NON-LINEARITIES IN EXCHANGE RATES

Michael P. Clements and Jeremy Smith

No.504

WARWICK ECONOMIC RESEARCH PAPERS

THE UNIVERSITY OF WARWICK

DEPARTMENT OF ECONOMICS

NON-LINEARITIES IN EXCHANGE RATES

Michael P. Clements and Jeremy Smith
Department of Economics
University of Warwick
Coventry, CV4 7AL
UK

No.504

March 1998

This paper is circulated for discussion purposes only and its contents should be considered preliminary.

Non-linearities in Exchange Rates

Michael P. Clements and Jeremy Smith*
Department of Economics,
University of Warwick.
Coventry CV4 7AL.

March 3, 1998

Abstract

We consider the forecasting performance of two SETAR exchange rate models proposed by Kräger and Kugler (1993). Assuming that the models are good approximations to the data generating process, we show that whether the non-linearities inherent in the data can be exploited to forecast better than a random walk depends on both how forecast accuracy is assessed and on the 'state of nature'. Evaluation based on traditional measures, such as (root) mean squared forecast errors, may mask the superiority of the non-linear models. Generalized impulse response functions are also calculated as a means of portraying the asymmetric response to shocks implied by such models.

1 Introduction

While there is clear evidence of non-linearities in the even-ordered conditional moments of post-War exchange rates (of the form that large /small changes tend to be followed by large /small changes, but of either sign), the evidence for dependence in first moments is less convincing. Thus ARCH (Engle, 1982) and GARCH (Bollerslev, 1986) models have claimed some success for interval prediction, while the literature on conditional mean prediction over the post-war period suggests it is difficult to better a random walk. For example, using a nearest-neighbor technique of locally-weighted regression (see, e.g., Cleveland, Devlin and Grosse, 1988), Diebold and Nason (1990) find no improvement over a simple random walk for predicting ten major dollar exchange rates over the post 1973 period. The use of a non-parametric prediction method is intended to guard against the failure to benefit from non-linearities due to the incorrect choice of functional form. From a different standpoint, Meese and Rose (1991) allow for non-linear extensions to a number of structural exchange rate models but with no significant improvement in forecast accuracy.

However, Kräger and Kugler (1993) find that self-exciting threshold autoregressive models (SETAR)[1] are able to describe the behaviour of five weekly dollar exchange rate series for 1980 – 90, and suggest that this might be expected in a period of managed floating. They argue that threshold models of this sort may arise (approximately) as the outcome of a rational expectations monetary model with stochastic intervention rules (see Hsieh, 1989). The suggestion is that the authorities react to large appreciations and depreciations (rates of change), whereas

*The first author acknowledges financial support under ESRC grant L116251015. Email: M.P.Clements@warwick.ac.uk and Jeremy.Smith@warwick.ac.uk, tel.: 01203 523055, FAX: 01203 523032.

[1]First proposed by Tong (1978), and subsequently developed in Tong and Lim (1980) and Tong (1983) (see also Tong, 1995a)

for the target zone approach to managed floating the level of the exchange rate (or rather, its proximity to ceilings or floors) is relevant for signalling interventions.

The aim of this paper is to assess the usefulness of their models for out-of-sample forecasting, and to address the closely related issue of the propagation of shocks implied by the models. 'Out-of-sample' is somewhat of a misnomer, in that as we explain in section 2, forecasting performance is assessed by a number of simulation experiments designed to favour the non-linear models. However, it distinguishes our approach from the in-sample evaluation of the models undertaken by Kräger and Kugler (1993), noting, following Diebold and Nason (1990), that a better forecast performance does not automatically follow from the rejection of the linear null in-sample. In passing, we note that Kräger and Kugler (1993) misinterpret their tests for the statistical significance of more than one regime, and in fact there is no evidence against the (linear) null. Their Table 3 which purportedly reports the 'Chan-Tong test for linearity in the estimated residuals of the SETAR models' actually reports the test outcomes for the hypothesis of a single regime (i.e., a linear model) versus more than one regime. We find that whether the non-linear models are capable of out-performing a random walk depends crucially on the way in which forecasts are evaluated. The results of the forecast comparison are presented in sections 2 and 3. The former draws on the method of analysis in Clements and Smith (1996), and the latter evaluates the whole forecast density.

The analysis of the propagation of shocks in section 4 is an application of the generalized non-linear response function analysis of Koop, Pesaran and Potter (1996) to the Kräger and Kugler (1993) models. This analysis allows us to measure the persistence of shocks implied by the empirical models. Of particular interest is the extent to which the non-linearity of the model gives rise to asymmetries, whereby the persistence of shocks depends on the regime in force. Just as the non-linear models of US GNP of Beaudry and Koop (1993), Potter (1995) and Pesaran and Potter (1997), for example, imply different degrees of persistence between shocks occurring in recessions and booms, so the exchange rate models imply asymmetric responses to shocks occurring at times of large appreciations (regime 1), large depreciations (regime 3) and more settled times (regime 2). Finally, section 5 briefly summarises our findings.

2 Forecasting performance of non-linear models relative to a random walk

We assess the forecast performance of the SETAR models relative to linear autoregressive (AR) alternatives by means of Monte Carlo simulations specifically designed to cast the non-linear models in their best possible light. By taking the estimated SETAR models as the data generating processes (DGPs) we are able to ensure that the future has the same non-linear imprint as the past, so that a poor performance of the non-linear models can not be attributed to a failure of the non-linearity to persist into the future (e.g., Granger and Teräsvirta, 1993, p.164, Teräsvirta and Anderson, 1992).

It is often claimed that 'how well we can predict depends on where we are' and that particularly for non-linear models there might be 'windows of opportunity for substantial reduction in prediction errors' (Tong, 1995b p. 409 – 410). In an empirical study of US GNP Tiao and Tsay (1994) find that the forecast performance of SETAR and AR models is broadly similar for data points in the 'boom' (as opposed to 'recession') regime but that the SETAR model records significant gains relative to the linear model at times when the economy is recovering from recession. In general, without large samples of historical data there may be few episodes of the phenomenon of interest and estimates of the gains achievable during these periods may lack precision. Our simulation approach allows us to better explore the dependence of forecast performance on the regime at the forecast origin.

Since regime-switching models may be better suited to predicting movements between regimes, rather than small movements within a regime, we employ forecast evaluation criteria based on the number of times the models correctly predict the right regime. The 'Non-Confusion' Rate (NCR) is the number of times the regimes are correctly predicted divided by the total number of predictions. We also report conditional probabilities of correctly predicting each regime, defined as the number of times the model correctly predicts a regime divided by the number of times the process was in that regime. These are referred to as CRPs – conditional regime predictions. A value of unity indicates that the regime is always correctly predicted (but note some caution is required – such an outcome would arise if the model always predicted that regime regardless of which regime the process is actually in). We report CRPs for the AR and SETAR models without conditioning on being in a particular regime.

Nevertheless, such criteria count one-for-one a forecast of a very small increase, when in fact a small decline occurred, with a forecast of a large increase when a large decline occurred, (if the regimes are positive and negative growth, say), so that we complement such measures with more traditional squared error loss measures, such as the mean squared forecast error (MSFE), for each horizon, using the variability over the replications of the Monte Carlo to obtain our estimates. We also calculate a test proposed by Diebold and Mariano (1995) of whether the differences in MSFEs between models are statistically significant.

Full details of the way in which the Monte Carlo was carried out are given in Clements and Smith (1996), but briefly, on each replication of the Monte Carlo data is generated from the empirical SETAR exchange rate models. We then estimate SETAR and AR models, holding back observations for which forecasts are then computed. We then repeat, and build up a Monte Carlo sample of multi-step forecast errors. The SETAR model forecasts can be generated under a number of assumptions about what is known about the model – in one extreme the model is assumed known and to coincide with the DGP (e.g., the number of regimes, lag orders, autoregressive coefficients and threshold values are correctly specified), whereas at the other a model selection strategy is adopted, assuming only the correct number of regimes. In the discussion of the results the former is referred to as the 'Known Model' case, and for the latter, 'Unknown Model' is used as a convenient shorthand.

The exchange rate models estimated by Kräger and Kugler (1993) are given in their Table 2, p. 200-2 for the German mark, the French franc, the Italian lira, the Japanese yen and the Swiss franc against the US dollar. To economise on space, we report results for the yen and deutschemark only. The models have three regimes and a delay of one period (so that the current regime at t is determined by the value of the process at $t-1$ in relation to the threshold values). For both the yen and deutschemark the middle regime is a third-order AR in the difference of the log of the exchange rate, and for the yen the first and third set the growth rate equal to a constant, whereas the mark differs in that the first regime is a second-order AR. The estimated standard deviations of the first and third regimes exceed that of the middle regime, which is explained by central bank interventions in response to large appreciations (regime 1) and depreciations (regime 3).

Table 1 summarises the results of our Monte Carlo evaluation of the multi-step forecast performance of the SETAR yen and deutschemark exchange rate models. The linear competitor is an AR(0) model with a constant for the differences (of the logs) – i.e., a random walk with drift. Higher-order linear models were generally dominated by the random walk (RW), consistent with the view that changes are not serially correlated.

Panel [A] reports, for the 'Known SETAR Model' case, unconditional (on the regime) MSFEs for the RW model divided by those for SETAR for 1, 2 and 5-steps ahead. For longer horizons the ratio is approximately unity. Panel [B] reports the ratio of the MSFEs, again for the 'Known SETAR Model' case, but this time conditioning the forecast origin on each of the three regimes: Lower, Middle and Upper. The p-values are of the Diebold and Mariano (1995) test of equal forecast accuracy (as measured by MSFE), and are the probabilities under the null (of equal accuracy) of obtaining lower test statistics than we record. The CRPs in panel [A] record the proportion of times the regimes were correctly predicted when we do not condition on the regime. The NCRs in panel [A] are the non-confusion rates. Finally, panels [C] and [D] repeat the information in panels [A] and [B] for the case when the SETAR model is not known. The CRPs and NCRs for the AR model are the same by construction in panels [A] and [C] and hence are reported only once.

Japanese Yen. There is a gain of over 40%, conditional on being in the Middle regime, when the SETAR model is known, and of 15% when it is estimated. However, even conditionally there is nothing to choose between the models at further steps ahead, on the basis of MSFE. The CRPs indicate that the SETAR is better at predicting the Middle regime even at 5 steps ahead, getting it right nearly 40% of the time. The NCRs indicate that the SETAR is less confused at 1-step ahead. For the Japanese Yen the qualitative and quantitative measures tell a similar story. If the yen exchange rate was generated by the proposed empirical SETAR model, then the inherent non-linearities could be exploited to do better than a random walk.

German mark. The same is not true for the mark. Relative to a RW there is nearly a 5% gain at $h = 1$ unconditionally when the SETAR model is known, but in the absence of this knowledge of the DGP, one would not be able to do better than a random walk. Then, even conditional on being in the middle regime, the SETAR is not statistically more accurate on MSFE at the 10% level, judged by the Diebold and Mariano (1995) test.

3 Evaluating density forecasts

In this section we report the results of a comparison of the SETAR and random walk models via an evaluation of the complete 1 and 2-step ahead forecast densities associated with each. The approach is due to Diebold, Gunther and Tay (1997), and amounts to calculating the probability integral transforms of the 'actual' values of the variable, over the forecast period, (e.g., $\{y_t\}_{t=1}^n$, $t = 1, \ldots, n$) with respect to the forecast densities of the SETAR and AR models, denoted by $\{p_t(y_t)\}_{t=1}^n$. That is, we calculate:

$$\{z_t\}_{t=1}^n = \left\{ \int_{-\infty}^{y_t} p_t(u)du \right\}_{t=1}^n . \tag{1}$$

When the model forecast density corresponds to the true predictive density (given by the DGP, and denoted by $f_t(y_t)$), i.e., $p_t(y_t) = f_t(y_t)$, then Diebold *et al.* (1997) show that $\{z_t\}_{t=1}^n \sim$ iid$U[0,1]$. The result that $z_t \sim U[0,1]$ can be found in, e.g., Kendall, Stuart and Ord (1987, sections 1.27 and 30.36). Diebold *et al.* (1997) make the result operational in the time series context by establishing that the z_t sequence are independent when the true densities are used at each t. Hence the idea is to evaluate the forecast density by assessing whether there is statistically significant evidence that the realizations do not come from that density – this amounts to testing whether the $\{z_t\}$ series depart from the iid uniform assumption.

The SETAR model densities are simulated, using 500 replications with gaussian errors, on each of the 1000 replications of the Monte Carlo. The z's are calculated for the 'Unknown Model' case.

We assess uniformity (conditional on the iid part) by plotting in figure 1 the actual cdf of the z_t's against the theoretical cdf for the 1-step ahead forecasts. The 95% confidence intervals drawn alongside the 45° lines are based on critical values tabulated by Miller (1956). [2]

Figure 1 shows clear violations of the uniformity assumptions for the RW models of the yen and DM (although visually this is less apparent for the latter due to scaling), while there is no evidence against the SETAR models. Similar figures were plotted for 2-step ahead forecasts, but in no instance was there any evidence against the null. So while there are minimal gains to the SETAR relative to the RW models on MSFE, the two models can be distinguished on the basis of their overall forecast densities.

4 Non-linear impulse response function analysis

Impulse response analysis traces the impact of a shock through time, assuming the model is correct. Consequently, this analysis parallels forecasting in the 'Known Model' case. In this section we illustrate the response to shocks implied by the SETAR models, and how these differ from those for the linear (random walk) models. Koop *et al.* (1996) develop generalized non-linear impulse response functions (GIs) (see also Gallant, Rossi and Tauchen, 1993) to analyse the response of non-linear models to shocks. Their analysis recognises that the impact of the shock is dependent upon the sign and size of the shock, and the position of the process when the shock hits. To illustrate, figure 2 plots the standard impulse responses for the empirical models of the yen and deutschemark estimated by Kräger and Kugler (1993), for both positive and negative one standard deviation shocks at time t. As can be seen, the responses depend on the regime the process is in at time $t-1$, and the positive and negative shocks are not mirror images. Different patterns may emerge if the magnitude of the shock is scaled up (or down).

Because linear response analysis is inappropriate, we analyse the distribution of responses using GIs. Linear impulse responses follow immediately from the moving-average representation of the process, but exact analytical representations are not available for SETAR processes. In the previous section we calculated multi-step forecasts by Monte Carlo, to allow for non-zero values of the future disturbances (see, e.g., Granger and Teräsvirta, 1993), and it is important to also integrate out the effects of future shocks in the present context.

The estimated time series model is used to produce simulated realisations of the series, $\{y_1, \ldots, y_{t-1}\}$. The maximum horizon of interest is assumed to be 5, so we analyse the response of the process in periods $\{t+1, \ldots, t+5\}$ to the shock ν_t, applied in period t. The shock, ν_t, and all future innovations for $\{t+1, \ldots, t+5\}$ are assumed to be drawn from a normal distribution,

[2] Miller (1956) reports exact critical values of Kolmogorov Statistics for small sample sizes, n. We have $n = 1000$, so use the asymptotic critical values of Smirnov reported in Miller (1956, eqn.3, p.115) – $\overline{(\ln(1/\alpha)/2 \times n)} = 0.0429$ for a 95% confidence level ($\alpha = 0.025$). The 95% confidence intervals are then the 45° line ± 0.0429.

with mean zero and variance given by the pooled variance across the three regimes. For a given shock, ν_t, the shocked series is obtained by averaging across all possible future realisations (estimated by $R = 500$ simulated paths) of the non-linear process, so the shocked series is obtained as:

$$\overline{y}^S_{R,t+i}(\nu_t, \omega_{t-1}) = \frac{1}{R}\sum_{j=1}^{R} y^j_{t+i}(\nu_t, \omega_{t-1}), \quad i = 1, \ldots, 5$$

where the notation denotes that we are conditioning on the shock ν_t and the past set of disturbances ω_{t-1}, and the variation over j arises from the different future sets of disturbances. The impulse response function is then obtained by comparing the shocked series with a base simulation. The base simulation uses the same innovations as $\overline{y}^S_{R,t+i}(\nu_t, \omega_{t-1})$ but averages across all possible realisations of the shock ν_t, in period t. That is, on each of the j replications, the period t shock is drawn along with the future shocks:

$$\overline{y}^B_{R,t+i}(\omega_{t-1}) = \frac{1}{R}\sum_{j=1}^{R} y^j_{t+i}(\omega_{t-1}), \quad i = 0, 1, \ldots, 5$$

The difference is then:

$$I_R(\nu_t, \omega_{t-1}) = \overline{y}^S_{R,t+i}(\nu_t, \omega_{t-1}) - \overline{y}^B_{R,t+i}(\omega_{t-1}), \quad i = 0, \ldots, 5,$$

which constitutes a single realisation of the GI. The Monte Carlo estimate of the distribution of the GI is given by repeating the above 1000 times, where on each replication we make a different drawing of $\{\nu_t, \omega_{t-1}\}$, and then another R drawings of the current shock and future disturbances.

The distributions of the GI for 1, 3 and 5-steps ahead are reported in figures 3 and 4 for the yen and deutschemark, respectively. The unconditional distributions as well as the distributions conditional on the process being in each of the three alternative regimes at time $t-1$ are also reported. In each case, as described above, we are assessing the persistence of the shock for all possible histories (unconditional), or for all histories for which $t-1$ falls in a particular regime (conditional). The Epanechnikov kernel, defined as:

$$K(z_t - z_j) = \begin{cases} \frac{3}{4\sqrt{5}}\{(z_t - z_j)^2/5\} & -\sqrt{5} \leq z_t - z_j \leq \sqrt{5} \\ 0 & otherwise \end{cases}$$

was used to estimate the graphed densities.

The impulse responses are calculated for the first differences of the logarithms of the exchange rates, so that the relevant linear model is of white noise plus a constant, for which the entire distribution of the GI is at zero for all horizons. The GIs for the non-linear models become less dispersed at 3, and then again at 5-steps ahead, compared to at 1-step. If the process were stationary then in the limit the distribution becomes concentrated on zero – see Potter (1994). Longer horizons would allow us to explore this further, which might be of interesting given that stationarity conditions for general SETAR models are not known (see, e.g., De Gooijer and De Bruin, 1997 for a review of what is known).

For the yen, 1-step GIs are bi-modal in the extreme regimes, but not in the middle regime, and the lower regime has a long left tail, and the upper a long right tail. By 3-steps ahead, the GIs are unimodal and appear symmetric and very similar no matter where the process was at $t-1$. By contrast, even at 3-steps ahead the lower regime GI for the deutschemark looks quite different than those for the other regimes, and all appear asymmetric.

5 Conclusions

We have shown that the Kräger and Kugler (1993) SETAR model of the dollar – yen exchange rate, if taken to be a good approximation to the process that generated the data, would suggest some predictability in the exchange rate over the period of managed floating in the 1980s. Our Monte Carlo study has quantified the size of the potential gains on MSFE, and the extent to which the achievable gains depend on where the process is at the time of forecasting. We have also shown that an evaluation of the whole forecast density may reveal gains to the non-linear model which are not apparent on MSFE measures.

In this paper the non-linear model is the data generating process by construction. The question we posed is whether the non-linearity is sufficiently pronounced that a non-linear model that has to be specified and estimated from the data can better a linear model. Judged by conventional criteria, the answer would appear to depend on where the process is at the time of forecasting. Evaluation of the whole density of the two models indicates a clearer discrimination may be possible.

The information recorded in this paper complements that in the original study, which concentrated on in-sample testing of the models and any residual (higher-moment) non-linearities in the models' residuals. As we noted, judged by some of those criteria the models have little to commend them.

The generalized impulse response analysis of Koop *et al.* (1996) was applied to the empirical SETAR models. The non-linearities were apparent in the dependence of the shapes of the estimated densities on the regime.

References

Beaudry, P., and Koop, G. (1993). Do recessions permanently affect output. *Journal of Monetary Economics*, **31**, 149–163.

Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics*, **51**, 307–327.

Clements, M. P., and Smith, J. (1996). A Monte Carlo study of the forecasting performance of empirical SETAR models. Warwick Economic Research Papers No. 464, Department of Economics, University of Warwick.

Cleveland, W. S., Devlin, S. J., and Grosse, E. (1988). Regression by local fitting: Methods, properties and computational algorithms. *Journal of Econometrics*, **37**, 87–114.

De Gooijer, J. G., and De Bruin, P. (1997). On SETAR forecasting. *Statistics and Probability Letters*. Forthcoming.

Diebold, F. X., Gunther, T. A., and Tay, A. S. (1997). Evaluating density forecasts. Manuscript, Department of Economics, University of Pennsylvania.

Diebold, F. X., and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, **13**, 253–263.

Diebold, F. X., and Nason, J. A. (1990). Nonparametric exchange rate prediction. *Journal of International Economics*, **28**, 315–332.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity, with estimates of the variance of United Kingdom inflation. *Econometrica*, **50**, 987–1007.

Gallant, A. R., Rossi, P. E., and Tauchen, G. (1993). Nonlinear dynamic structures. *Econometrica*, **61**, 871–907.

Granger, C. W. J., and Teräsvirta, T. (1993). *Modelling Nonlinear Economic Relationships.* Oxford: Oxford University Press.

Hsieh, D. A. (1989). A non-linear stochastic rational expectations model of exchange rates. Unpublished manuscript, Fuqua School of Business, Duke University.

Kendall, M. G., Stuart, A., and Ord, J. K. (1987). *Advanced Theory of Statistics,* 5th edn., Vol. 1 and 2. London: Charles Griffin and Co.

Koop, G., Pesaran, M. H., and Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, **74**, 119–147.

Kräger, H., and Kugler, P. (1993). Non-linearities in foreign exchange markets: a different perspective. *Journal of International Money and Finance*, **12**, 195–208.

Meese, R. A., and Rose, A. K. (1991). An empirical assessment of non-linearities in models of exchange rate determination. *Review of Economic Studies*, **58**, 603–619.

Miller, L. H. (1956). Table of percentage points of Kolmogorov statistics. *Journal of the American Statistical Association*, **51**, 111–121.

Pesaran, M. H., and Potter, S. M. (1997). A floor and ceiling model of US Output. *Journal of Economic Dynamics and Control*, **21**, 661–695.

Potter, S. (1994). Nonlinear impulse response functions. Department of economics working paper, University of California, Los Angeles, CA.

Potter, S. (1995). A nonlinear approach to U.S. GNP. *Journal of Applied Econometrics*, **10**, 109–125.

Teräsvirta, T., and Anderson, H. M. (1992). Characterizing nonlinearities in business cycles using smooth transition autoregressive models. *Journal of Applied Econometrics*, **7**, 119–139.

Tiao, G. C., and Tsay, R. S. (1994). Some advances in non-linear and adaptive modelling in time-series. *Journal of Forecasting*, **13**, 109–131.

Tong, H. (1978). On a threshold model. In Chen, C. H. (ed.), *Pattern Recognition and Signal Processing*, pp. 101–141. Amsterdam: Sijhoff and Noordoff.

Tong, H. (1983). *Threshold Models in Non-Linear Time Series Analysis*: Springer-Verlag, New York.

Tong, H. (1995a). *Non-linear Time Series. A Dynamical System Approach.* Oxford: Clarendon Press. First published 1990.

Tong, H. (1995b). A personal overview of non-linear time series analysis from a chaos perspective. *Scandinavian Journal of Statistics*, **22**, 399–445.

Tong, H., and Lim, K. S. (1980). Threshold autoregression, limit cycles and cyclical data. *Journal of The Royal Statistical Society*, B **42**, 245–292.

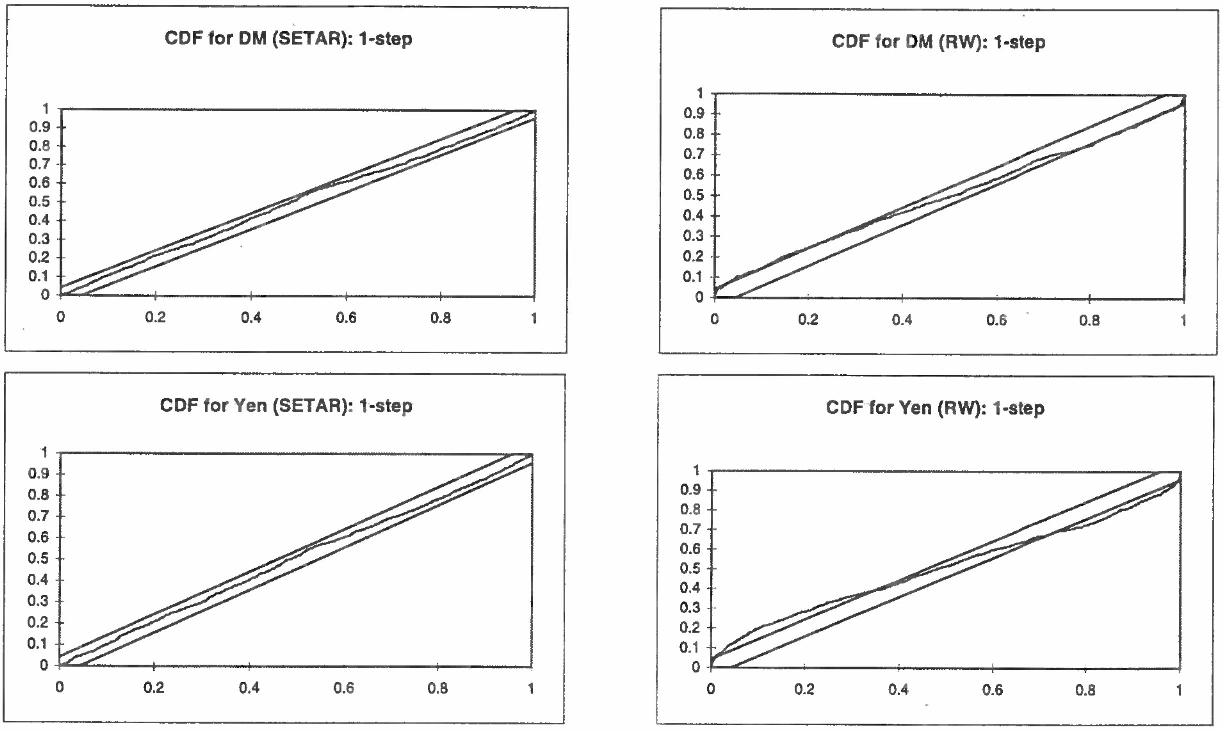

Figure 1 CDF's for probability integral transforms.

Table 1 Forecast evaluation of Kräger and Kugler (1993) exchange rate models.

Japanese Yen

[A] Known SETAR Model – Unconditional

Horizon	MSFE : AR/SETAR		CRPs : AR			CRPs : SETAR			NCRs	
	RW	p-value	Lower	Middle	Upper	Lower	Middle	Upper	AR	SETAR
1	1.046	0.000	0.634	0.344	0.000	0.643	0.521	0.069	.367	.415
2	1.003	0.340	0.647	0.333	0.000	0.611	0.455	0.000	.372	.370
5	0.997	0.774	0.643	0.318	0.000	0.663	0.374	0.000	.347	.363

[B] Known SETAR Model – Conditional. MSFE : AR/SETAR

	Lower		Middle		Upper	
	RW	p-value	RW	p-value	RW	p-value
1	1.028	0.000	1.411	0.000	1.024	0.010
2	1.006	0.141	1.005	0.210	1.001	0.438
5	1.011	0.021	1.004	0.198	0.999	0.605

[C] Unknown SETAR Model – Unconditional

Horizon	MSFE : AR/SETAR		CRPs : SETAR			NCR
	RW	p-value	Lower	Middle	Upper	SETAR
1	1.011	0.192	0.603	0.469	0.072	.390
2	0.993	0.877	0.602	0.402	0.011	.362
5	0.995	0.958	0.620	0.393	0.000	.344

[D] Unknown SETAR Model – Conditional. MSFE : AR/SETAR

	Lower		Middle		Upper	
	RW	p-value	RW	p-value	RW	p-value
1	1.001	0.466	1.149	0.000	0.994	0.673
2	1.000	0.478	1.005	0.204	1.003	0.271
5	1.004	0.085	0.999	0.594	1.000	0.557

Deutschemark

[A] Known SETAR Model – Unconditional

Horizon	MSFE : AR/SETAR		CRPs : AR			CRPs : SETAR			NCRs	
	RW	p-value	Lower	Middle	Upper	Lower	Middle	Upper	AR	SETAR
1	1.048	0.000	0.000	1.000	0.000	0.014	0.995	0.000	.556	.557
2	1.010	0.109	0.000	1.000	0.000	0.000	1.000	0.000	.567	.567
5	0.997	0.734	0.000	1.000	0.000	0.000	1.000	0.000	.544	.544

[B] Known SETAR Model – Conditional. MSFE : AR/SETAR

	Lower		Middle		Upper	
	RW	p-value	RW	p-value	RW	p-value
1	1.063	0.000	1.073	0.000	1.001	0.426
2	1.015	0.029	1.005	0.254	1.017	0.008
5	1.005	0.130	1.002	0.346	0.996	0.816

[C] Unknown SETAR Model – Unconditional

Horizon	MSFE : AR/SETAR		CRPs : SETAR			CRP
	RW	p-value	Lower	Middle	Upper	SETAR
1	0.978	0.913	0.024	0.829	0.007	.551
2	0.987	0.950	0.002	1.000	0.000	.567
5	0.995	0.945	0.000	1.000	0.000	.544

[B] Unknown SETAR Model – Conditional. MSFE : AR/SETAR

	Lower		Middle		Upper	
	RW	p-value	RW	p-value	RW	p-value
1	1.018	0.104	1.019	0.124	0.969	0.997
2	0.987	0.958	1.004	0.287	0.995	0.695
5	0.999	0.647	1.001	0.255	0.997	0.808

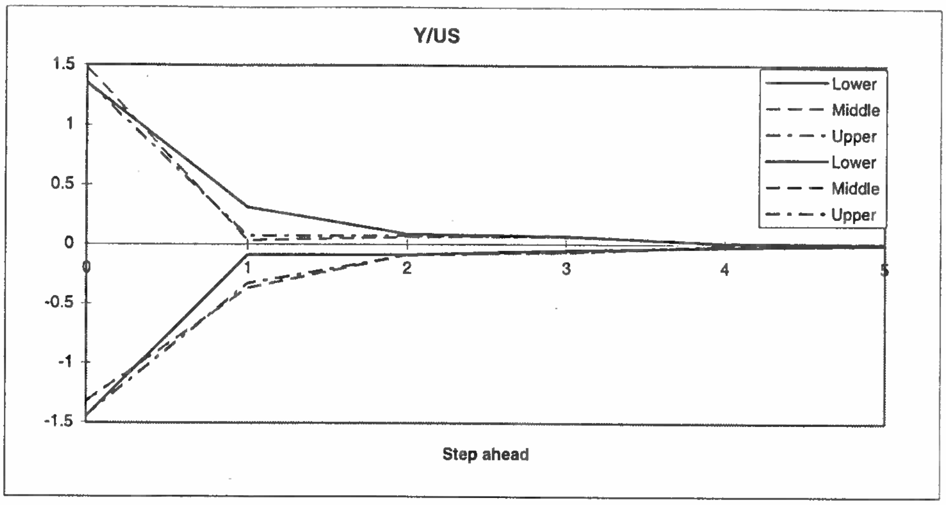

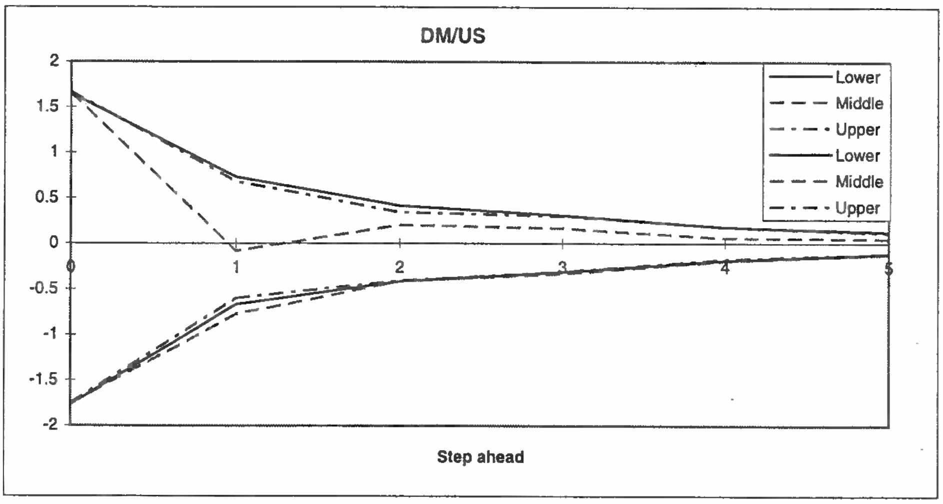

Figure 2 Traditional impulse responses to positive and negative one standard deviation shocks.

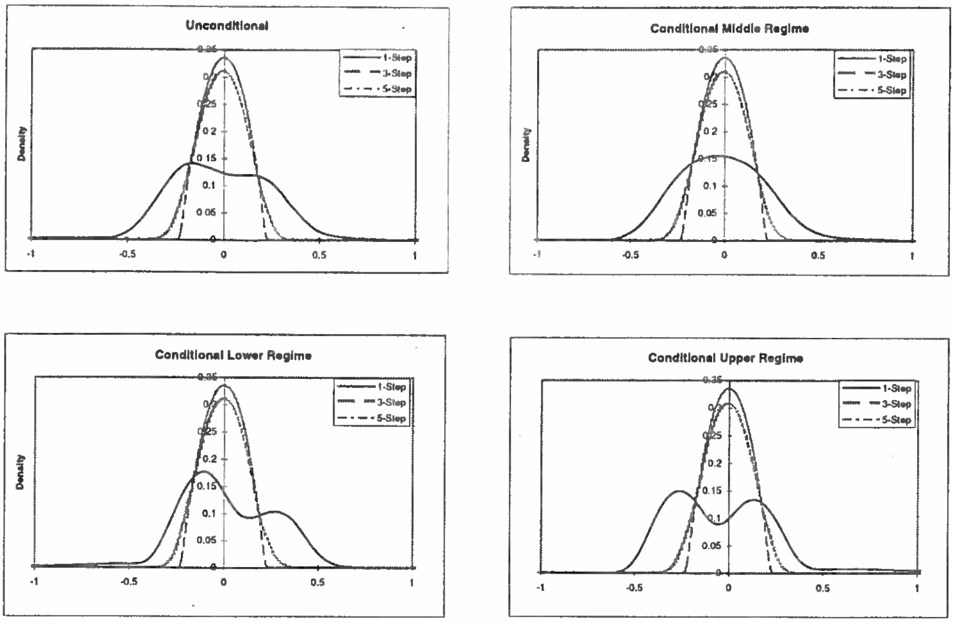

Figure 3 Generalized impulse responses for the Y/US.

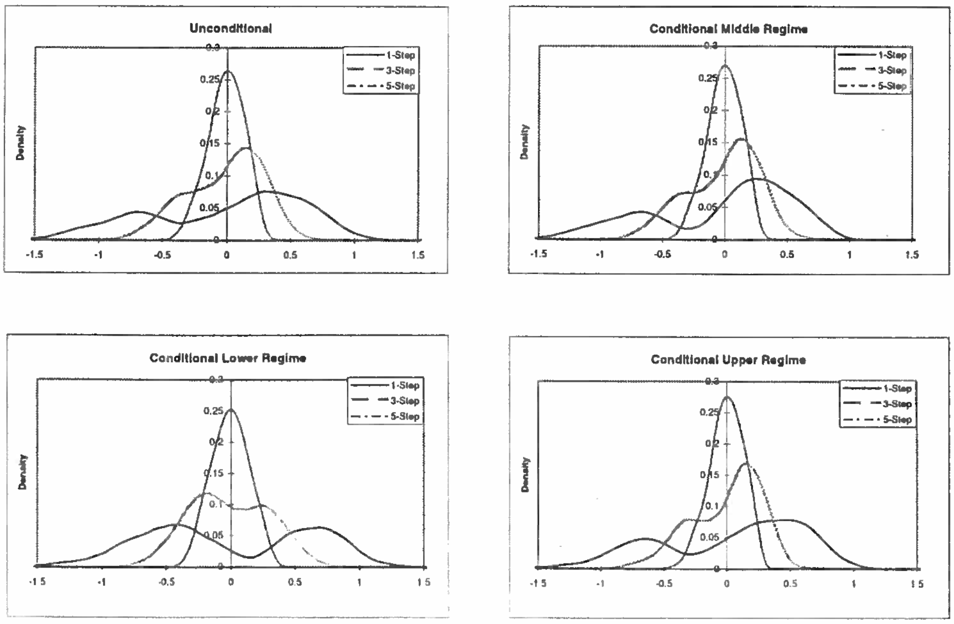

Figure 4 Generalized inpulse responses for the DM/US.