695	0.1215	3.6859
97.5%	1.4468	2.4703	-0.3645	0.2816	0.1311	4.3033

	$\beta_4(\text{BOS}^2)$	$\beta_5(\text{SIZE}_{-3})$	ϕ
mean	10.9702	0.0456	0.0599
2.5%	-0.3092	-0.0109	0.0071
5.0%	1.4785	0.0007	0.0157
95.0%	20.4869	0.0920	0.1043
97.5%	22.1121	0.1009	0.1132

Table 6: Posterior mean estimates and estimated posterior quantiles for the basic model together with SIZE_{-3}.

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)
mean	1.3624	2.3528	-0.7865	0.2037	0.0810	-0.1659
2.5%	1.2844	2.2187	-1.1452	0.1334	0.0240	-4.1364
5.0%	1.2955	2.2398	-1.0861	0.1444	0.0324	-3.4634
95.0%	1.4312	2.4677	-0.4858	0.2645	0.1296	3.1457
97.5%	1.4438	2.4912	-0.4248	0.2764	0.1387	3.7471

	$\beta_4(\text{BOS}^2)$	$\beta_5(\text{BOS}_{-1})$	ϕ
mean	10.3964	1.3617	0.0439
2.5%	-0.8205	0.4053	-0.0118
5.0%	0.9833	0.5358	-0.0027
95.0%	19.7331	2.1930	0.0898
97.5%	21.6042	2.3427	0.0986

Table 7: Posterior mean estimates and estimated posterior quantiles for the basic model together with BOS_{-1}.

examination of the lagged covariates SIZE, BOS, and BOS^2. First we add SIZE_{-1} to the basic model with covariates TD, SIZE, BOS, BOS^2. The quantile estimates however show that this lagged covariate is not significant (90% credibility interval for SIZE_{-1}: $[-0.0590, 0.0373]$). The same holds for SIZE_{-2} (90% credibility interval for SIZE_{-2}: $[-0.0358, 0.0598]$). However, SIZE with lag 3 is significant on the 90% level, but not on the 95% level. The parameter estimates are given in Table 6 together with quantile estimates. We then added BOS_{-1} to the basic model. This lagged covariate is also significant, as Table 7 shows. However, this is the first of the considered models where

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)
mean	1.3727	2.3482	-0.3903	0.0221	-0.2284	0.4785
2.5%	1.3000	2.2192	-0.7738	-0.0922	-0.3837	-3.3538
5.0%	1.3119	2.2404	-0.7122	-0.0745	-0.3590	-2.7660
95.0%	1.4371	2.4585	-0.0702	0.1186	-0.0966	3.7637
97.5%	1.4505	2.4815	-0.0128	0.1380	-0.0702	4.4162

	β_4(BOS^2)	β_5(SIZE$_{-3}$)	β_6(TD$*$SIZE)	ϕ
mean	10.5530	0.0442	0.1645	0.0624
2.5%	-0.6003	-0.0118	0.0846	0.0110
5.0%	1.2464	-0.0035	0.0972	0.0194
95.0%	19.8143	0.0921	0.2316	0.1067
97.5%	21.4507	0.1013	0.2432	0.1143

Table 8: Posterior mean estimates and estimated posterior quantiles for the basic model together with interaction TD$*$SIZE and lagged covariate SIZE$_{-3}$.

the credibility intervals for the autoregressive parameter contain zero and therefore the autoregressive component must be viewed as non-significant.
Further computations showed that neither BOS$^2_{-1}$ nor BOS$_{-2}$, BOS$^2_{-2}$, BOS$_{-3}$, BOS$^2_{-3}$ have significant impact on the absolute price changes.

2.2.4 Models with interactions and lagged covariates

We consider the basic model together with the most significant interaction, which was that one of TD and SIZE, and together with the most significant lagged covariates, SIZE$_{-3}$ and BOS$_{-1}$. The results are presented in Tables 8 and 9. We see that the lagged covariate SIZE$_{-3}$ becomes non-significant when the interaction is included. However, when BOS$_{-1}$ is added to the model, all covariates and the interaction of TD and SIZE remain in the model. Even the autoregressive component remains significant on the 90% level.

2.2.5 Interpretation and summary

Summarizing the results from the previous sections we prefer the model with the covariates TD, SIZE, BOS, BOS^2, BOS$_{-1}$ and the interaction of TD and SIZE among all other considered models. It is the model that takes into account the greatest number of covariates under the restriction that the credibility intervals for all the regression coefficients (except those that are only used for the hierarchical structure) do not contain

	α_2	α_3	β_0	β_1(TD)	β_2(SIZE)	β_3(BOS)
mean	1.3717	2.3606	-0.4517	0.0121	-0.2285	-0.0807
2.5%	1.2957	2.2341	-0.8458	-0.1009	-0.3821	-4.0492
5.0%	1.3085	2.2530	-0.7846	-0.0830	-0.3576	-3.4288
95.0%	1.4372	2.4768	-0.1268	0.1088	-0.0963	3.2174
97.5%	1.4502	2.4988	-0.0666	0.1286	-0.0739	3.8201

	β_4(BOS^2)	β_5(BOS$_{-1}$)	β_6(TD∗SIZE)	ϕ
mean	10.0123	1.4221	0.1679	0.0454
2.5%	-1.0789	0.4502	0.0885	-0.0074
5.0%	0.7265	0.6084	0.1005	0.0001
95.0%	19.4674	2.2384	0.2334	0.0911
97.5%	21.3088	2.4115	0.2471	0.0994

Table 9: Posterior mean estimates and estimated posterior quantiles for the basic model together with interaction TD∗SIZE and lagged covariate BOS$_{-1}$.

zero.

We now consider the posterior mean estimates of the model with the covariates TD, SIZE, BOS, BOS^2, BOS$_{-1}$ and the interaction of TD and SIZE in detail, and refer in the following to Table 9.

The impact of the covariates TD and SIZE on the fitted mean of the latent variables Y_t^* is given in Table 10 and in Figure 7. For example, a (logarithm of) TD of about 0.2 and a SIZE of 1 (500 or less stocks) leads to a estimated decrease of $0.0121 \cdot 0.2 - 0.2285 \cdot 1 + 0.1679 \cdot 0.2 \cdot 1 = -0.1925$ of the mean of Y_t^*, which is the top left value in the table. From the table we conclude that for fixed SIZE a longer TD

SIZE	TD (scored)						
	0.2	0.6	1.0	1.4	1.8	2.2	2.6
1	-0.1925	-0.1205	-0.0485	0.0235	0.0955	0.1675	0.2395
2	-0.3874	-0.2483	-0.1091	0.0301	0.1692	0.3084	0.4475
3	-0.5823	-0.3760	-0.1697	0.0366	0.2429	0.4493	0.6556

Table 10: Impact of TD and SIZE on the fitted mean of the latent variables Y_t^*.

increases the probability for a higher price jump. Furthermore we can see that the more stocks are traded the more dramatical the impact of TD is: For SIZE 1 the difference between the maximum and the minimum is $0.2395 + 0.1925 = 0.432$, whereas for SIZE 3 this difference is $0.6556 + 0.5823 = 1.2379$, nearly three times as high.

The impact of SIZE, given a fixed TD, depends on the value of TD. For example, when TD $= 0.2$, we can see: The more stocks are traded, the lower is the probability for a

high price jump. When TD = 2.6, we can see: The more stocks are traded, the bigger is the probability for a high price jump.
The impacts of the last bid-offer-spread BOS and the lagged covariate BOS_{-1} on the fitted mean of the latent variables Y_t^* are given in Table 11. The bigger the last bid-

bid-offer-spread (BOS)	0.0625	0.1250	0.1875	0.2500	0.3125
impact on fitted mean of Y_t^*	0.0341	0.1463	0.3369	0.6056	0.9525

lagged spread (BOS_{-1})	0.0625	0.1250	0.1875	0.2500	0.3125
impact on fitted mean of Y_t^*	0.0889	0.1778	0.2666	0.3555	0.4444

Table 11: Impact of BOS and BOS_{-1} on the fitted mean of the latent variables Y_t^*.

offer-spread the higher the probability for a high price jump. The same holds for the lagged spread BOS_{-1}.
Comparing the magnitude of the effects on the fitted means of Y_t^* for the different covariates and the interaction, we see that BOS has the highest maximal effect, while the maximal effect of BOS_{-1} is only half that of BOS. The effect of the interaction between TD and SIZE is considerable compared to BOS and BOS_{-1}.

2.3 Predictions

In the following only the model with covariates TD, SIZE, BOS, BOS^2, BOS_{-1} and the interaction of TD and SIZE is considered. We now try to predict the absolute price changes at time 1501 to 2000 using posterior mean estimates of all parameters based only on the first 1500 observations of the data set. These posterior mean estimates are quite similar to that in Table 9 and are therefore not shown explicitly.
We now show how to calculate predictions where we use ideas contained in Müller and Czado (2002).
Let $t_{max} = 1500$. Starting from the posterior mean estimates $\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\phi}$ and the posterior mean estimate $\overline{Y}^*_{t_{max}}$ for $Y^*_{t_{max}}$, we compute recursively the estimates $\hat{Y}^*_{t_{max}+n}$ for $Y^*_{t_{max}+n}$, $n = 1, ..., 500$, as follows:

$$\begin{aligned} \hat{Y}^*_{t_{max}+1} &:= \mathbf{X}'_{t_{max}+1}\overline{\boldsymbol{\beta}} + \overline{\phi}\,\overline{Y}^*_{t_{max}} \\ \hat{Y}^*_{t_{max}+n} &:= \mathbf{X}'_{t_{max}+n}\overline{\boldsymbol{\beta}} + \overline{\phi}\,\hat{Y}^*_{t_{max}+n-1}, \quad n \geq 2. \end{aligned}$$

true category	predicted category			
	1	2	3	4
1	102	96	0	0
2	95	142	0	0
3	7	49	0	0
4	0	9	0	0
SSE	304			

Table 12: Frequencies of predicted categories classified by true categories for the model with covariates TD, SIZE, BOS, BOS^2, BOS$_{-1}$ and TD$*$SIZE.

Then we calculate recursively the following probability estimates for $k = 1, \ldots, 4$ (set $\overline{\alpha}_0 = -\infty$ and $\overline{\alpha}_4 = \infty$):

$$
\begin{aligned}
\hat{P}(Y_{t_{max}+1} = k \mid \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\phi}, \overline{Y}^*_{t_{max}}) &:= \Phi(\overline{\alpha}_k - \mathbf{X}'_{t_{max}+1}\overline{\boldsymbol{\beta}} - \overline{\phi}\,\overline{Y}^*_{t_{max}}) - \\
&\quad -\Phi(\overline{\alpha}_{k-1} - \mathbf{X}'_{t_{max}+1}\overline{\boldsymbol{\beta}} - \overline{\phi}\,\overline{Y}^*_{t_{max}}) \\
\hat{P}(Y_{t_{max}+n} = k \mid \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\phi}, \hat{Y}^*_{t_{max}+n-1}) &:= \Phi(\overline{\alpha}_k - \mathbf{X}'_{t_{max}+n}\overline{\boldsymbol{\beta}} - \overline{\phi}\,\hat{Y}^*_{t_{max}+n-1}) - \\
&\quad -\Phi(\overline{\alpha}_{k-1} - \mathbf{X}'_{t_{max}+n}\overline{\boldsymbol{\beta}} - \overline{\phi}\,\hat{Y}^*_{t_{max}+n-1}), \quad n \geq 2.
\end{aligned}
$$

This leads to the predictions

$$\hat{Y}_{t_{max}+n} := \text{argmax}_{k=1,\ldots,4}\hat{P}(Y_{t_{max}+n} = k \mid \ldots), \quad n = 1, \ldots, 500.$$

Table 12 shows how often each category was predicted, classified by the true categories. Obviously only categories 1 and 2 are predicted. This may be due to the fact that the categories 1 and 2 occur much more than categories 3 and 4 (cf. Table 1). Nevertheless we can see: The higher the true category is, the more likely it is, that category 2 is predicted, and when the true category is 4, we never predict category 1. The sum of squared errors, SSE=$\sum_{n=1}^{500}(\hat{Y}_{1500+n} - Y_{1500+n})^2$, is 304 which might be satisfying since 500 values were predicted.

3 Sleep Data

The data set considered here contains 1024 sleep state measurements of a newborn child together with its heart rate and temperature sampled every 30 seconds. This data set was also discussed in Kedem and Fokianos (2002). The sleep states are classified as shown in Table 13. The temperature was measured with an accuracy of 0.05 degrees Centigrade. Figure 8 shows the progression of sleep state, heart rate, and temperature, while Figure 9 provides histograms for these. The frequencies of the different sleep

sleep state	depth of sleep	category
quiet sleep	low	1
indeterminate sleep	medium	2
active sleep	high	3
awake	-	4

Table 13: Classification of sleep states.

Sleep State	Frequency
1	404
2	94
3	237
4	289

Heart Rate	Frequency
100-110	13
111-120	160
121-130	368
131-140	181
141-150	120
151-160	104
161-160	64
171-180	11
181-190	3

Temperature	Frequency
36.85	5
36.90	80
36.95	123
37.00	44
37.05	177
37.10	80
37.15	137
37.20	96
37.25	45
37.30	123
37.35	86
37.40	20
37.45	8

Table 14: Frequencies of sleep states, heart rates and temperatures.

states, heart rates, and temperatures measured are listed in Table 14. Table 15 shows the empirical quantiles.

3.1 Explorative Data Analysis

In our model we assume a linear influence of the covariates on the latent variables Y_t^*. Because of that, we first want to examine, whether the covariates heart rate and temperature have a linear influence on sleep state (assuming nearly equidistant cutpoints) and, if not, which transformations could be appropriate. Therefore we categorize the covariates and compute the average sleep state for each category. Of course, the more frequencies we use, the less smooth the relationsships between average response and covariate will be and the more likely categories with only few observations will occur. Figure 10 shows the relationships for 9 and 4 categories of heart rate, respectively, and for 13 and 5 categories of temperature, respectively. Tables 16 and 17 show which

Variable	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Sleep State	1	1	3	2.401	4	4
Heart Rate	100	123	130	134.1	145	187
Temperature	36.85	37.05	37.15	37.128	37.25	37.45

Table 15: Table of empirical quantiles, Min., Max. and Mean.

heart rate	100-110	111-120	121-130	131-140	141-150	151-160
9 categories	1	2	3	4	5	6
4 categories	1		2		3	

heart rate	161-170	171-180	181-190
9 categories	7	8	9
4 categories	4		

Table 16: Categories of heart rate when 9 resp. 4 categories are used.

interval corresponds to each category in each case.
Whereas the influence of temperature on sleep state seems to be quadratic (cf. Figure 10), the relationship between sleep state and heart rate is not evident when 9 categories are used. When only 4 categories for heart rate are used the relationship between heart rate level and average sleep state is more linear. We therefore decided to model the influence of heart rate linearly.
The sleep depth ranges from 4 (awake) over 1 (light sleep) to the deepest sleep state 3 (active sleep). Therefore it would be reasonable to rename the categories ($4 \rightarrow 1$, $1 \rightarrow 2$, $2 \rightarrow 3$, $3 \rightarrow 4$) so that 1 means awake and 4 the deepest sleep state. Then a higher category always would indicate a deeper sleep. However, this does not lead to a more evident relationship between sleep state and heart rate (cf. Figure 11). Therefore in Sections 3.2 and 3.3 we will use the original coding of the sleep state given in Table 13, but we will return to the renamed categories in Section 3.4.

temperature	36.85	36.90	36.95	37.00	37.05	37.10	37.15
13 categories	1	2	3	4	5	6	7
5 categories	1			2		3	

temperature	37.20	37.25	37.30	37.35	37.40	37.45
13 categories	8	9	10	11	12	13
5 categories	4		5			

Table 17: Categories of temperature when 13 resp. 5 categories are used.

measured heart rate	100-120	121-140	141-160	161-180
covariate 'heart rate'	20	40	60	80

Table 18: Definition of covariate heart rate.

measured temperature	36.85	36.90	36.95	37.00	37.05	37.10	37.15
covariate 'temp.'	0.025	0.075	0.125	0.175	0.225	0.275	0.325

measured temperature	37.20	37.25	37.30	37.35	37.40	37.45
covariate 'temp.'	0.375	0.425	0.475	0.525	0.575	0.625

Table 19: Definition of covariate temperature.

3.2 Model with covariates heart rate, temperature and (temp.)^2 and AR(1)-component

Because of the results of the explorative data analysis we take heart rate, temperature, and (temperature)^2 as covariates. We use scaled and scored values for heart rate which arise after substracting 90 from the scored value. The covariate temperature will be considered to be metric. Here we take the true temperature minus 36.825 as covariate. These transformations are done to avoid large estimates for the intercept. Please note that an heart rate of more than 180 was measured only three times. These three values are set to 180 to omit a category with only 3 observations. The definitions for the covariates used in the statistical analysis compared to the original measurements are given in Tables 18 and 19.

3.2.1 Estimation of parameters

We use again a Bayesian approach and employ the GM-MGMC Gibbs sampler by Müller and Czado (2002) to get posterior mean estimates and credibility intervals for the cutpoints α_1 and α_2, the intercept β_0, the regression parameters β_1, β_2 and β_3, and the autoregressive parameter ϕ. We use 25000 iterations and a burn-in of 10000 iterations. The time series plots of the Gibbs sampler are displayed in Figure 12. Since we are interested primarily in point estimates such as the mean or specific quantiles higher autocorrelations (cf. Figure 13) are acceptable.
Table 20 contains posterior mean and quantile (2.5%, 5%, 95% and 97.5%) estimates for the cutpoints α_1 and α_2, for the intercept β_0, for the regression parameters β_1 (heart rate), β_2 (temperature), β_3 (temperature^2), and for the autoregressive parameter ϕ. The posterior mean estimate of the autoregressive component ϕ is 0.9907. That means that

	α_1	α_2	β_0	β_1 (H. Rate)	β_2 (Temp.)	β_3 (Temp.^{2})	ϕ
mean	0.9834	3.3243	0.0892	0.0040	-0.9684	0.3052	0.9907
2.5%	0.5598	2.5117	-0.2424	-0.0038	-2.8197	-2.7077	0.9788
5.0%	0.6164	2.6428	-0.1856	-0.0025	-2.5546	-2.2749	0.9810
95.0%	1.4208	4.0745	0.3611	0.0104	0.6112	2.8756	0.9985
97.5%	1.5361	4.2720	0.4125	0.0117	0.8853	3.3082	0.9992

Table 20: Posterior mean and quantile (2.5%, 5%, 95% and 97.5%) estimates for the parameters α_1, α_2, β_0, β_1, β_2, β_3 and ϕ.

the development of the sleep state is explained nearly only by the autoregression in this model, such that the sleep state behaves similar to a random walk. Therefore it is no surprise that for the regression parameters β_0, β_1, β_2, and β_3 the 90% and the 95% credibility intervals include zero and therefore all used covariates have to be considered not to be significant.

This result is due to the relatively long lasting periods of constant sleep states (cf. Figure 8). Because of these periods of constant sleep states the probability that the next sleep state will be the same as the actual one is very high which, of course, explains the estimated high value of the autogressive parameter.

Therefore, in the next section, we will drop the autoregressive component and use the resulting well-known ordered probit model to search for significant covariates. Before that we now investigate shortly how predictions behave in the model with AR(1)-component.

3.2.2 Predictions

We predict the last 274 data points of the data set using posterior mean estimates of all parameters based only on the first 750 observations of the data set. Note that we have $1024 = 750 + 274$ observations in total. The predictions are made in the same manner as in Section 2.3, for more details see Müller and Czado (2002). Table 21 shows the frequencies of predicted and true values. Mostly the sleep state 4 (awake) was predicted, and never the sleep states 1 and 2 were predicted. The true and predicted values are displayed in Figure 14 from which the behaviour of the predictions is evident: Given the 750^{th} observation, which has sleep state 4, we predict this state for a long time since the autoregressive component is nearly 1. Finally, to compare this model to those discussed later we compute the sum of squared errors, which is 1098 in this case, showing that this model is not satisfactory.

true category	predicted category			
	1	2	3	4
1	0	0	29	85
2	0	0	0	35
3	0	0	0	65
4	0	0	12	48
SSE	1098			

Table 21: Frequencies of predicted categories classified by true categories, and the sum of squared errors (SSE).

3.3 Application of the ordered probit model

The previous model includes an autoregressive component. The nature of the data leads to an autoregressive parameter near to 1 whereas covariates do not show any significant influence. Hence we drop now the autoregressive component, which leads to the well-known ordered probit model. The ordered probit model was originally developed by Aitchison and Silvey (1957) and Ashford (1959). We use the same covariates as before, namely heart rate, temperature and (temperature)^2.

3.3.1 Estimation of parameters

For the estimation of the interesting parameters we employ a modified version of the GM-MGMC Gibbs sampler used above. This modification is straight-forward and consists only in dropping the estimation of the autoregressive parameter in each iteration and to use GM steps with the total scale group instead of the partial scale group. Further details concerning GM steps for the ordered probit model can be found in Liu and Sabatti (2000). We use again 25000 iterations with a burn-in of 10000 iterations. The output of the Gibbs sampler is displayed in Figure 15. Because of the easier structure of the ordered probit model the autocorrelations of the produced chains (cf. Figure 16) behave much better than in the case with AR(1) component.
The examination of the significance of the cutpoints α_1 and α_2, of the intercept β_0, of the regression parameters β_1 (Heart Rate), β_2 (Temperature) and β_3 ((Temp.)^2) is again done by considering credibility intervals. From Table 22 one can see that for all parameters the 90% and the 95% credibility intervals do not contain zero. Therefore they can all be considered to be significant, and we expect to get better results in prediction.

	α_1	α_2	β_0	β_1 (Heart Rate)	β_2 (Temp.)	β_3 (Temp.^{2})
mean	0.2470	0.8915	0.5167	0.0078	-6.6223	12.6949
2.5%	0.2026	0.8113	0.1572	0.0030	-9.2101	8.6232
5.0%	0.2092	0.8234	0.2143	0.0038	-8.7822	9.2837
95.0%	0.2874	0.9616	0.8133	0.0117	-4.4708	16.1288
97.5%	0.2960	0.9762	0.8739	0.0125	-4.0478	16.8379

Table 22: Posterior mean and 2.5%, 5%, 95% and 97.5% quantile estimates of the parameters α_1, α_2, β_0, β_1, β_2 and β_3.

3.3.2 Predictions

To predict the last 274 data points of the data set we use again posterior mean estimates of all parameters based only on the first 750 observations of the data set. Table 23 shows the relation of predicted versus true values of sleep state. As for the model with autoregressive component only two categories are predicted (4 and 1). As expected the model without autoregressive component leads to more changes in the sleepstate as the model with autoregressive component. Comparing the predictions with and without autoregressive component (Tables 21, 23 and Figures 14, 17) one can see that the latter model leads to much better predictions, also indicated by the sum of squared errors, which is 622 in this case in contrast to 1098 in the model with AR(1) component.

true category	predicted category			
	1	2	3	4
1	95	0	0	19
2	14	0	0	21
3	36	0	0	29
4	20	0	0	40
SSE	622			

Table 23: Ordered probit model: Frequencies of predicted categories classified by true categories, and the sum of squared errors (SSE) using coding of Table 13.

3.4 Ordered probit model with renamed categories

We now rename the categories in the natural way ($4 \to 1$, $1 \to 2$, $2 \to 3$, $3 \to 4$), so that a higher category always indicates a deeper sleep, in particular that 1 respresents being awake and 4 the deepest sleep (active sleep). Again we apply the ordered probit model and use the same covariates as before.

3.4.1 Estimation of parameters

We use the modified Gibbs sampler with again 25000 iterations and a burn-in of 10000. From Table 24 we see that only for the heart rate with corresponding parameter β_1 the 90% credibility interval does contain zero. The covariates temperature and (temp.)^2 are both significant.

	α_1	α_2	β_0	β_1 (Heart Rate)	β_2 (Temp.)	β_3 (Temp.^{2})
mean	1.0552	1.3311	0.5717	-0.0012	2.2221	-5.4201
2.5%	0.9732	1.2380	0.2240	-0.0059	-0.2924	-9.3975
5.0%	0.9858	1.2522	0.2830	-0.0053	0.1137	-8.7342
95.0%	1.1286	1.4125	0.8621	0.0028	4.3094	-2.0624
97.5%	1.1433	1.4291	0.9193	0.0035	4.6952	-1.3948

Table 24: Renamed sleep categories: Posterior mean and 2.5%, 5%, 95% and 97.5% quantile estimates of the parameters α_1, α_2, β_0, β_1, β_2 and β_3.

3.4.2 Predictions

To predict the last 274 data points of the data set we use again posterior mean estimates of all parameters based on the first 750 observations of the data set. Table 25 shows the relation of predicted versus true values of sleep state. Again only two categories

true category	predicted category			
	1	2	3	4
1	30	30	0	0
2	11	103	0	0
3	18	17	0	0
4	18	47	0	0
SSE	480			

Table 25: Renamed sleep categories: Frequencies of predicted categories classified by true categories, and the sum of squared errors (SSE).

are predicted (1 and 2), which are the same as in Section 3.3 in consideration of the renaming procedure.
The sum of squared errors is 480 in this case, which can be considered as an indicator that this model fits the data better in comparison to the other two models.

Significance of	Model with AR	Ordered Probit	Ord. Probit & Renaming
Intercept	-	+	+
Heart Rate	-	+	-
Temperature	-	+	+
$(\text{Temperature})^2$	-	+	+
AR component	+		
Predictions: SSE	1098	622	480

Table 26: Comparison of the model with AR(1) component, the ordered probit model, and the ordered probit model applied on the process with renamed categories. SSE is the sum of squared errors. + significant, - nonsignificant on 90% level.

true temperature	covariate 'temp.'	influence	
36.90	0.075	$2.2221 \cdot 0.075 - 5.4201 \cdot 0.075^2 =$	0.1362
37.00	0.175	$2.2221 \cdot 0.175 - 5.4201 \cdot 0.175^2 =$	0.2229
37.10	0.275	$2.2221 \cdot 0.275 - 5.4201 \cdot 0.275^2 =$	0.2012
37.20	0.375	$2.2221 \cdot 0.375 - 5.4201 \cdot 0.375^2 =$	0.0711
37.30	0.475	$2.2221 \cdot 0.475 - 5.4201 \cdot 0.475^2 =$	-0.1674
37.40	0.575	$2.2221 \cdot 0.575 - 5.4201 \cdot 0.575^2 =$	-0.5143

Table 27: Renamed sleep categories: Influence of the temperature on the sleep state.

3.5 Interpretation and summary

In Sections 3.2, 3.3, and 3.4 we applied three different models to the sleep data. Table 26 shows, which covariates turned out to be significant and which sum of squared errors we got in predicting the last 274 observations of the data set. The first model was dominated by the autoregressive component. Because of this the prediction was very poor and had a sum of squared errors of 1098. Dropping the autoregressive component led to the ordered probit model, which fits the data much better, indicated by the sum of squared errors of 622. The best fit, however, we achieved by renaming the sleep state categories in a natural way, so that state 1 represents being awake and states 2 to 4 the 3 sleep states where a higher category indicates a deeper sleep. Here the sum of squared errors was only 480. Furthermore in this model the parameters are much better interpretable than in the other two models because of the natural order of the sleep states. Therefore now we consider the posterior mean estimates of this model in detail, and refer in the following to Table 24.

In this model the covariate heart rate is not significant from the viewpoint of the 90% credibility interval which contains zero. Nevertheless the negative sign of the posterior mean estimate -0.0012 for β_1 indicates that a lower heart rate tends to lead to a higher sleep state, i.e. to a deeper sleep.

For interpreting the influence of the temperature on the sleep state we have to take both the linear and the quadratic component into account. In Table 27 we investigate the effect of temperature assuming a quadratic relationship on the fitted mean of the latent variable Y_t^*. From this table we can conclude that for a deep sleep a temperature of about 37.00 degress Centigrade is best (influence = 0.2229), and a lower or higher temperature tends to counteract a deep sleep. This holds especially for temperatures higher than 37.30 degress Centigrade (influence $\leq$ -0.1674). Note the nonsymmetric behavior.

ACKNOWLEDGEMENT

This work was supported by the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 386 *Statistical Analysis of Discrete Structures*. We like to thank Konstantinos Fokianos for providing the sleep data.

References

Aitchison, J. and S. Silvey (1957). The generalization of probit analysis to the case of multiple responeses. *Biometrika 44*, 131–140.

Ashford, J. (1959). An approach to the analysis of data for semi-quantal responses in biological response. *Biometrics 26*, 535–581.

Bauwens, L. and P. Giot (2001). *Econometric Modelling of Stock Market Intraday Activity*. Boston: Kluwer Academic Publishers.

Dacorogna, M. M., R. Gençay, U. A. Müller, R. B. Olsen, and O. V. Pictet (2001). *An Introduction to High Frequency Finance*. San Diego: Academic Press.

Kedem, B. and K. Fokianos (2002). *Regression Models for Time Series Analysis*. New York: Wiley.

Liu, J. S. and C. Sabatti (2000). Generalized Gibbs sampler and multigrid Monte Carlo for Bayesian computation. *Biometrika 87*, 353–369.

Müller, G. and C. Czado (2002). Regression Models for Ordinal Valued Time Series with Application to High Frequency Financial Data. *SFB 386, Discussion Paper 301*.

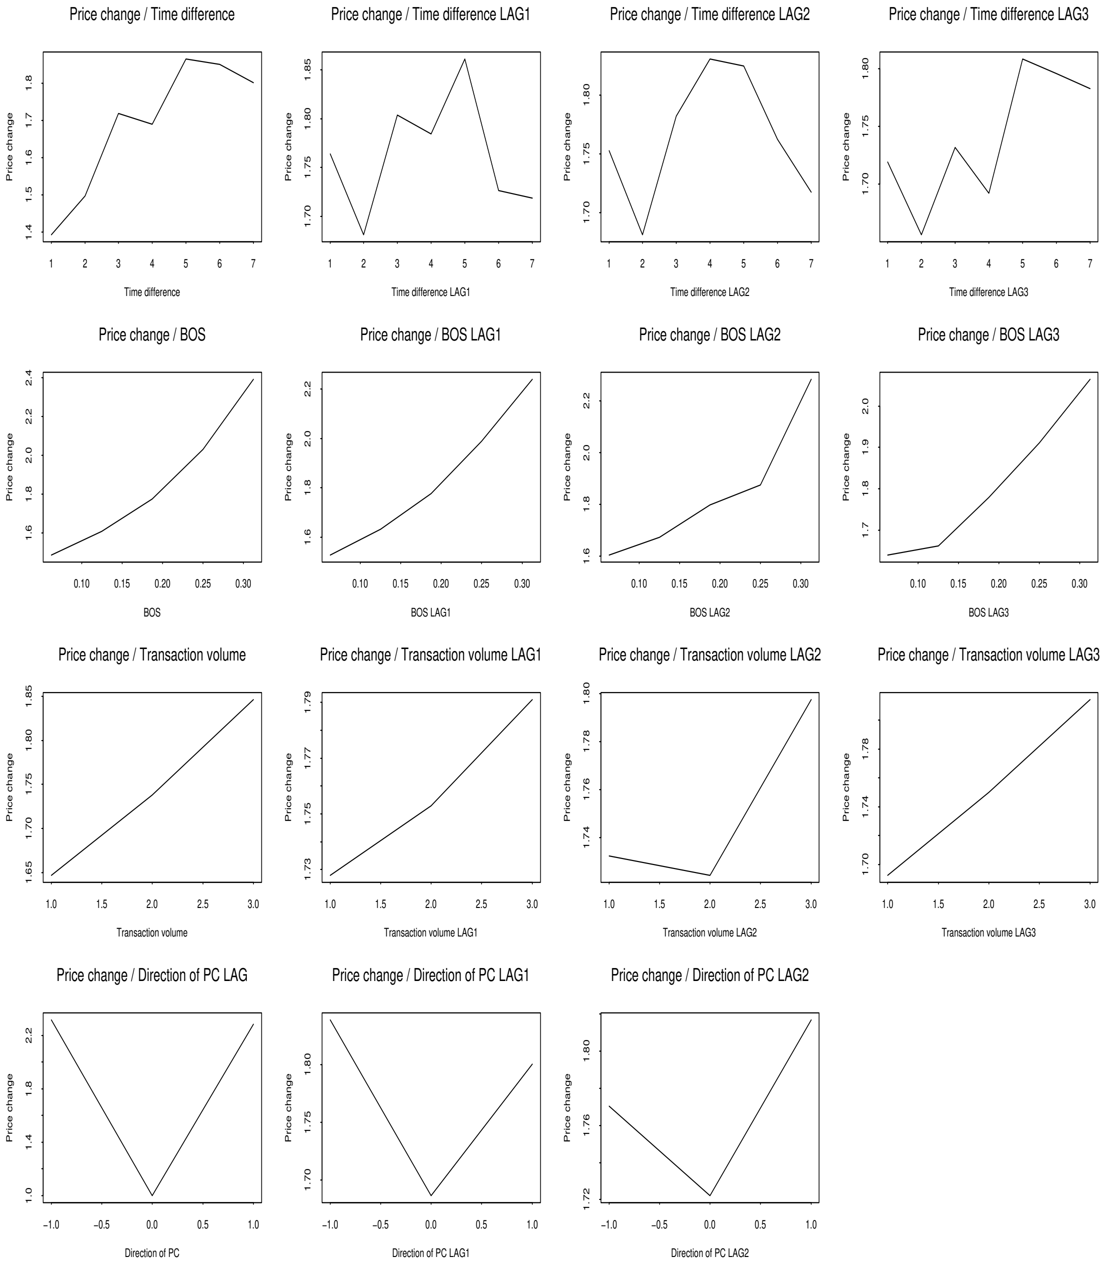

Figure 1: Observed relationships between average PC values per category and different levels of the covariates TD, SIZE, BOS, and DIR.

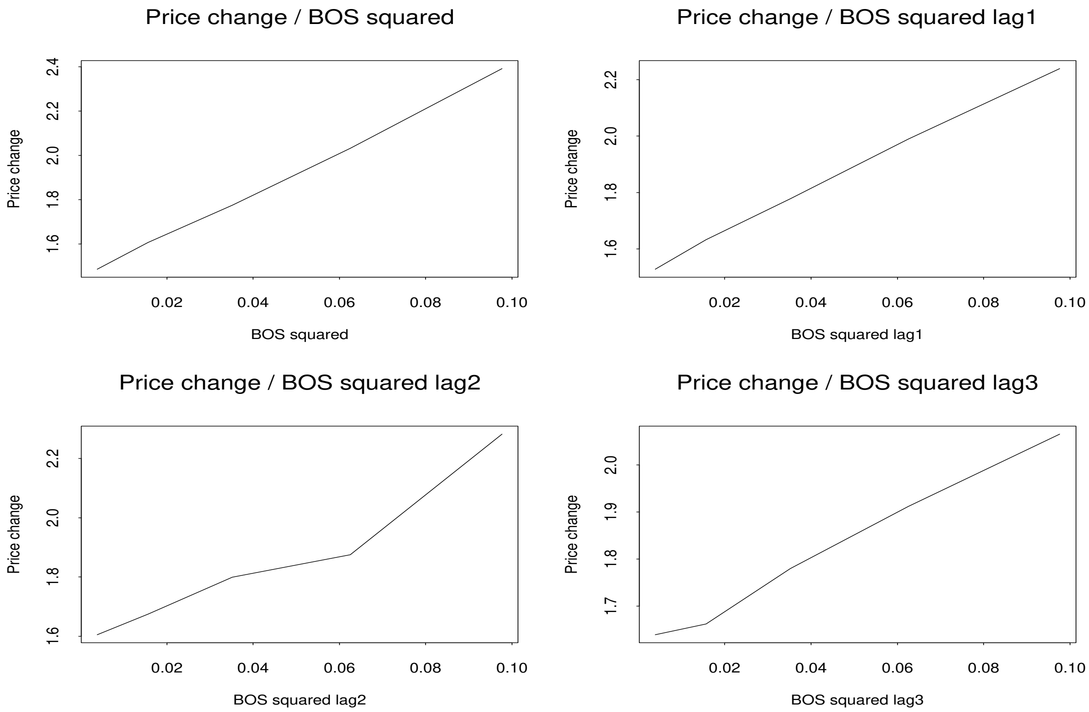

Figure 2: Observed relationships between average PC values per category and different levels of the covariate BOS^2.

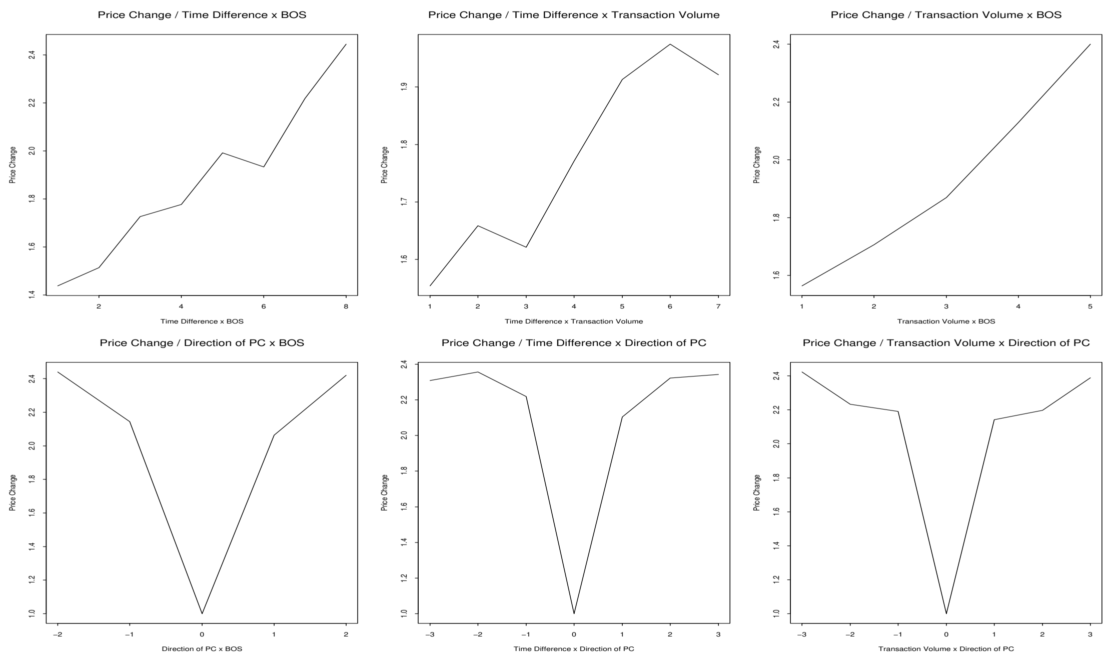

Figure 3: Observed relationships between average PC values and all twoway interaction terms formed by the covariates TD, BOS, SIZE, and DIR.

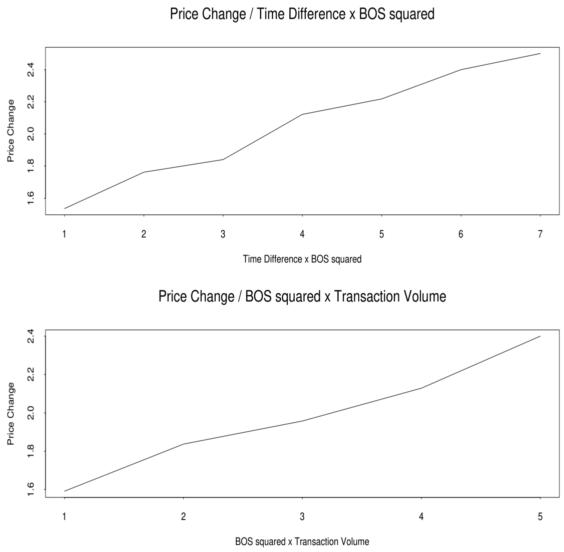

Figure 4: Relationship between Price Change and interactions with BOS^2.

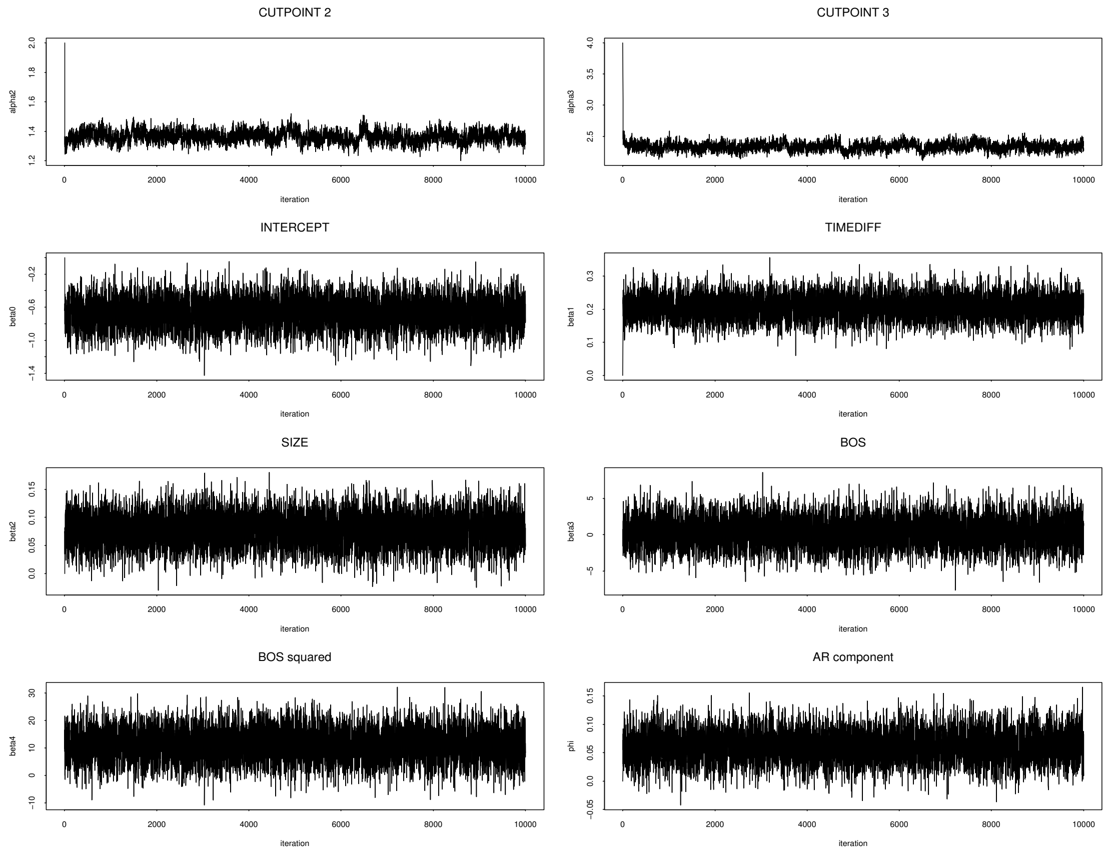

Figure 5: Output produced by the GM-MGMC Gibbs Sampler with covariates TD, SIZE, BOS, and BOS^2.

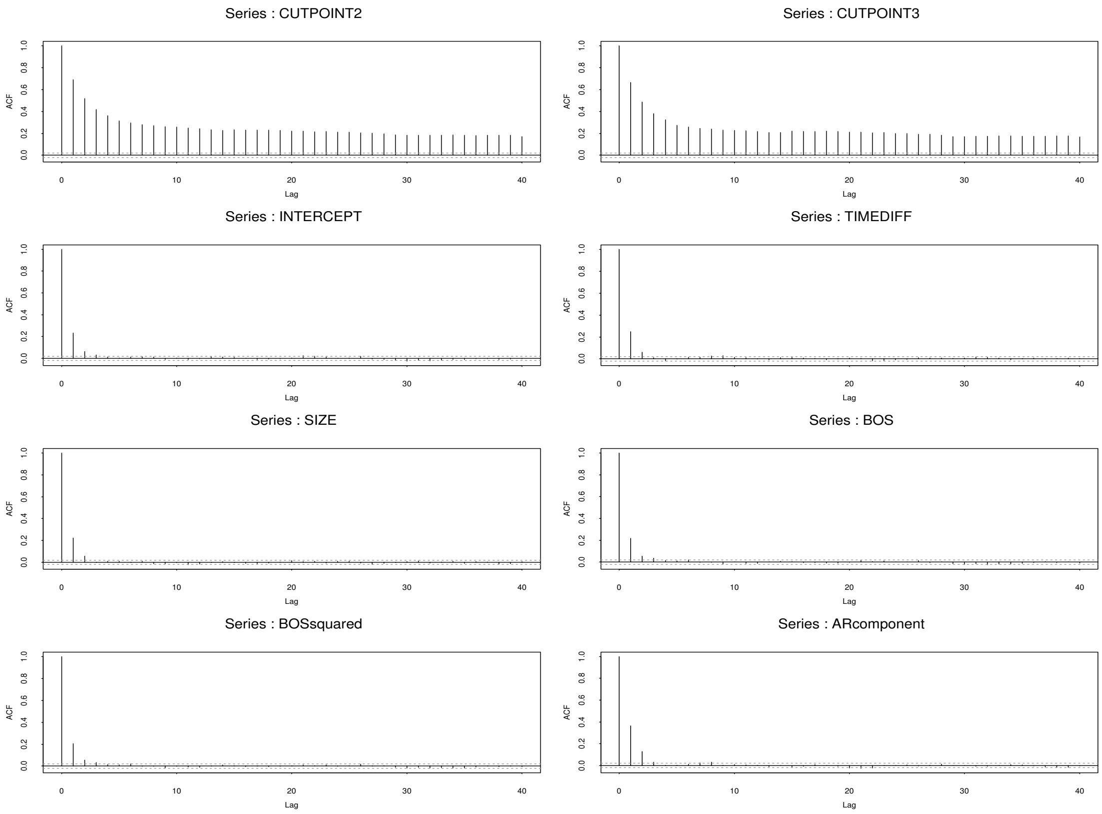

Figure 6: Sample autocorrelations for the time series in Figure 5.

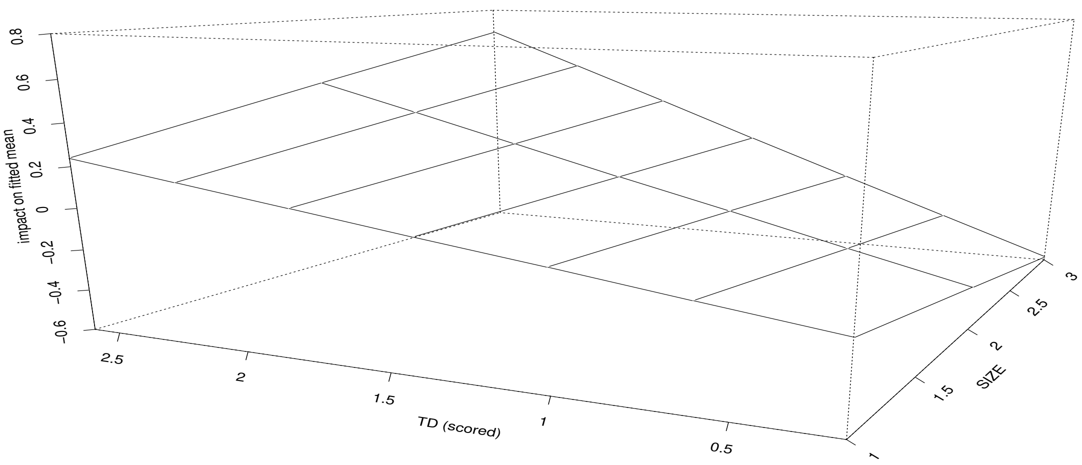

Figure 7: Impact of covariates TD and SIZE on fitted mean of the latent variables Y_t^*.

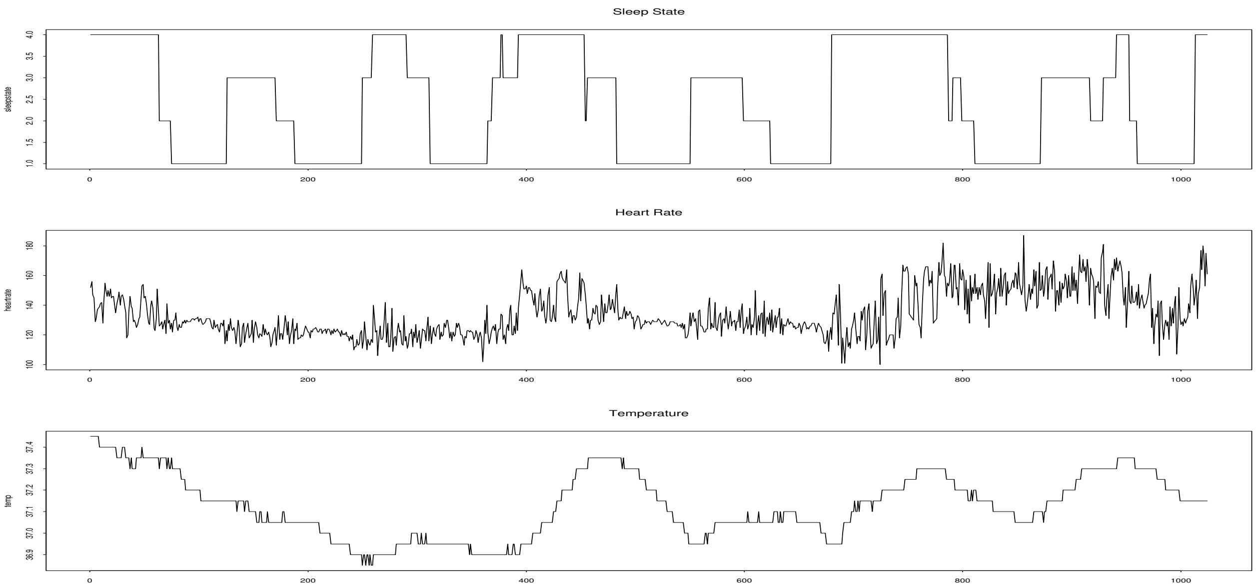

Figure 8: Progression of Sleep State, Heart Rate and Temperature.

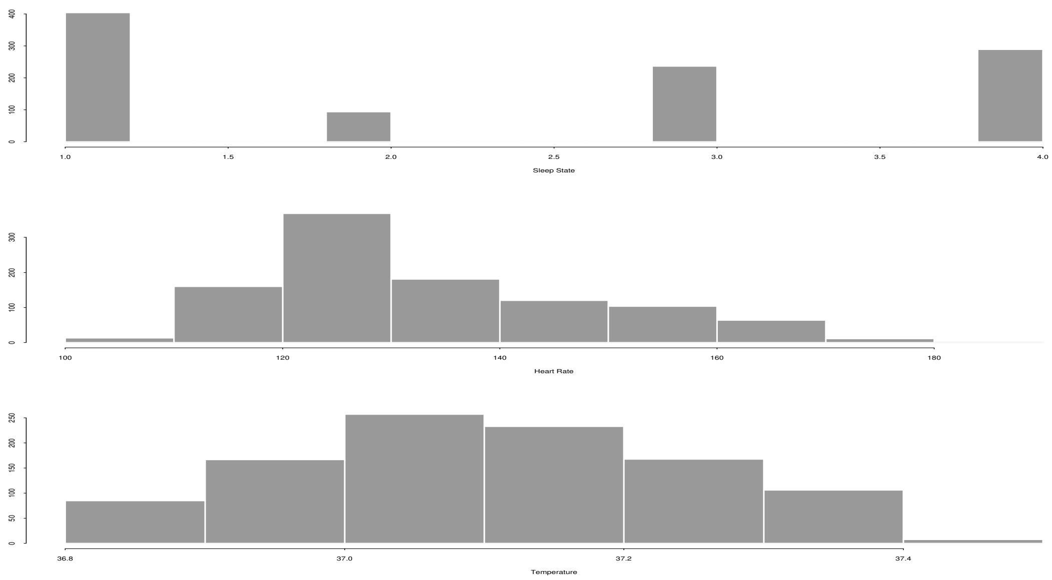

Figure 9: Histograms of the response Sleep State and of the covariates Heart Rate and Temperature.

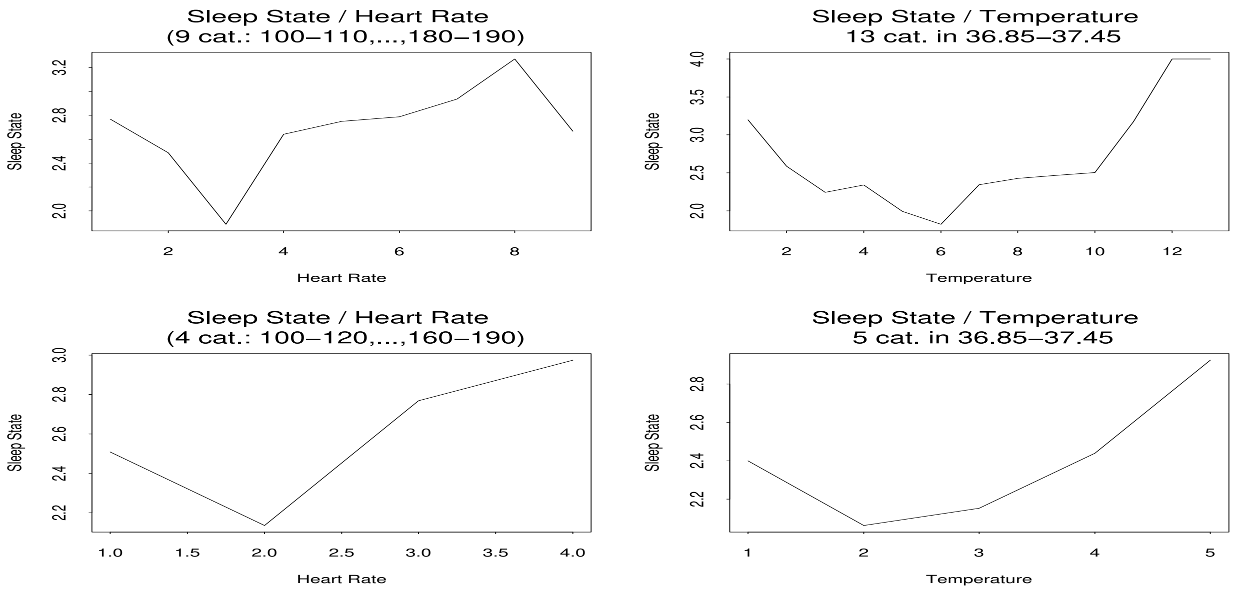

Figure 10: Observed relationships between average Sleep State values per category and different levels of the covariates Heart Rate and Temperature.

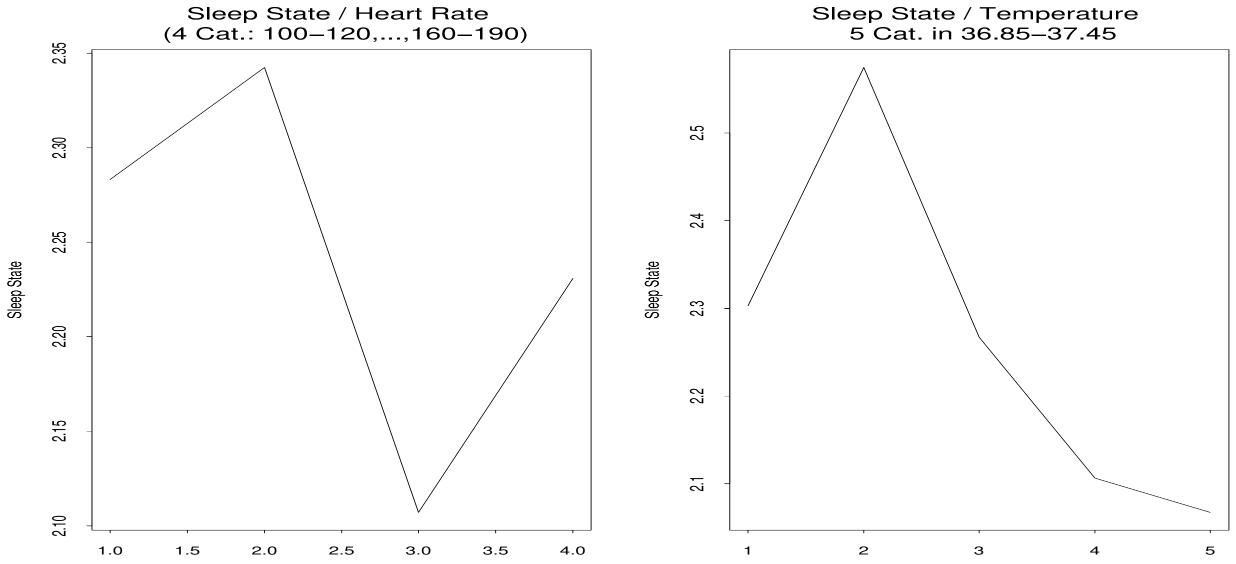

Figure 11: Influence of Heart Rate and Temperature on Sleep State using the renamed categories $4 \to 1$, $1 \to 2$, $2 \to 3$, $3 \to 4$.

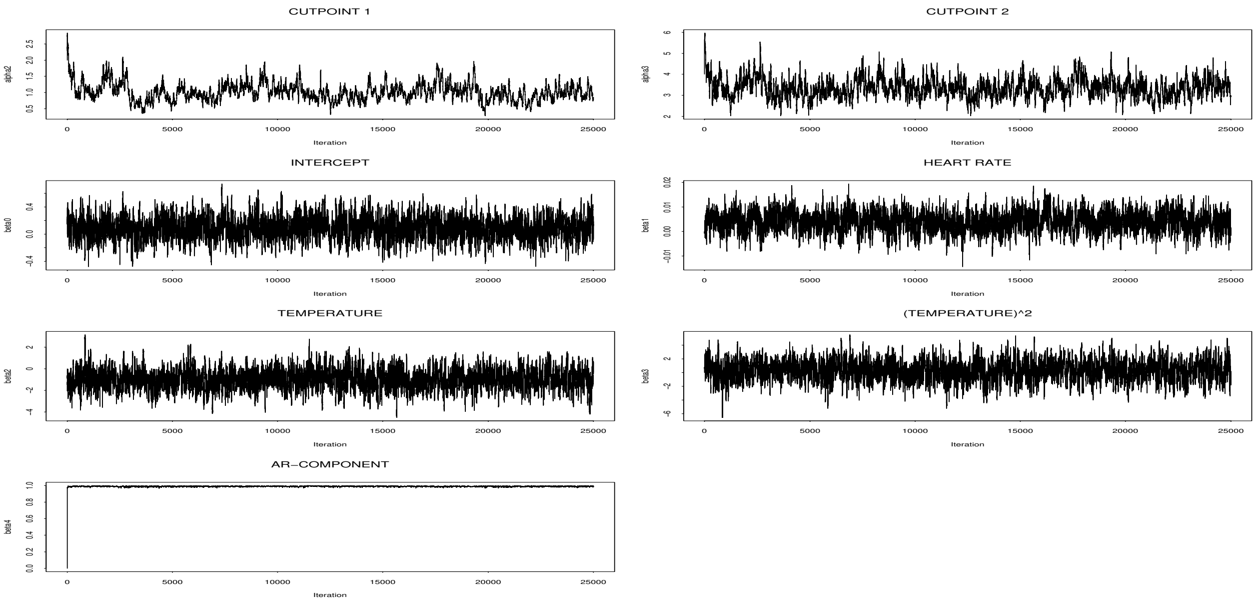

Figure 12: 25000 iterations of Gibbs-Sampler for approach with covariates Heart Rate, Temperature, $(\text{Temperature})^2$ and autoregressive component.

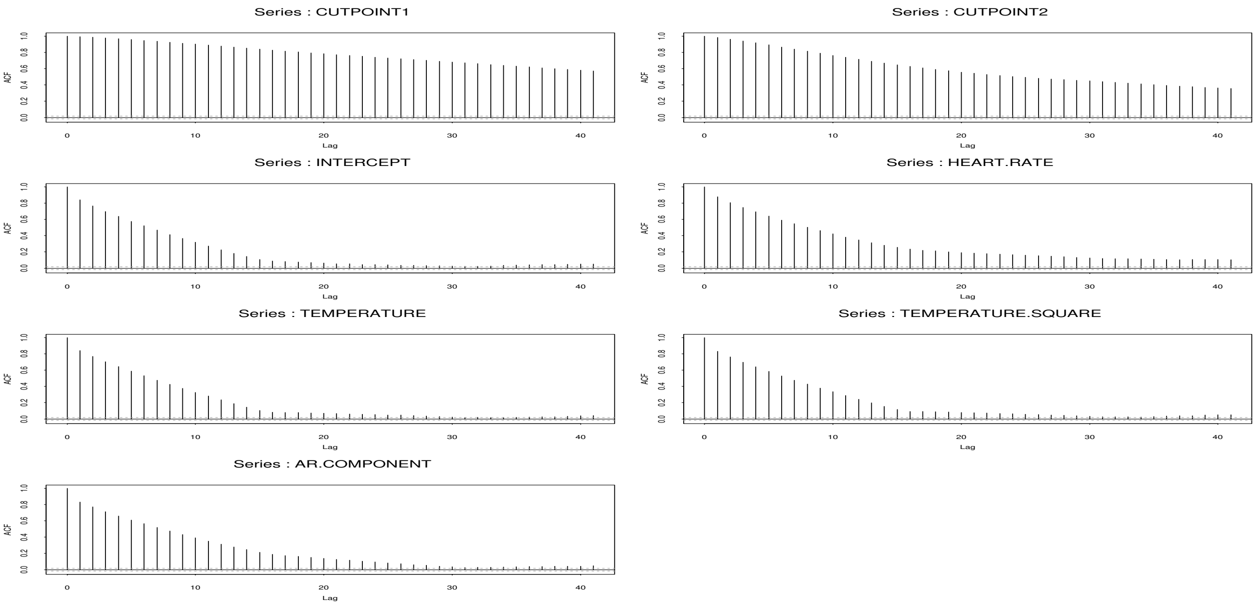

Figure 13: Autocorrelations of covariates α_1, α_2, β_0, β_1, β_2, β_3 and β_4 in model with autoregressive component.

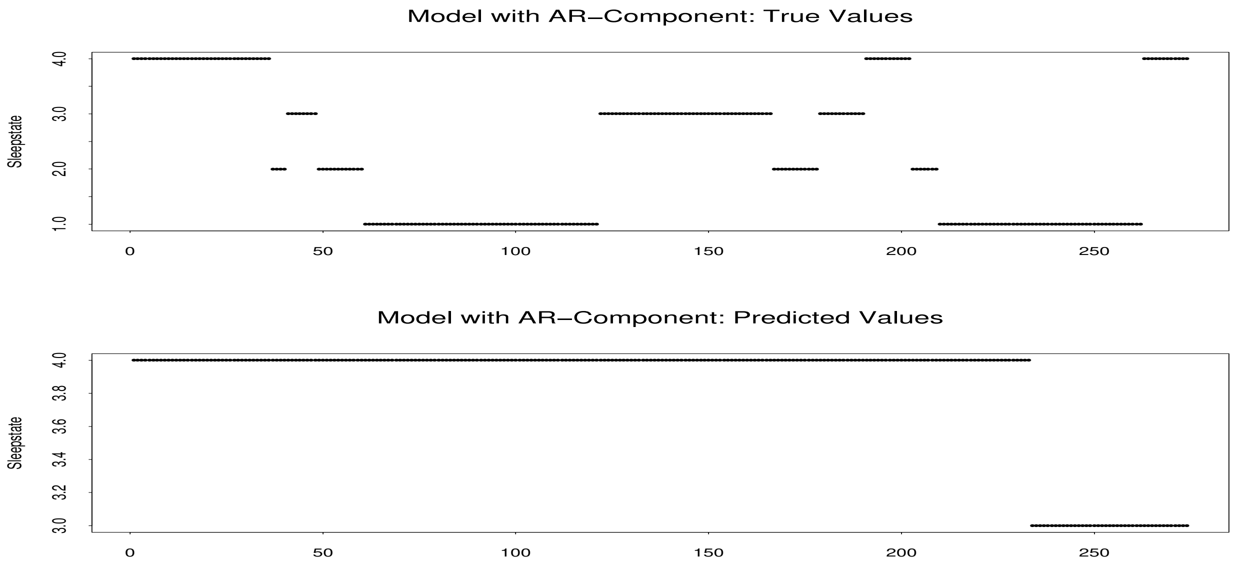

Figure 14: True Values and Prediction of the last 274 data points using the model with AR-Component.

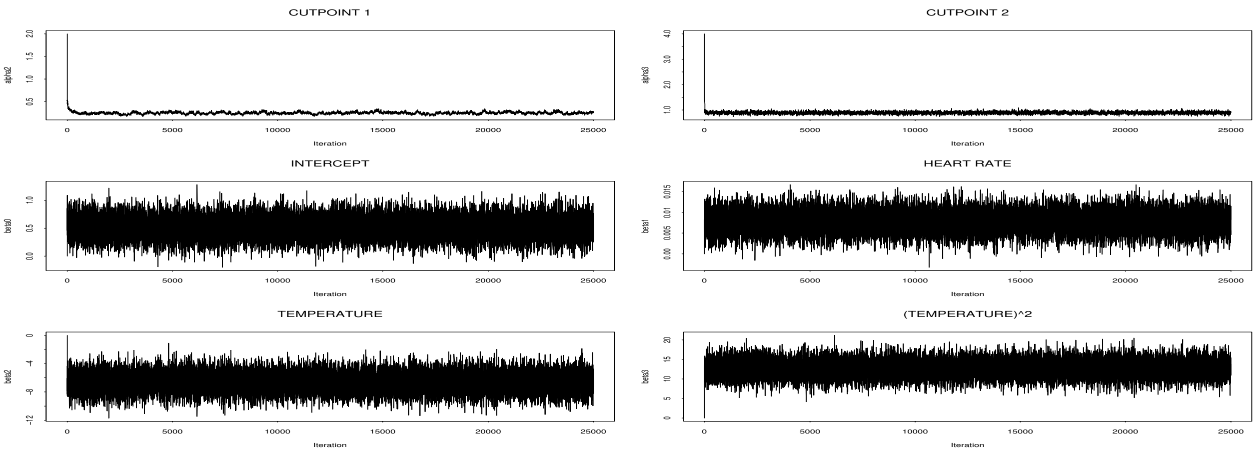

Figure 15: Time series plots of 25000 iterations of Gibbs-Sampler (ordered probit model).

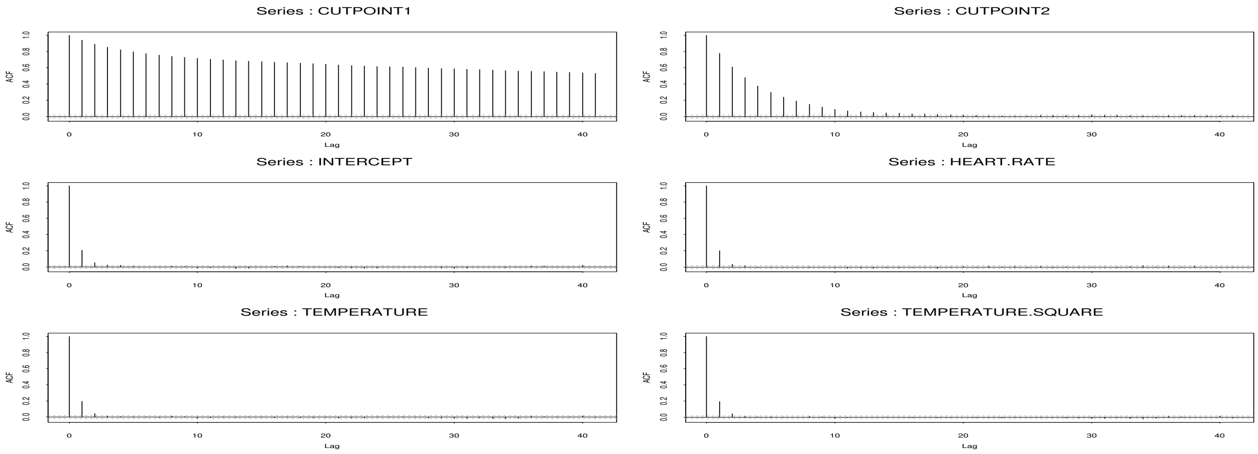

Figure 16: Autocorrelations of sampled values for α_1, α_2, β_0, β_1, β_2 and β_3 in model without autoregressive component.

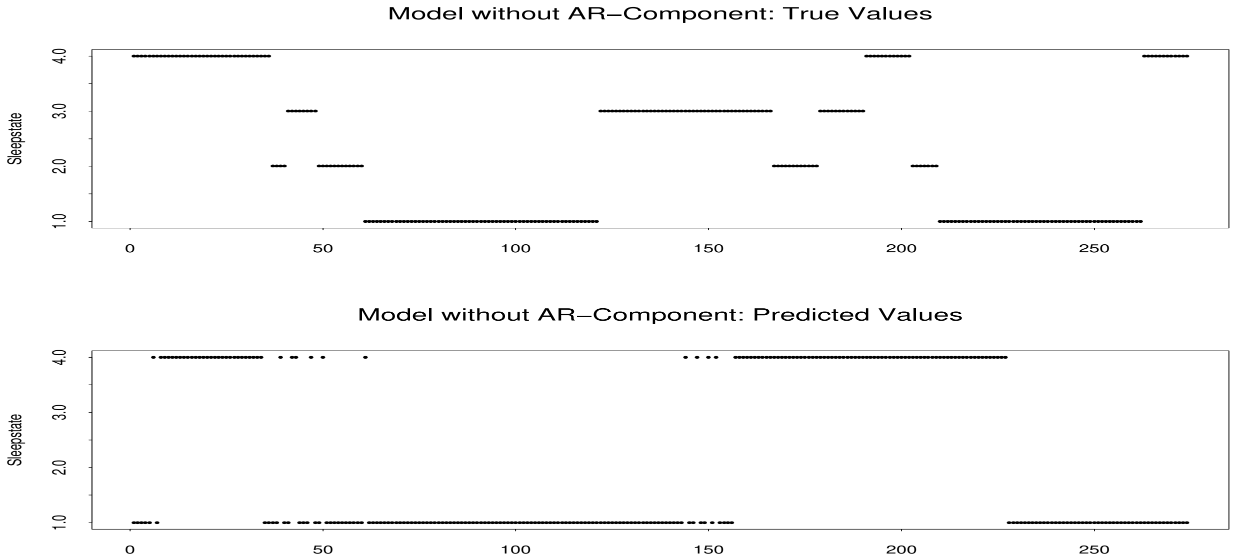

Figure 17: True Values and Prediction of the last 274 data points using the model without AR-Component.