

The Role of the Log Transformation in Forecasting Economic Variables

HELMUT LÜTKEPOHL
FANG XU

CESIFO WORKING PAPER NO. 2591
CATEGORY 12: EMPIRICAL AND THEORETICAL METHODS
MARCH 2009

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

The Role of the Log Transformation in Forecasting Economic Variables

Abstract

For forecasting and economic analysis many variables are used in logarithms (logs). In time series analysis this transformation is often considered to stabilize the variance of a series. We investigate under which conditions taking logs is beneficial for forecasting. Forecasts based on the original series are compared to forecasts based on logs. It is found that it depends on the data generation process whether the former or the latter are preferable. For a range of economic variables substantial forecasting improvements from taking logs are found if the log transformation actually stabilizes the variance of the underlying series. Using logs can be damaging for the forecast precision if a stable variance is not achieved.

JEL Code: C22.

Keywords: autoregressive moving average process, forecast mean squared error, instantaneous transformation, integrated process, heteroskedasticity.

Helmut Lütkepohl
Department of Economics
European University Institute
Via delle Piazzola 43
50133 Florence
Italy
helmut.luetkepohl@eui.eu

Fang Xu
Max Weber Programme
European University Institute
Via delle Fontanelle 10
50014 San Domenico di Fiesole
Italy
fang.xu@eui.eu

March 10, 2009

1 Introduction

Many variables are used in logarithms (logs) in economic studies because this transformation is deemed appropriate for one reason or another. From the point of view of univariate time series modelling the log transformation may be used because a transformed version of the variable of interest may be better modelled with standard linear autoregressive integrated moving average (ARIMA) processes. For example, a logarithmic transformation is often employed to obtain a more homogeneous variance of a series. In this study we investigate the implications for forecasting the original variable. If the log series is well described by an ARIMA model, optimal forecasts can be obtained easily for the log series. Of course, one may reverse the log transformation by applying the exponential function to the forecasts and thereby obtain a forecast of the original variable. This approach has a drawback in the present situation, however. It is well-known that an instantaneous nonlinear transformation applied to the optimal forecast of a variable may not result in the optimal forecast of the transformed variable (Granger and Newbold (1976)). In particular, if optimal forecasts of the logs are available, converting them to forecasts for the original variable by applying the exponential function is in general not optimal.

In this study we compare different forecasts for variables which are typically used in logs in economic models. The following forecasts are compared: (1) An ARIMA forecast for the original variable without the log transformation. This forecast is not implausible because ARIMA models capture the conditional mean and this is what is important for point forecasts. The log transformation is typically used to stabilize the variance and hence has an impact on the second moments. (2) An ARIMA forecast based on the logs of the series, where the forecast of the original series is obtained by applying the exponential function to the forecast of the log series. (3) The forecast for the log series obtained under (2) is converted to a forecast for the original series by a more sophisticated transformation which gives a more efficient forecast under ideal conditions.

We conduct a simulation experiment to investigate the performance of the different forecasts under controlled conditions and we also use the three predictors to forecast a range of economic variables. It is found that the log transformation can lead to substantial reductions in forecast mean squared error (MSE) if taking logs really leads to a more stable variance of a series of interest. On the other hand, if the log transformation is applied although it does not make the variance more homogeneous, using it can be damaging to the forecast precision.

The plan of the study is as follows. In the next section the models and

forecasts are summarized formally and some related results regarding the efficiency of different forecasts are reviewed. In Section 3 the results of simulation experiments are reported and in Section 4 a forecast comparison for a set of economic variables is presented. Finally, Section 5 concludes and detailed data sources are given in an Appendix.

2 The Predictors

2.1 Stationary Processes

Let $x_t = \log y_t$ be the natural logarithm of the univariate time series variable y_t and suppose that x_t is generated by a stationary ARMA(p, q) process,

$$\alpha(L)x_t = \nu + \theta(L)\varepsilon_t, \quad (2.1)$$

where ν is a constant, $\alpha(L)$ and $\theta(L)$ are polynomials in the lag operator, L , of orders p and q , respectively, and $\varepsilon_t \sim \text{i.i.d.}\mathcal{N}(0, \sigma_\varepsilon^2)$ is Gaussian white noise. In practice, there may be other deterministic terms such as seasonal dummy variables or deterministic trends. Although the analysis can be easily generalized to account for such terms, we ignore them because in this study we focus on the stochastic part of the data generation process (DGP).

Granger and Newbold (1976) show that the process y_t is in fact stationary if x_t has this property. They also show that y_t is a finite order MA process of order at most q if $x_t \sim \text{MA}(q)$. If, however, x_t is a mixed ARMA(p, q) process with nontrivial AR part ($p > 0$), then the covariance structure becomes more complicated. In any case, it is possible that an h -step forecast for y_{t+h} given y_t, y_{t-1}, \dots , is based on an ARMA model fitted to the variable of interest, y_t . The forecast obtained in this way by using the usual forecasting formula is denoted by

$$y_{t+h|t}^{\text{lin}}$$

It may be the optimal *linear* forecast for y_{t+h} , e.g., if x_t is a zero-mean finite order MA process. In the following we refer to a forecast based on an ARMA or ARIMA model for the original y_t variable as *linear forecast*.

Another plausible forecast for y_{t+h} may be obtained via x_t . Because x_t is a stationary Gaussian ARMA process, the usual forecasting formulas result in the conditional expectation which is the optimal (minimum MSE) predictor. We use the notation

$$x_{t+h|t} = E(x_{t+h}|x_t, x_{t-1}, \dots).$$

This forecast is unbiased, i.e., the forecast error has mean zero and its variance, denoted by $\sigma_x^2(h)$, equals the forecast MSE. A naive h -step forecast for

y_{t+h} may be based on $x_{t+h|t}$ by reversing the log transformation,

$$y_{t+h|t}^{nai} = \exp(x_{t+h|t}). \quad (2.2)$$

Granger and Newbold (1976) call this forecast naive because it is not the optimal forecast. They show that in this case the optimal forecast can be obtained as

$$y_{t+h|t}^{opt} = \exp(x_{t+h|t} + \frac{1}{2}\sigma_x^2(h)). \quad (2.3)$$

The forecasts $y_{t+h|t}^{lin}$, $y_{t+h|t}^{nai}$ and $y_{t+h|t}^{opt}$ are compared in the following.

2.2 Integrated Processes

As usual we define the variable y_t to be integrated of integer order d ($I(d)$), if its DGP is nonstationary but the DGP of the d times differenced variable is stationary, while differencing $d - 1$ times will not suffice to achieve stationarity. Suppose that the original variable, y_t , is $I(d)$. In principle one may in this case use all the predictors that we considered in the foregoing. Rather than ARMA models one shall use ARIMA models, however. Apart from that the same predictors can be used. In the following we only consider the $I(1)$ case, that is, all integrated variables are $I(1)$ for convenience because that is the most important case from a practical point of view.

If the transformed variable $x_t = \log y_t$ is integrated, the situation becomes a bit more complicated. Typically, if y_t needs to be transformed to obtain a Gaussian ARIMA process, it may not have the usual characteristics of an $I(1)$ variable, as pointed out by Granger and Hallman (1991). These authors note that the autocorrelations of $y_t = \exp x_t$ may decay more quickly than for an $I(1)$ variable if x_t is a random walk, i.e., a simple $I(1)$ variable. Moreover, the usual Dickey-Fuller and augmented Dickey-Fuller tests for unit roots have quite different distributions from the standard ones when applied to y_t . Thus, it may well be that y_t is not classified as $I(1)$ if $x_t = \log y_t$ is clearly found to be $I(1)$.

Although this problem may occur in practice and may lead a forecaster to proceed differently in some cases, in general, the leading approach still seems to be that an ARIMA model is fitted and there is also a good chance that y_t is classified as $I(1)$ if x_t is $I(1)$. Therefore, in Section 4 where we consider real economic time series, we focus on series for which the original and the logs are likely to be classified as integrated. Before we look at actual economic series, we study the performance of the predictors in a controlled simulation environment in the next section to get a feeling for what to expect in actual applications.

3 Simulation Comparison of Forecasts

In this section we use a simulation experiment to compare the three predictors introduced in Section 2. We consider two different situations. In the first set of simulations we generate x_t by an ARI process so that the variances of the logs of $y_t = \exp(x_t)$ are indeed homogeneous. In contrast, in a second set of simulations we generate y_t by an ARI process. Hence, the log transformation is applied although y_t already has a homogeneous variance.

3.1 Linear DGP of the log Series

We first use an AR(1) process for the first differences of x_t , i.e., an ARI(1,1) process, to simulate x_t . In other words, denoting the differencing operator by $\Delta (= 1 - L)$, our DGP has the form

$$\Delta x_t = \nu + \rho \Delta x_{t-1} + \varepsilon_t, \quad t = 1, 2, \dots, \quad (3.1)$$

with $x_0 = x_{-1} = 0$, $\rho = -0.9, -0.5, 0, 0.5, 0.9$, ν is a constant term which induces a drift in the levels of the integrated x_t series and ε_t is independent, identically normally distributed with zero mean and variance σ_ε^2 , i.e., $\varepsilon_t \sim \text{i.i.d.}\mathcal{N}(0, \sigma_\varepsilon^2)$. Samples of size $T = 40$ and 80 are considered and 4 post-sample values are generated additionally to evaluate the forecasts. Furthermore, we have discarded 50 values at the beginning of each sample to alleviate start-up effects and we add as many pre-sample values to each sample as needed for model selection and estimation. Thus, T is the net sample size for the levels series. Note, however, that the net sample size for the differenced series is $T - 1$ because one observation is lost by taking first differences. The variable $y_t = \exp(x_t)$ is computed from the generated x_t series.¹

To simulate an approach which is used in applied work, we fit only AR(p) processes with an intercept to the first differences of x_t and y_t . The three forecasts for y_t as summarized in Section 2 are computed for forecast horizons up to $h = 4$. The AR orders are chosen by model selection criteria. More precisely we use the very parsimonious SC (Schwarz (1978)) and the more profligate AIC (Akaike (1973)) to choose the lag orders (see also Lütkepohl (2005, Section 4.3.3) for a more detailed discussion of the model selection criteria). We use maximum lag orders of 4 and 6 for samples of size $T = 40$ and 80 , respectively, in the selection procedure and we also experimented with other maximum orders. For our DGPs small changes in the maximum AR order did not change our main results.

¹All computations are performed with MATLAB programs.

The forecast error variances required for the optimal forecasts are estimated as follows. Let $\hat{\alpha}_1, \dots, \hat{\alpha}_{p+1}$ be the estimated coefficients of an AR($p+1$) model for the levels variables. The AR operator is $1 - \hat{\alpha}_1 L - \dots - \hat{\alpha}_{p+1} L^{p+1} = \hat{\rho}(L)(1 - L)$, where $\hat{\rho}(L)$ is the estimated AR(p) polynomial of the series in differences. Furthermore, let $\hat{\sigma}_\varepsilon^2 = (T - 1 - p)^{-1} \sum_{t=2}^T \hat{\varepsilon}_t^2$ be the corresponding estimator of the residual variance. Then we set $\hat{\phi}_0 = 1$, compute

$$\hat{\phi}_i = \sum_{j=1}^{\min(i, p+1)} \hat{\phi}_{i-j} \hat{\alpha}_j$$

recursively for $i = 1, 2, \dots$, and determine the estimator for the h -steps ahead forecast error variance as

$$\hat{\sigma}_x^2(h) = \hat{\sigma}_\varepsilon^2 \sum_{i=0}^{h-1} \hat{\phi}_i^2 \quad (3.2)$$

(see Lütkepohl (2005, Section 6.5) for a justification). We compute forecast MSEs on the basis of 10 000 replications of the experiment.

Results of simulation experiments with different parameter values are presented in Tables 1 and 2. The parameter values are chosen so that they are roughly in line with some of the AR models used for the actual economic variables in Section 4. In particular, small values of the residual variance σ_ε^2 and the drift parameter ν are typical in practice when ARIMA models are fitted to logs of economic time series. In Table 1 MSEs of naive forecasts relative to linear forecasts are given for sample sizes of $T = 40$. The AR order is selected by SC. Numbers greater than one indicate that the MSE of the naive forecast is larger than the one of the linear forecast. Clearly, there are some numbers greater than one in Table 1. Notably for a zero drift term using logs does not help to improve the forecasts except that the AR coefficient has a large positive value, $\rho = 0.9$. The losses due to using logs are minimal, however. The largest loss in Table 1 for $\nu = 0$ is obtained when $\rho = 0.5$, $\sigma_\varepsilon^2 = 0.001$ and 4-steps ahead forecasts are considered. Even in that case the MSE of the naive forecast is only about 5% larger than the one of the linear forecast. Of course, if the drift is zero, it is possible that x_t is a time series of negative values and applying the exponential function may actually reduce the variability of the series. In turn, applying the log transformation to y_t may not result in sizable improvements in the homogeneity of the variance. In that case, using the log series may not improve the forecasts.

The situation is quite different if the drift term is positive. Then, depending on the residual variance, the AR coefficient ρ and the forecast horizon, the efficiency gains from using logs can be dramatic. For example, for $\sigma_\varepsilon^2 = 0.001$,

$\rho = 0.9$ and forecast horizon $h = 4$, the MSE of the linear forecast is more than 10 times as large as that of the naive forecast if the drift parameter $\nu = 0.02$. In fact, in general, if using logs is beneficial for the forecast precision, the gains tend to increase with the forecast horizon. Also, for a given residual variance, a larger drift value and a larger ρ tend to make the log transformation more beneficial or, in other words, they improve the forecast MSE of the naive forecast relative to the linear one. This result may not be too surprising because a larger drift term induces more irregularity in the variances of $y_t = \exp(x_t)$. For example, x_t tends to increase more rapidly if the drift term is larger. Hence, the magnifying effect of the exponential function for positive values induces more heteroskedasticity in y_t . Thus, the overall conclusion from the results in Table 1 is that the log transformation helps to improve forecasts a bit or even substantially if it actually has a sizable stabilizing effect on the variance of the series of interest, y_t .

This conclusion turns out to be robust in various dimensions. For example, we have also used the AIC criterion for AR order selection. The resulting relative forecast MSEs are very similar to those in Table 1. This may not be very surprising given the simple AR structure of the DGP. Although we have not checked this, one may guess that AIC and SC selected the same AR orders in the vast majority of all cases. We also considered a larger sample size of $T = 80$. For zero drift term, $\nu = 0$, the results are again similar to the corresponding ones in Table 1. For $\nu > 0$, the relative performance of the naive forecast tends to improve, however, in some situations.

Of course, so far we just compared the linear forecast to its naive competitor which is theoretically inferior to the optimal forecast. In practice it is not clear that the optimal forecast actually outperforms the other two predictors because the forecast error variance which is used in the forecast transformation is unknown and has to be replaced by an estimate. Using the estimate from (3.2) does not in fact improve the forecast precision much over the naive one if at all. This can be seen in Table 2 where the MSEs of the optimal forecasts relative to the naive ones are presented. Clearly, the numbers tend to be very close to but still slightly larger than one. This holds across all forecast horizons, drift terms and error variances. In fact, the results in Table 2 are invariant to the value of the drift parameter ν . Therefore we report results only for $\nu = 0$. Given these simulation results, in applied forecasting, the optimal forecast may not be of great value. At least it does not improve the forecast MSEs in the experimental situations which we consider here.

The DGPs considered so far favor the forecasts based on logs because this transformation has the potential to make the variance more stable. One may, of course, also wonder how much can be lost by applying the log trans-

formation when the variance is stable already. This question is considered next.

3.2 Linear DGP of the Series of Interest

To investigate whether the log transformation can be damaging to the forecast precision if it does not stabilize the variance, we perform another experiment where we generate the variable of interest, y_t , by an ARI process and we apply logs to obtain x_t . More precisely, the DGP of Δy_t is an AR(1),

$$\Delta y_t = \nu + \rho \Delta y_{t-1} + \varepsilon_t, \quad t = 1, 2, \dots, \quad (3.3)$$

with $y_0 = y_{-1} = 0$ and all other quantities are specified as in (3.1). The x_t 's are generated as $x_t = \log y_t$. We use only processes with positive drift term, $\nu > 0$, to ensure positive values of y_t at least after the initial burn-in period of 50 observations which are dropped as in the previous simulation setup. Some results based on samples of size $T = 40$ and AR order selection by SC are reported in Table 3. Again we show MSEs of naive forecasts relative to linear forecasts.

Now all entries are greater than one and some substantially so. Consider, for instance, the relative MSEs associated with 4-step ahead forecasts when $\nu = 0.05$ and $\sigma_\varepsilon^2 = 0.0001$. Taking logs can lead to MSEs which are more than six times as large as those of the linear forecasts. Thus, quite a bit of damage can be done by taking logs if the variance of a series is stable across the sample already. In particular, longer-term forecasts based on logs can be quite poor relative to the linear ones. But even for short-term forecasts (1-step ahead) sizable losses are possible.

Also in this case we checked the robustness of these results. For instance, we used AIC instead of SC for AR order selection and we increased the sample size to $T = 80$. The results remain qualitatively the same.

Thus, the overall conclusion from our simulations so far is that forecasting logs of a variable of interest first and then converting back to the original variable can lead to substantial improvements in MSE if the log transformation indeed stabilizes the variance of the DGP. If however, the variance is stable already, quite some damage can be done by forecasting logs. Our simulations also show that no substantial gains can be expected from using the optimal instead of the naive forecast. In fact, the optimal forecast with estimated forecast error variance often has a slightly larger MSE than the naive forecast. The differences are usually minimal, however. Thus, there is no compelling reason for considering the more elaborate optimal forecast.

4 Forecast Comparison Based on Economic Data

We consider a range of different economic time series which are often used in logs in economic modelling and compute the three different forecasts for the original variables. More precisely, we use monthly series of different stock indices as well as quarterly series of gross domestic product (GDP) and consumption for a range of countries. Although stock returns are often of interest, the level of stock indices is also of interest because the returns of many certificates are linked to the level of specific stock indices. If the returns are of interest, first differences of logs are typically considered. Hence, using logs of the stock index series is not uncommon. Similarly, in economic analyses, logs of GDP and consumption are often considered. Although forecasts of the rates of change of these series may be more important in practice, we focus on forecasting the levels because thereby we hope to cover a good range of different DGPs which come up in applied work. Having a range of different DGPs is important in this context because the characteristics of the DGP are crucial for the performance of the different forecasts, as we have seen in the simulations. We first discuss the results for the stock indices and then consider the GDP and consumption series.

4.1 Forecasting Stock Indices

We consider nine well-known stock indices, the Dow Jones Euro Stoxx 50, FTSE, DAX, CAC 40, Dow Jones, Nasdaq, S&P 500, Nikkei and HangSeng. The indices are related to important stock exchanges from all over the world. They measure stock prices from different regions and sectors. Moreover, they differ in how many stock prices they incorporate. Overall they cover a good range of the stock markets in the world. We use monthly series from 1990M1 to 2007M12 based on end-of-month index values. Details on the data sources are provided in Appendix B.

The first differences and first differences of logs of all series are plotted in Figure 1. Apparently the variation in volatility is larger for the series without logs in most cases. Thus, using logs can be interpreted as a means for stabilizing the variance. However, even the first differences of logs still show considerable variation in the volatility, as one would expect for monthly financial market series. In such series conditional heteroskedasticity is often diagnosed and modelled. To ensure that our results are not driven purely by the specific period used for the forecast comparison, we report results for different forecast periods. Also, we vary the sample on which model

specification and estimation are based.

Forecasts are computed by fitting AR models to the first differences of the original index series and to the first differences of logs. The AR orders are chosen by model selection criteria, as in the simulations, that is, we use the SC and the AIC for AR order selection. The maximum lag order is 4 in the selection procedure because no seasonality or higher order AR dynamics are expected in the differenced series. In fact, for the current set of series, the AR orders chosen by both model selection criteria are often zero. In efficient markets it is of course not surprising to find no predictability in the returns. Since the choice of selection criterion did not make any difference for the results qualitatively, we report relative forecast MSEs only for SC models in Table 4. In fact, the AIC results are identical for most countries and samples and very close to the SC results when there are differences. The reason is, of course, that the AR(0) for the first differences is the dominating model. For our purposes, computing the forecasts from those models means that estimation uncertainty has only a limited impact on the results. The only estimated parameters are the drift term and the residual variance which enters the estimator for the forecast error variance in the optimal forecast formula.

In Table 4 we report relative forecast MSEs for two alternative sample periods, three different forecasting periods and forecast horizons $h = 1, 3$ and 6. Thus, the forecast horizons refer to one month, one quarter and half a year. In this table the forecast MSEs of the naive forecasts relative to the linear forecasts are displayed. An asterisk indicates that the difference between the forecast MSEs is significant at the 5% level based on a two-sided modified Diebold-Mariano (DM) test (Diebold and Mariano (1995)). We use a modification which was proposed by Harvey, Leybourne and Newbold (1997) and give the precise form of the test statistic in Appendix A.

The two sample periods used in Table 4 begin in 1990M1 and 1995M1. The forecast periods start in 2001M1, 2003M1, and 2005M1 and they all end in 2007M12. Accordingly the forecast MSEs are based on 79, 55 and 31 forecasts. Forecasts up to six steps ahead are computed based on estimated models fitted to samples of increasing length. For example, for the sample starting in 1990M1 and the forecast period 2001M1 - 2007M12, we first fit AR models using data from 1990M1 - 2000M12 (sample size 131 if the first observation is not counted which is used for forming differences) and we use these models to produce up to six steps ahead forecasts and corresponding squared forecast errors. Then we extend the sample length by one and perform a new AR order selection and estimation to produce the next set of forecasts etc.. The largest sample for estimation is achieved when the period ends in 2007M6. Thereby we generate 79 squared forecast errors for each

forecast horizon. For the shorter forecast periods starting from 2003M1 and 2005M1, only 55 and 31 forecasts, respectively, are computed in this way. We divide the sums of the squared forecast errors by the corresponding sum of squared forecast errors obtained for the original series. Thus, a value greater than one in Table 4 means that the MSE of the linear forecast is smaller than the one of the naive forecast based on logs, that is, logs do not improve the forecasts, whereas a number smaller than one implies that taking logs improves the forecasts.

Obviously, numbers smaller than one dominate in Table 4. Thus, producing forecasts on the basis of the log series is clearly beneficial. In fact, all numbers greater than one typically exceed one only by very little, while gains from taking logs can be considerable. This is similar to the simulation results reported in Table 1. For example, the largest relative forecast MSE in Table 4 is about 1.24 which occurs for the Euro Stoxx index when the sample period starts in 1995M1 (smallest sample size 71 for model selection and estimation) and the longest forecast period 2001M1-2007M12 is considered. Thus, a potential loss in forecast efficiency of about 24% is incurred by using logs. Note, however, that the difference in the two MSEs is not significant at the 5% level based on the DM test. In fact, the only significantly larger MSE of the naive forecast is obtained for the 1-step ahead forecast of the Euro Stoxx for the shorter sample period and forecast period 2001M1-2007M12. On the other hand, there are many cases where the naive forecasts produce significantly smaller MSEs than the linear forecasts. In a number of cases the relative forecast MSEs are smaller than 0.80, that is, forecast efficiency gains of more than 20% are found. For example, for the Dow Jones index, using the longer sample and the shortest forecast period 2005M1-2007M12, the relative forecast MSE is 0.7445 meaning that the MSE of the naive forecast is only about 3/4th of the one of the linear forecast.

In fact, most forecast MSEs above one occur for the longest forecast period 2001M1-2007M12 which covers the general downturn in the stock markets in the early years of the current millennium. Had we eliminated this forecast period, the advantage of the forecasts based on logs would have been overwhelming. The results for the long forecast period show that the precision of specific forecasts in practice is quite dependent on the sample and forecast periods. We account for this fact by reporting results for different periods.

The optimal forecasts typically have MSEs close to the naive forecasts or are slightly better. The MSEs of the optimal forecasts relative to the naive ones are shown in Table 5. Obviously, they are all very close to one, as in the simulations. In no case has the optimal forecast more than a 10% higher MSE than the naive one. On the other hand, forecast efficiency gains

from using the optimal forecast over the naive one of more than 10% are also rare. Notice, however, that with one exception, all significant differences are obtained in situations when the optimal forecast is better. The only exception occurs for the Euro Stoxx when the shorter sample and a forecast period from 2001M1 - 2007M12 is used. In that case the MSE improvement from using the naive forecast is only about 1.5%. Thus, overall the optimal forecasts are about equally good or even slightly better than the naive forecasts for these samples and forecast periods. Hence, the gains from using logs in forecasting the stock indices would be even greater in some occasions than is seen in Table 4 had we used the optimal forecast in that table instead of the naive one. It may also be worth noting that the logs of the stock indices are not well modelled by normal distributions whereas our optimal forecast is derived under the normality assumption.

As mentioned earlier, many of the models underlying the forecasts in Tables 4 and 5 are AR(0) models for the first differences. Hence, these results do not tell us much about the impact of estimation uncertainty in the AR parameters. Such estimation uncertainty may in particular have an impact on the optimal forecast which involves an estimate of the forecast error variance and is, thus, based on the AR coefficient estimates. Therefore we now consider series for which more dynamics in the differences and the differences of logs can be expected.

4.2 GDP

Quarterly, seasonally adjusted GDP series from seven different OECD countries for the period 1980Q1 - 2006Q4 are investigated. The precise data sources are again given in Appendix B. Our choice of series is determined by the objective of our forecast comparison, namely to see whether taking logs is beneficial for forecasting even if forecasts of the original series are of interest. The countries are Belgium, Canada, Denmark, France, Japan, Norway and the U.S.. Thus, we have a range of smaller, medium size and larger countries in our set. The first differences of the original variables and the logs are plotted in Figure 2. The main criterion for including these countries is to ensure that there are no major distortions and data irregularities during the sample period. For instance, we exclude Germany because the unification in 1990 resulted in a series with a shift in that year. The shift is due to the fact that the GDP series refers to West Germany only before the unification. Of course, we could have adjusted the series in one way or another. We did not want to include series for which manual adjustments on our side were necessary to ensure a reasonably good fit, however. Thereby we try to safeguard against the critique that our results may be driven by

our adjustments. Thus, we only include countries for which the GDP series both in first differences and in first differences of logs can be modelled reasonably well by low order AR processes for the entire sample period. This is also why we use data which are adjusted by some kind of official seasonal adjustment procedure. The residuals of the AR processes fitted to the series of changes (first differences of the original series) may still be heteroskedastic, however, which is why logs are often used. In some cases there is still some heteroskedasticity left in the series if the rates of change are considered, as can be seen in Figure 2. Again the first differences of logs on the right hand side of the figure generally appear to have a more stable variance over the sample period, although that is not obvious in all cases. In fact, taking logs in some cases seems to lead to a change in variance which may, e.g., be attributed to the great moderation. There is a substantial literature which discusses the possibility of a reduction in the volatility of U.S. series from the middle of the 1980s onward (e.g., Sims and Zha (2006), Sims, Waggoner and Zha (2008), Lanne and Lütkepohl (2008)). This phenomenon is also seen to some extent in the last panel in Figure 2 which shows the first differences of logs of U.S. GDP. However, overall the log transformation seems to stabilize the variability in the GDP series considered.

The forecasts are computed as explained for the stock indices. That is, we fit AR models to the original GDP series and to their logs. The AR orders are chosen by AIC and SC model selection criteria based on increasing sample sizes for each fixed sample beginning. The maximum lag order considered is now 8 because we expect some more serial dynamics in the first differences of the series. Again we just report SC results because the AIC results are qualitatively identical. MSEs of naive forecasts relative to MSEs for linear forecasts for samples starting in 1980Q1, 1985Q1 and 1990Q1 are reported in Table 6. The MSEs for the optimal forecasts are either very close to those of the naive forecasts or they are slightly smaller. Thus, whenever a number smaller than one appears in the table, the log transformation improves the forecasts. Had the optimal forecasts been used, the improvements may have been even slightly larger. The forecast periods start either in 2000Q1 or in 2003Q1 and results for forecast horizons from one to four quarters are reported. We also used different estimation and forecast periods and found similar results so that we believe that the results reported in Table 6 provide a good summary of the overall outcome.

The general picture in Table 6 is again in favor of using logs. In fact, Japan is the only country for which the linear forecasts have significantly smaller MSEs than the naive ones in some cases, where significance is again assessed by (two-sided) DM tests with 5% level. For the other countries there are quite substantial improvements due to the log transformation notably for

longer term forecasts. Consider, for example, Belgium where in all cases the log transformation improves the forecasts and in a couple of cases the relative MSE is less than 10%, that is, a dramatic improvement is obtained (see the figures for forecast horizon 4 and forecast period 2003Q1-2006Q4). In fact, in most cases where the forecasts based on the original variables are superior, the gains are only small, typically a few percent improvements. The only exception is Denmark for which the original variables deliver more than 50% improvements in the forecast period 2000Q1 - 2006Q4. This improvement is not significant at the 5% level according to the DM test, however. More generally for this forecast period the forecasts based on the original variables are often superior although not significantly so when checking with the DM test. In fact, if one would eliminate the first column of the table which shows the relative forecast MSEs for this period and the rows associated with Japan, the picture would be even more clearly in favor of taking logs.

Thus, in summary, we can conclude from the results for the GDP series that taking logs is beneficial from a forecasting point of view. Forecasting the logs first and then converting to forecasts for the original variables can lead to dramatic MSE improvements. On the other hand, for those cases where the forecasts based on the original variables exhibit smaller MSEs, the gains are typically very limited. This situation is similar to the simulation results reported in Table 1. In the next section we check whether similar results are obtained for consumption series from different countries.

4.3 Consumption Forecasts

The next set of variables we use for a forecast comparison are quarterly, seasonally adjusted aggregate private consumption expenditures for a range of OECD countries, more precisely for Australia, Belgium, Canada, Japan, Norway, United Kingdom (UK) and the U.S.. Thus, the countries overlap with those for which we considered GDP series in the previous subsection, but they are not identical. The sample period is the same as for GDP, that is, 1980Q1-2006Q4. The precise data sources are again given in Appendix B and the first differences and first differences of logs are depicted in Figure 3. For some of the series the log transformation apparently leads to a series with a clearly more homogeneous variance in first differences than without logs. Norway is a particularly clear case. On the other hand, there are also series such as the one from the UK where the advantage of taking logs for stabilizing the variance is not apparent. Whereas the variability of the first differences seems to increase over the sample period, it appears to decrease for the first differences of logs. Clearly, one may question the log transformation in such a case. Still, the fact remains that logs of consumption series are

often considered in economic modelling and, hence, it is of interest to check whether the log transformation can be beneficial for forecasting as well.

We produce forecasts in the same way as for the GDP series, using the same sample and forecasting periods, and report MSEs of naive forecasts relative to linear forecasts in Table 7. They are again based on SC models because the AIC results are qualitatively similar. Also the optimal forecast MSEs are again similar to those of the naive forecasts. For five countries there are partly substantial gains from using logs. For example, for Norway the 4-steps ahead forecasts based on the logs for the sample period starting in 1985 and forecasting period 2003Q1-2006Q4 produce an MSE which is less than 20% of the corresponding linear forecast MSE. Many of the MSEs of the naive forecasts are significantly smaller than those of the linear forecasts.

In contrast, using logs for forecasting Japanese and UK consumption results in considerable and significant efficiency losses. In particular, for the UK the MSEs of the naive forecasts based on logs can be more than five times the corresponding MSEs of the linear forecast. These figures are similar to the simulation results in Table 3 which are obtained by simulating the y_t series with a linear DGP. Thus, for series where a stabilization of the variance is not achieved by taking logs, the log transformation may be quite damaging for the forecast precision.

5 Conclusions

In this study we have investigated whether and under which conditions using logarithms can help improving forecasts of economic variables. More precisely, if forecasts of a variable y_t are of interest, the question is under what conditions forecasting $x_t = \log y_t$ and then converting the forecast of x_t to a forecast of y_t may lead to a more precise forecast than predicting y_t directly. To explore this question, we have compared three predictors: (1) a linear forecast based on an ARIMA model for y_t , (2) a naive forecast which converts an ARIMA forecast of x_t by the exponential transformation to a forecast of y_t and (3) an ‘optimal forecast’ which adjusts the ARIMA forecast of x_t to account for the nonlinearity of the log transformation when converting to a forecast of y_t . The MSE has been used as a measure for forecast precision.

In a simulation study based on ARI processes for x_t as well as for y_t we found that using logs can result in dramatic gains in forecast precision if the log transformation indeed makes the variance more homogeneous throughout the sample. In other words, forecasts based on $x_t = \log y_t$ and then converting to y_t can be much better than direct predictions of y_t if x_t has a more stable

variance than y_t . On the other hand, directly forecasting y_t is preferable in terms of forecast precision if y_t has a more homogeneous variance than x_t . Generally, the so-called optimal forecast based on x_t is typically not better or at least not much better than the naive forecast based on x_t . Although the optimal predictor minimizes the forecast MSE in theory, it involves the forecasts error variance which is unknown in practice and has to be replaced by an estimator. Using the usual estimator for this quantity, the optimal predictor does not appear to have an advantage over the naive predictor in samples of common size, at least for the DGPs used in our simulation experiment. In this context it may also be worth noting that Granger and Newbold (1976) do not report large gains in theoretical forecast MSE in their examples if a log transform is used. Hence, our simulation findings are in line with their results.

We have also considered a range of economic series which are typically used in logarithmic form in economic analysis and compared the three predictors using different sample and forecast periods. The overall results from the empirical forecast comparisons are the same as those of the simulations. In other words, series whose variability becomes more homogeneous by taking logs, can be forecasted better by the naive or optimal predictors. The gains in forecast MSE can be dramatic. On the other hand, if the log transformation does not stabilize the variance of a series, it is preferable to base forecasts directly on ARIMA models for the original series. In that case, using forecasts based on the log series can be damaging to the forecast precision.

These results can be potentially important for forecasting aggregated series. If disaggregate data is available it was found that forecasting the disaggregate series and then aggregating the forecasts may be preferable to forecasting the aggregate series directly. Such results were found for both temporal as well as contemporaneous aggregation (e.g., Amemiya and Wu (1972), Wei (1978), Lütkepohl (1986, 1987, 2006), Silvestrini, Salto, Moulin and Veredas (2008)). Many of the available results relate to linear aggregation, however. Since the log transformation is a nonlinear one, the question arises whether forecasting the logs of the disaggregate series and aggregating the forecasts is still preferable to forecasting the aggregate directly based on the original series or based on its logs. For the case of contemporaneous aggregation this would require multivariate extensions of the results regarding optimal prediction of nonlinearly transformed series. Such extensions were discussed by Ariño and Franses (2000). These issues are left for future research.

Acknowledgements

We thank Gunnar Bardsen, Antoni Espasa and the participants of a workshop at CORE, where this paper was presented, for useful comments and discussions of an earlier draft.

Appendix A. Modified Diebold-Mariano Test

The following version of the Diebold-Mariano test from Harvey et al. (1997) for equality of the MSEs of different forecasts is used. Let (e_{1i}, e_{2i}) , $i = 1, \dots, N$, be a set of errors from two different procedures for computing h -step ahead forecasts and define $d_i = e_{1i}^2 - e_{2i}^2$, $i = 1, \dots, N$. The modified DM statistic has the form,

$$\text{DM} = \left(\frac{N + 1 - 2h + N^{-1}h(h-1)}{N} \right)^{1/2} \hat{V}^{-1/2} \bar{d},$$

where $\bar{d} = N^{-1} \sum_{i=1}^N d_i$ is the mean of the d_i 's and

$$\hat{V} = \frac{1}{N} \left(\hat{\gamma}_0 + 2 \sum_{k=1}^{h-1} \hat{\gamma}_k \right)$$

with $\hat{\gamma}_k = N^{-1} \sum_{i=k+1}^N (d_i - \bar{d})(d_{i-k} - \bar{d})$, is an estimator of the variance of \bar{d} . The statistic is used with a t distribution with $N - 1$ degrees of freedom and the significance level refers to a two-sided alternative.

Appendix B. Data Sources

All data considered are obtained directly from Thomson Datastream where data from international sources such as International Monetary Fund (IMF) or national sources such as Banque Nationale de Belgique are collected.

Stock Indices

Nine price indices from stock markets all over the world are investigated. The corresponding codes (DS Mnemonics) in the Datastream database are: DJES50I for Dow Jones Euro Stoxx 50, FTSE100 for FTSE 100, DAXINDEX for DAX 30 Performance, FRCAC40 for CAC 40, DJINDUS for Dow Jones Industrials, NASCOMP for Nasdaq Composite, S&PCOMP for Standard and Poors 500 Composite, JAPDOWA for Nikkei 225 Stock Average, and HNGKNGI for Hang Seng Index.

GDP

The GDP series are seasonally adjusted in current prices for seven countries. The corresponding codes (DS Mnemonics) are: BGGDP...B for Belgium, CNI99B.CB for Canada, DKESNGDPB for Denmark, FRL99B.CB for France, JPI99B.CB for Japan, NWGDP...B for Norway, and USI99B.CB for U.S..

The GDP of Canada is measured in billions of Canadian dollars, French GDP is in billions of French franc, Japanese GDP is in billions of Japanese yen, and U.S. GDP is in billions of U.S. dollars. These series are from IMF International Financial Statistics. Norway's GDP, from Statistics Norway, is in millions of Norwegian kroner. The data for Denmark (in billions of euros) are from Statistical Office of the European Communities and the GDP series for Belgium is given in millions of euros and is provided by Banque Nationale de Belgique.

Consumption

Seasonally adjusted time series of private consumption in current prices for seven countries are considered. The corresponding codes (DS Mnemonics) are: AUI96F.CB for Australia, BGCNPER.B for Belgium, CNI96F.CB for Canada, JPI96F.CB for Japan, NWCNPER.B for Norway, UKI96F.CB for UK, and USI96F.CB for U.S..

The household consumption expenditures of Australia (in billions of Australian dollars), Canada (in billions of Canadian dollars), Japan (in billions of Japanese yen), UK (in billions of UK sterling pounds), and U.S. (in billions of U.S. dollars) are from IMF International Financial Statistics. For Belgium, private consumption expenditures (in millions of euros) are from Banque Nationale de Belgique and the data for Norway (in millions of Norwegian kroner) are from Statistics Norway.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle, in B. Petrov and F. Csáki (eds), *2nd International Symposium on Information Theory*, Akadémiai Kiadó, Budapest, pp. 267–281.
- Amemiya, T. and Wu, R. Y. (1972). The effect of aggregation on prediction in the autoregressive model, *Journal of the American Statistical Association* **67**: 628–632.

- Ariño, M. A. and Franses, P. H. (2000). Forecasting the levels of vector autoregressive log-transformed time series, *International Journal of Forecasting* **16**: 111–116.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy, *Journal of Business & Economic Statistics* **13**: 253–263.
- Granger, C. W. J. and Hallman, J. (1991). Nonlinear transformations of integrated time series, *Journal of Time Series Analysis* **12**: 207–224.
- Granger, C. W. J. and Newbold, P. (1976). Forecasting transformed series, *Journal of the Royal Statistical Society B* **38**: 189–203.
- Harvey, D., Leybourne, S. and Newbold, P. (1997). Testing the equality of prediction mean squared errors, *International Journal of Forecasting* **13**: 281–291.
- Lanne, M. and Lütkepohl, H. (2008). Identifying monetary policy shocks via changes in volatility, *Journal of Money, Credit and Banking* **40**: 1131–1149.
- Lütkepohl, H. (1986). Forecasting temporally aggregated vector ARMA processes, *Journal of Forecasting* **5**: 85–95.
- Lütkepohl, H. (1987). *Forecasting Aggregated Vector ARMA Processes*, Springer-Verlag, Berlin.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.
- Lütkepohl, H. (2006). Forecasting with VARMA models, in G. Elliott, C. W. J. Granger and A. Timmermann (eds), *Handbook of Economic Forecasting, Volume 1*, Elsevier, Amsterdam, pp. 287–325.
- Schwarz, G. (1978). Estimating the dimension of a model, *Annals of Statistics* **6**: 461–464.
- Silvestrini, A., Salto, M., Moulin, L. and Veredas, D. (2008). Monitoring and forecasting annual public deficit every month: The case of France, *Empirical Economics* **34**: 493–524.
- Sims, C. A., Waggoner, D. F. and Zha, T. (2008). Methods for inference in large multiple-equation Markov-switching models, *Journal of Econometrics* **146**: 255–274.
- Sims, C. A. and Zha, T. (2006). Were there regime switches in U.S. monetary policy?, *American Economic Review* **96**: 54–81.
- Wei, W. W. S. (1978). Some consequences of temporal aggregation in seasonal time series models, in A. Zellner (ed.), *Seasonal Analysis of Economic Time Series*, U.S. Department of Commerce, Bureau of the Census, pp. 433–444.




Figure 1: First differences (left column) and first differences of logs (right column) of stock indices.




Figure 2: First differences (left column) and first differences of logs (right column) of GDP series.




Figure 3: First differences (left column) and first differences of logs (right column) of consumption series.

Table 1: Forecast MSEs of Naive Forecast Relative to Linear Forecast for Simulated Series, DGP $\Delta x_t = \nu + \rho \Delta x_{t-1} + \varepsilon_t$, $y_t = \exp x_t$, Sample Size 40

ν	h	$\sigma_\varepsilon^2 = 0.001$					$\sigma_\varepsilon^2 = 0.0001$				
		$\rho =$					$\rho =$				
		-0.9	-0.5	0	0.5	0.9	-0.9	-0.5	0	0.5	0.9
0	1	1.001	1.001	1.004	1.010	0.582	1.001	1.000	1.001	1.004	1.013
	4	1.002	1.007	1.017	1.051	0.515	1.000	1.001	1.002	1.006	1.037
0.01	1	1.008	1.006	1.014	0.980	0.392	1.000	0.997	0.984	0.947	0.659
	4	1.016	1.040	1.059	1.008	0.103	0.988	0.979	0.935	0.793	0.541
0.02	1	1.002	0.992	0.976	0.849	0.322	0.904	0.864	0.767	0.767	0.459
	4	0.987	0.987	0.933	0.691	0.025	0.747	0.598	0.439	0.461	0.156

Note: AR order selection based on SC with maximum lag order of 4. The number of replications is 10000.

Table 2: Forecast MSEs of Optimal Forecast Relative to Naive Forecast for Simulated Series, DGP $\Delta x_t = \rho \Delta x_{t-1} + \varepsilon_t$, $y_t = \exp x_t$, Sample Size 40

h	$\sigma_\varepsilon^2 = 0.001$					$\sigma_\varepsilon^2 = 0.0001$				
	$\rho =$					$\rho =$				
	-0.9	-0.5	0	0.5	0.9	-0.9	-0.5	0	0.5	0.9
1	1.001	1.001	1.002	1.003	0.998	1.000	1.000	1.000	1.000	1.001
4	1.001	1.003	1.006	1.019	1.059	1.000	1.000	1.001	1.002	1.011

Note: AR order selection based on SC with maximum lag order of 4. The number of replications is 10000.

Table 3: Forecast MSEs of Naive Forecast Relative to Linear Forecast for Simulated Series, DGP $\Delta y_t = \nu + \rho \Delta y_{t-1} + \varepsilon_t$, $x_t = \log y_t$, Sample Size 40

ν	h	$\sigma_\varepsilon^2 = 0.001$					$\sigma_\varepsilon^2 = 0.0001$				
		$\rho =$					$\rho =$				
		-0.9	-0.5	0	0.5	0.9	-0.9	-0.5	0	0.5	0.9
0.02	1	1.083	1.086	1.085	1.101	1.130	1.376	1.365	1.395	1.283	1.128
	4	1.233	1.334	1.358	1.346	1.392	2.175	2.612	2.731	2.276	1.521
0.05	1	1.255	1.256	1.268	1.224	1.135	1.943	1.915	1.916	1.380	1.114
	4	1.786	2.099	2.172	1.960	1.535	4.845	6.491	6.033	3.230	1.480

Note: AR order selection based on SC with maximum lag order of 4. The number of replications is 10000.

Table 4: Forecast MSEs of Naive Forecast Relative to Linear Forecast for Stock Index Series

Index	Forecast horizon	Sample 1990M1- Forecast period			Sample 1995M1- Forecast period		
		2001 -2007	2003 -2007	2005 -2007	2001 -2007	2003 -2007	2005 -2007
Euro Stoxx	1	1.0353	0.9771	0.9641	1.0570*	0.9823	0.9706
	3	1.0860	0.9325	0.9108	1.1371	0.9460	0.9262
	6	1.1520	0.8126*	0.7691	1.2417	0.8306	0.7839
FTSE	1	1.0182	0.9844	0.9789	1.0209	0.9902	0.9862
	3	1.0481	0.9393	0.9310	1.0548	0.9617	0.9535
	6	1.0781	0.8492*	0.8217	1.0906	0.8981	0.8687
DAX	1	1.0098	0.9700	0.9551	1.0224	0.9678	0.9509
	3	1.0274	0.9282	0.9051	1.0606	0.9208	0.8927
	6	1.0430	0.8624	0.8303*	1.1010	0.8393	0.7964
CAC 40	1	1.0182	0.9830	0.9772	1.0446	0.9816	0.9747
	3	1.0465	0.9584	0.9521	1.1119	0.9541	0.9501
	6	1.0815	0.9074	0.8961	1.1960	0.8878	0.8858
Dow Jones	1	1.0284	0.9795	0.9874	1.0366	0.9919	0.9984
	3	1.0786	0.9101	0.9008	1.1034	0.9495	0.9415
	6	1.1675	0.8234	0.7445	1.2183	0.9021	0.8275
Nasdaq	1	1.0283	0.9888	1.0169	1.0277	0.9965	1.0143
	3	1.0889	0.9621	0.9947	1.0855	0.9847	1.0027
	6	1.1692	0.9010	0.8419	1.1614	0.9500	0.8941
S&P 500	1	1.0340	0.9757	0.9975	1.0419	0.9883	1.0077
	3	1.0907	0.9140	0.9394	1.1122	0.9501	0.9776
	6	1.1682	0.8038	0.7815	1.2086	0.8790	0.8780
Nikkei	1	0.9430*	0.9022*	0.9340*	0.9907	0.9806*	0.9926*
	3	0.8783	0.8177*	0.8766	0.9758	0.9563*	0.9828
	6	0.8141	0.7541*	0.8350	0.9574	0.9314	0.9712
HangSeng	1	0.9971	0.9275*	0.9034	0.9950	0.9681*	0.9523*
	3	0.9779	0.8397	0.8031	0.9819	0.9288	0.9040
	6	0.9228	0.7944	0.7732	0.9557	0.9089	0.8918

Note: AR order selection based on SC with maximum lag order of 4.

* significant at 5% level according to DM test with two-sided alternative.

Table 5: Forecast MSEs of Optimal Forecasts Relative to Naive Forecasts for Stock Index Series

Index	Forecast horizon	Sample 1990M1- Forecast period			Sample 1995M1- Forecast period		
		2001 -2007	2003 -2007	2005 -2007	2001 -2007	2003 -2007	2005 -2007
Euro Stoxx	1	1.0108	0.9893	0.9844	1.0146*	0.9924	0.9881
	3	1.0251	0.9637	0.9598	1.0328	0.9713	0.9703
	6	1.0419	0.8867	0.8782	1.0529	0.9009	0.8989
FTSE	1	1.0073	0.9901	0.9891	1.0061	0.9908	0.9906
	3	1.0185	0.9567	0.9644	1.0153	0.9601	0.9692
	6	1.0287	0.8884*	0.8974	1.0231	0.8985*	0.9101
DAX	1	1.0038	0.9719	0.9636	1.0072	0.9730	0.9654
	3	1.0098	0.9268*	0.9175	1.0179	0.9264	0.9179
	6	1.0137	0.8504*	0.8387*	1.0279	0.8388*	0.8252*
CAC 40	1	1.0104	0.9820	0.9789	1.0141	0.9896	0.9872
	3	1.0259	0.9493	0.9541	1.0338	0.9686	0.9752
	6	1.0438	0.8804	0.8980	1.0552	0.9176	0.9451
Dow Jones	1	1.0075	1.0005	1.0043	1.0106	1.0027	1.0071
	3	1.0208	0.9940	1.0026	1.0290	1.0000	1.0115
	6	1.0432	0.9931	0.9958	1.0575	1.0091	1.0174
Nasdaq	1	1.0143	1.0027	1.0256	1.0204	1.0014	1.0311
	3	1.0444	1.0023	1.0485	1.0630	0.9977	1.0573
	6	1.0961	1.0073	1.0895	1.1352	0.9980	1.0967
S&P 500	1	1.0083	0.9972	1.0060	1.0113	0.9986	1.0092
	3	1.0214	0.9858	1.0078	1.0285	0.9889	1.0172
	6	1.0394	0.9686	1.0087	1.0512	0.9783	1.0362
Nikkei	1	0.9846	0.9646*	0.9659	0.9918	0.9767*	0.9783
	3	0.9690	0.9339*	0.9389	0.9838	0.9553	0.9606
	6	0.9562	0.9144*	0.9268	0.9782	0.9416	0.9534
HangSeng	1	1.0103	0.9697	0.9629	1.0035	0.9563	0.9500
	3	1.0142	0.9208	0.9041	0.9999	0.9023	0.8912
	6	0.9911	0.8854	0.8743	0.9711	0.8735	0.8705

Note: AR order selection based on SC with maximum lag order of 4.

* significant at 5% level according to DM test with two-sided alternative.

Table 6: Forecast MSEs of Naive Forecast Relative to Linear Forecast for GDP Series

Country	Forecast horizon	Sample 1980Q1- Forecast period		Sample 1985Q1- Forecast period		Sample 1990Q1- Forecast period	
		2000Q1 -2006Q4	2003Q1 -2006Q4	2000Q1 -2006Q4	2003Q1 -2006Q4	2000Q1 -2006Q4	2003Q1 -2006Q4
Belgium	1	0.9485	0.5219	0.9001	0.5174	0.8489	0.5946*
	2	0.8919	0.2939	0.8137	0.2753*	0.7276	0.3925*
	3	0.8523	0.1764*	0.7643	0.1723*	0.6601	0.3206*
	4	0.8242	0.0930*	0.7270	0.0975*	0.5992	0.2625*
Canada	1	0.7016	0.6226	0.6432	0.5515	0.6706	0.5563
	2	0.7564	0.6445	0.6795	0.5402	0.7579	0.6234
	3	0.7901	0.6254	0.6961	0.5243	0.8114	0.7011
	4	0.7874	0.5622	0.6728	0.4654	0.7792*	0.6444*
Denmark	1	1.2308	1.1222	1.0483	0.9875	1.0289	1.0135
	2	1.2764	0.8814	1.0070	0.7981	0.9676	0.8285
	3	1.3627	0.5792	0.9461	0.5125	0.8491	0.5779
	4	1.5162	0.3651	0.9429	0.2850	0.8278	0.4218
France	1	1.0624	0.6920	0.9574	0.7469	0.8451	0.6874*
	2	1.1257	0.3834	0.7983	0.3775	0.6875*	0.5169*
	3	1.2773	0.3327	0.6850	0.2275	0.5624*	0.4211*
	4	1.2678	0.3133	0.5927	0.2062	0.4953	0.3997*
Japan	1	1.0148	1.0348	1.0146*	1.0245	1.0272	1.0096
	2	1.0486	1.0794	1.0289*	1.0683	1.0481	0.9976
	3	1.0836	1.1142	1.0435*	1.0923*	1.0655	0.9914
	4	1.1117	1.1486	1.0536*	1.1117	1.0770	0.9592
Norway	1	0.9578	0.8317	1.0816	0.9980	1.0772	0.9370
	2	0.9781	0.7438	1.1390	0.9932	1.0890	0.8897
	3	0.8681	0.5344*	0.9847	0.7664	0.9875	0.7507*
	4	0.8272	0.4596*	0.8818	0.6703*	0.9032	0.7048*
U.S.	1	1.1019	0.6882	1.0549	0.6883	1.0632	0.7690
	2	0.9840	0.2254*	1.1369	0.3975*	0.9948	0.5648
	3	0.9916	0.1331*	1.0530	0.3122*	0.9016	0.4532
	4	0.9396	0.1084*	0.9720	0.2683*	0.8134	0.4046

Note: AR order selection based on SC with maximum lag order of 8.

* significant at 5% level according to DM test with two-sided alternative.

Table 7: Forecast MSEs of Naive Forecast Relative to Linear Forecast for Consumption Series

Country	Forecast horizon	Sample 1980Q1- Forecast period		Sample 1985Q1- Forecast period		Sample 1990Q1- Forecast period	
		2000Q1 -2006Q4	2003Q1 -2006Q4	2000Q1 -2006Q4	2003Q1 -2006Q4	2000Q1 -2006Q4	2003Q1 -2006Q4
Australia	1	0.9167	0.8351	0.8728	0.7471	0.6223*	0.5081
	2	0.8236	0.8079	0.7022	0.6413	0.4621*	0.4035
	3	0.7521	0.7532	0.5694	0.5207	0.3634*	0.3111
	4	0.7554	0.7165	0.4947	0.4484	0.2861*	0.2591
Belgium	1	0.8823	0.7683	0.9443	0.8812	0.8619	0.6263
	2	0.8977	0.7138	0.8986	0.7358	0.8050	0.5908
	3	0.9512	0.5009	0.8593	0.4859	0.7923	0.4352
	4	1.0271	0.2722	0.8319	0.2834	0.7446	0.2947
Canada	1	0.7766	0.8363	0.7314	0.6338	0.7396	0.6342
	2	0.6130	0.5491	0.5810	0.3672	0.6358	0.5076
	3	0.5536	0.3767	0.4479	0.2186	0.5428	0.4142
	4	0.5103	0.2707	0.3494	0.1382	0.4661*	0.3487
Japan	1	1.0403	1.0942*	1.0457	1.0658	1.0637*	1.0717
	2	1.0526	1.0952	1.0537	1.0738	1.1064*	1.0897
	3	1.1007*	1.1765*	1.0796*	1.1343*	1.1554*	1.1534
	4	1.1184*	1.2665*	1.0994*	1.2002*	1.2042*	1.2254
Norway	1	1.0253	0.9317	0.8493	0.8094	0.9002	0.7414
	2	1.0274	0.7857	0.6680	0.5676	0.6204	0.4870
	3	0.9936	0.4574	0.4634	0.2760*	0.4595	0.3083*
	4	0.8791	0.3241	0.3699	0.1871*	0.3540	0.2199*
UK	1	1.8154	2.0701	1.8468	2.0323	1.5834	2.5313
	2	3.3214*	3.3397	3.5009	2.5500	2.7726*	3.1062
	3	4.0529*	5.7324*	4.1098*	3.3790*	3.0556*	3.6966
	4	5.2223*	7.8226*	5.1315*	4.6133*	3.7520*	4.6226*
U.S.	1	1.0235	0.7678	1.0685	0.6963	1.1045	0.8015
	2	0.9416	0.4251	0.9251	0.4385	0.9587	0.5327*
	3	0.8956	0.2590	0.8348	0.3175	0.8810	0.4467
	4	0.8437	0.2879	0.7458	0.3156	0.7302	0.4228

Note: AR order selection based on SC with maximum lag order of 8.

* significant at 5% level according to DM test with two-sided alternative.

CESifo Working Paper Series

for full list see www.cesifo-group.org/wp

(address: Poschingerstr. 5, 81679 Munich, Germany, office@cesifo.de)

- 2529 Bas Jacobs, Dirk Schindler and Hongyan Yang, Optimal Taxation of Risky Human Capital, January 2009
- 2530 Annette Alstadsæter and Erik Fjærli, Neutral Taxation of Shareholder Income? Corporate Responses to an Announced Dividend Tax, January 2009
- 2531 Bruno S. Frey and Susanne Neckermann, Academics Appreciate Awards – A New Aspect of Incentives in Research, January 2009
- 2532 Nannette Lindenberg and Frank Westermann, Common Trends and Common Cycles among Interest Rates of the G7-Countries, January 2009
- 2533 Erkki Koskela and Jan König, The Role of Profit Sharing in a Dual Labour Market with Flexible Outsourcing, January 2009
- 2534 Tomasz Michalak, Jacob Engwerda and Joseph Plasmans, Strategic Interactions between Fiscal and Monetary Authorities in a Multi-Country New-Keynesian Model of a Monetary Union, January 2009
- 2535 Michael Overesch and Johannes Rincke, What Drives Corporate Tax Rates Down? A Reassessment of Globalization, Tax Competition, and Dynamic Adjustment to Shocks, February 2009
- 2536 Xenia Matschke and Anja Schöttner, Antidumping as Strategic Trade Policy Under Asymmetric Information, February 2009
- 2537 John Whalley, Weimin Zhou and Xiaopeng An, Chinese Experience with Global 3G Standard-Setting, February 2009
- 2538 Claus Thustrup Kreiner and Nicolaj Verdelin, Optimal Provision of Public Goods: A Synthesis, February 2009
- 2539 Jerome L. Stein, Application of Stochastic Optimal Control to Financial Market Debt Crises, February 2009
- 2540 Lars P. Feld and Jost H. Heckemeyer, FDI and Taxation: A Meta-Study, February 2009
- 2541 Philipp C. Bauer and Regina T. Riphahn, Age at School Entry and Intergenerational Educational Mobility, February 2009
- 2542 Thomas Eichner and Rüdiger Pethig, Carbon Leakage, the Green Paradox and Perfect Future Markets, February 2009
- 2543 M. Hashem Pesaran, Andreas Pick and Allan Timmermann, Variable Selection and Inference for Multi-period Forecasting Problems, February 2009

- 2544 Mathias Hoffmann and Iryna Shcherbakova, Consumption Risk Sharing over the Business Cycle: the Role of Small Firms' Access to Credit Markets, February 2009
- 2545 John Beirne, Guglielmo Maria Caporale, Marianne Schulze-Ghattas and Nicola Spagnolo, Volatility Spillovers and Contagion from Mature to Emerging Stock Markets, February 2009
- 2546 Ali Bayar and Bram Smeets, Economic and Political Determinants of Budget Deficits in the European Union: A Dynamic Random Coefficient Approach, February 2009
- 2547 Jan K. Brueckner and Anming Zhang, Airline Emission Charges: Effects on Airfares, Service Quality, and Aircraft Design, February 2009
- 2548 Dolores Messer and Stefan C. Wolter, Money Matters – Evidence from a Large-Scale Randomized Field Experiment with Vouchers for Adult Training, February 2009
- 2549 Johannes Rincke and Christian Traxler, Deterrence through Word of Mouth, February 2009
- 2550 Gabriella Legrenzi, Asymmetric and Non-Linear Adjustments in Local Fiscal Policy, February 2009
- 2551 Bruno S. Frey, David A. Savage and Benno Torgler, Surviving the Titanic Disaster: Economic, Natural and Social Determinants, February 2009
- 2552 Per Engström, Patrik Hesselius and Bertil Holmlund, Vacancy Referrals, Job Search, and the Duration of Unemployment: A Randomized Experiment, February 2009
- 2553 Giorgio Bellettini, Carlotta Berti Ceroni and Giovanni Prarolo, Political Persistence, Connections and Economic Growth, February 2009
- 2554 Steinar Holden and Fredrik Wulfsberg, Wage Rigidity, Institutions, and Inflation, February 2009
- 2555 Alexander Haupt and Tim Krieger, The Role of Mobility in Tax and Subsidy Competition, February 2009
- 2556 Harald Badinger and Peter Egger, Estimation of Higher-Order Spatial Autoregressive Panel Data Error Component Models, February 2009
- 2557 Christian Keuschnigg, Corporate Taxation and the Welfare State, February 2009
- 2558 Marcel Gérard, Hubert Jayet and Sonia Paty, Tax Interactions among Belgian Municipalities: Does Language Matter?, February 2009
- 2559 António Afonso and Christophe Rault, Budgetary and External Imbalances Relationship: A Panel Data Diagnostic, February 2009
- 2560 Stefan Krassa and Mattias Polborn, Political Competition between Differentiated Candidates, February 2009

- 2561 Carsten Hefeker, Taxation, Corruption and the Exchange Rate Regime, February 2009
- 2562 Jiahua Che and Gerald Willmann, The Economics of a Multilateral Investment Agreement, February 2009
- 2563 Scott Alan Carson, Demographic, Residential, and Socioeconomic Effects on the Distribution of 19th Century US White Statures, February 2009
- 2564 Philipp Harms, Oliver Lorz and Dieter Urban, Offshoring along the Production Chain, February 2009
- 2565 Patricia Apps, Ngo Van Long and Ray Rees, Optimal Piecewise Linear Income Taxation, February 2009
- 2566 John Whalley and Shunming Zhang, On the Arbitrariness of Consumption, February 2009
- 2567 Marie-Louise Leroux, Endogenous Differential Mortality, Non-Contractible Effort and Non Linear Taxation, March 2009
- 2568 Joanna Beza-Bojanowska and Ronald MacDonald, The Behavioural Zloty/Euro Equilibrium Exchange Rate, March 2009
- 2569 Bart Cockx and Matteo Picchio, Are Short-Lived Jobs Stepping Stones to Long-Lasting Jobs?, March 2009
- 2570 David Card, Jochen Kluve and Andrea Weber, Active Labor Market Policy Evaluations: A Meta-analysis, March 2009
- 2571 Frederick van der Ploeg and Anthony J. Venables, Harnessing Windfall Revenues: Optimal Policies for Resource-Rich Developing Economies, March 2009
- 2572 Ondřej Schneider, Reforming Pensions in Europe: Economic Fundamentals and Political Factors, March 2009
- 2573 Jo Thori Lind, Karl Ove Moene and Fredrik Willumsen, Opium for the Masses? Conflict-Induced Narcotics Production in Afghanistan, March 2009
- 2574 Silvia Marchesi, Laura Sabani and Axel Dreher, Agency and Communication in IMF Conditional Lending: Theory and Empirical Evidence, March 2009
- 2575 Carlo Altavilla and Matteo Ciccarelli, The Effects of Monetary Policy on Unemployment Dynamics under Model Uncertainty - Evidence from the US and the Euro Area, March 2009
- 2576 Falko Fecht, Kjell G. Nyborg and Jörg Rocholl, The Price of Liquidity: Bank Characteristics and Market Conditions, March 2009
- 2577 Giorgio Bellettini and Filippo Taddei, Real Estate Prices and the Importance of Bequest Taxation, March 2009

- 2578 Annette Bergemann and Regina T. Riphahn, Female Labor Supply and Parental Leave Benefits – The Causal Effect of Paying Higher Transfers for a Shorter Period of Time, March 2009
- 2579 Thomas Eichner and Rüdiger Pethig, EU-Type Carbon Emissions Trade and the Distributional Impact of Overlapping Emissions Taxes, March 2009
- 2580 Antonios Antypas, Guglielmo Maria Caporale, Nikolaos Kourogenis and Nikitas Pittis, Selectivity, Market Timing and the Morningstar Star-Rating System, March 2009
- 2581 António Afonso and Christophe Rault, Bootstrap Panel Granger-Causality between Government Budget and External Deficits for the EU, March 2009
- 2582 Bernd Süßmuth, Malte Heyne and Wolfgang Maennig, Induced Civic Pride and Integration, March 2009
- 2583 Martin Peitz and Markus Reisinger, Indirect Taxation in Vertical Oligopoly, March 2009
- 2584 Petra M. Geraats, Trends in Monetary Policy Transparency, March 2009
- 2585 Johannes Abeler, Armin Falk, Lorenz Götte and David Huffman, Reference Points and Effort Provision, March 2009
- 2586 Wolfram F. Richter, Taxing Education in Ramsey's Tradition, March 2009
- 2587 Yin-Wong Cheung, Menzie D. Chinn and Eiji Fujii, China's Current Account and Exchange Rate, March 2009
- 2588 Alexander Haupt and Silke Uebelmesser, Voting on Labour-Market Integration and Education Policy when Citizens Differ in Mobility and Ability, March 2009
- 2589 Hans Jarle Kind, Marko Koethenbueger and Guttorm Schjelderup, Should Utility-Reducing Media Advertising be Taxed?, March 2009
- 2590 Alessandro Cigno, How to Avoid a Pension Crisis: A Question of Intelligent System Design, March 2009
- 2591 Helmut Lutkepohl and Fang Xu, The Role of the Log Transformation in Forecasting Economic Variables, March 2009