Bos, Charles S.; Mahieu, Ronald J.; van Dijk, Herman K.

Working Paper

On the Variation of Hedging Decisions in Daily Currency Risk Management

Tinbergen Institute Discussion Paper, No. 01-018/4

Provided in Cooperation with:
Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Bos, Charles S.; Mahieu, Ronald J.; van Dijk, Herman K. (2001) : On the Variation of Hedging Decisions in Daily Currency Risk Management, Tinbergen Institute Discussion Paper, No. 01-018/4, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:
https://hdl.handle.net/10419/86108

TI 2001-018/4
Tinbergen Institute Discussion Paper

On the Variation of Hedging Decisions in Daily Currency Risk Management

Charles S. Bos
Ronald J. Mahieu
Herman K. van Dijk

Tinbergen Institute
The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Keizersgracht 482
1017 EG Amsterdam
The Netherlands
Tel.: +31.(0)20.5513500
Fax: +31.(0)20.5513555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31.(0)10.4088900
Fax: +31.(0)10.4089031

Tinbergen Institute Discussion Paper TI-01-018/4

On the Variation of Hedging Decisions in Daily Currency Risk Management*

Charles S. Bos †
Econometric and Tinbergen Institutes

Ronald J. Mahieu
Rotterdam School of Management

Herman K. van Dijk
Erasmus University Rotterdam, The Netherlands

November 3rd, 2000

Abstract

Internationally operating firms naturally face the decision whether or not to hedge the currency risk implied by foreign investments. In a recent paper, Bos, Mahieu and van Dijk (2000) evaluate the returns from optimal and alternative currency hedging strategies, for a series of 7 models, using Bayesian inference and decision analysis. The models differ in the way time-varying means, variances or the unconditional error distributions are incorporated. In this extension, we compare the hedging decisions and financial returns and utilities as they result from the modelling assumptions and the attitudes towards risk.

JEL classification: C11, C44, E47, G15
Keywords: Exchange rates, risk management, Bayesian analysis

*This is a preprint of an article accepted for publication in the proceedings of ISBA 2000, the Sixth World Meeting of the International Society for Bayesian Analysis

†Correspondence to Charles S. Bos, Tinbergen Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands. Email: cbos@few.eur.nl

1 Introduction

Every firm with foreign investments faces the risk of a depreciation of the foreign currency, leading to a lower value of the foreign investment expressed in the home currency. In many firms the decision whether or not to hedge currency risk is revised regularly, independently of the original investment decision. In the finance industry this approach to currency hedging is called currency overlay management.

In Bos et al. (2000) a series of models for the DMark/US Dollar daily exchange rate are constructed. Based on the models, optimal hedge ratios for a utility-optimizing currency overlay manager are calculated, and the risk and return of the optimal and several alternative strategies are evaluated using Bayesian inference and decision analysis with Markov chain Monte Carlo techniques. In this brief extension, we describe the sensitivity and variation of the hedging decision in more detail.

Section 2 summarizes the models, the hedging strategy and the data that are used. Also, the main findings from a Bayesian analysis of the posterior distribution of the parameters are reported here. The paper continues in section 3 with a detailed description of the hedging decisions that are taken, with special attention for the link between modelling decisions and the variation of the hedge ratio over the period January 1998 until December 1999. The financial returns and utilities that are derived from applying the hedging strategies are examined both over the first year of the evaluation period and over both years jointly. A summary of the results is given in section 4.

2 Preliminaries

We concentrate on the hedging decision that a manager may take in order to hedge the currency risk. A detailed description of the setting is given in Bos et al. (2000), here we limit ourselves to a basic overview of the modelling framework, and on a summary of the results concerning the Bayesian posterior density.

Let s_{t+1} be the return on the exchange rate S over the time interval $[t, t+1]$, defined as $s_{t+1} = 100 \ln(S_{t+1}/S_t)$. The investor may choose to hedge a fraction $H \in [0, 1]$,[1] leading to a continuously compounded gross hedged currency return $\exp(r_{t+1})$ equal to a weighted average of returns concerning the exchange rate s_{t+1} and the difference between the home and foreign risk free interest rates r_t^h and r_t^f,

$$\exp(r_{t+1}) = (1 - H_t) \exp(s_{t+1}) + H_t \exp(r_t^h - r_t^f). \tag{1}$$

The overlay manager is interested in optimizing his wealth $W_{t+1} = W_t \exp(r_{t+1})$, according to, as we assume here, a power utility function $U(W_{t+1}) = (W_{t+1}^\gamma - 1)/\gamma, \gamma < 1$. For a background on international portfolios and risk, see Jorion (1985).

In order to find the optimal hedging decision we need to derive a predictive density $P(s_{t+1}|\mathcal{I}_t)$, with $\mathcal{I}_t = \{s_t, s_{t-1}, ..\}$, marginal with respect to the posterior density of the vector of parameters θ in the model. Bos et al. (2000) consider 7 models, describing the evolution of the exchange rate return over time. The baseline model is a state space model (see Harvey 1989),

$$s_t = \mu_t + \epsilon_t, \qquad \epsilon_t \sim i.i.d.\ (0, \sigma_{\epsilon,t}^2), \tag{2}$$

$$\mu_t = \rho \mu_{t-1} + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{3}$$

which allows for a varying mean μ_t.[2] We also allow the variance $\sigma_{\epsilon,t}^2$ in the observation equation (2) to be time time varying. Special cases of this general model include: the White Noise model (**WN**), with $\rho = 1, \sigma_\eta = 0, \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$; the Local Level model (**LL**) which equals the **WN** model but with σ_η unrestricted; the Generalized Local Level model (**GLL**), both ρ and σ_η free. Four models allow for different disturbance distributions: the **GLL-Student t** model allows for ϵ_t to be Student

[1] We do not allow the hedging position to exceed the underlying exposure, i.e. $H < 0$ or $H > 1$.

[2] The uncovered interest rate parity, which prescribes to introduce the interest rate differential as the expectation of s_t, does not hold on a daily timescale. The interest rates are introduced in the hedged return equation (1).

t distributed (with ν degrees of freedom); a **GLL-GARCH** model allows for a GARCH evolution of the variances with the ϵ_t normally distributed; a **GLL-GARCH-Student t** model combines the GARCH and the Student t effects; and a GLL-Stochastic Volatility model (**GLL-SV**). The disturbance of the observation equation (2) of the **GLL-SV** model is normally distributed, but with a randomly evolving variance, i.e. $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2 \exp(h_t))$ and $h_t = \phi h_{t-1} + \xi_t, \xi_t \sim \mathcal{N}(0, \sigma_\xi^2)$.

Our inference is Bayesian. We compute posterior distributions of the parameters through Markov chain Monte Carlo sampling. For the models without Student t-distributed disturbances or Stochastic Volatility, a Metropolis-Hastings sampler can be applied (see e.g. Chib and Greenberg 1995). For the other models, a Gibbs chain was constructed, sampling successively from the full conditional distributions (Carter and Kohn 1994, Kim, Shephard and Chib 1998, Koop and van Dijk 2000). In constructing the posterior, care was taken to use only mildly informative conjugate priors. We take the sample distribution of the parameters into account in the analysis of the predictive density and the hedging decision.

Before we can evaluate the expected utility, we need to know the predictive density of the exchange rate returns. The predictive density $P(s_{t+1}|\mathcal{I}_t)$ was constructed by integrating the conditional predictive density $P(s_{t+1}|\theta, \mathcal{I}_t)$ with respect to the posterior $\theta|\mathcal{I}_t$[3] (see Geweke (1989) and Bauwens, Bos and van Dijk (1999)). Effectively, we calculate

$$P(s_{t+1}|\mathcal{I}_t) \approx \frac{1}{N} \sum P(s_{t+1}|\theta^{(i)}, \mathcal{I}_t) \tag{4}$$

with $\theta^{(i)}$ the i-th drawing of the vector of parameters in the model. This predictive density is calculated over a fine grid of values of s_{t+1}, and is used in optimizing the expected utility $E_{s_{t+1}|\mathcal{I}_t}(U(W_{t+1}))$. It is sufficient to optimize the expected utility using only the wealth increases $w_{t+1} = W_{t+1}/W_t = \exp(r_{t+1})$. The function to be optimized is

$$E_{s_{t+1}|\mathcal{I}_t}(U(w_{t+1})) = \int \frac{\exp\left(\gamma\ r_{t+1}(H_t)\right) - 1}{\gamma} P(s_{t+1}|\mathcal{I}_t)\ ds_{t+1}, \tag{5}$$

with equation (1) substituted for the hedged currency return. The optimization is again implemented evaluating the expected utility increase over a fine grid of values for the hedge ratio H_t, and choosing the optimal H_t.

The analysis is carried out using daily observations on the DMark/US Dollar exchange rate over the period 1/1/1982-31/12/1999 (4695 observations). We look at the case of a manager based in Germany, seeking to hedge currency risk connected to investments in a US Dollar denomination. Estimation[4] of the posterior distribution is done using the first 16 years of data, leaving two years (523 days) for evaluating the hedging decision. For the interest rates, the 1-month Eurocurrency middle rates are used.[5,6] Figure 1 depicts both the exchange rate S_t (over both the estimation and evaluation periods) and the interest rate differential $r_t^{\mathrm{DM}} - r_t^{\mathrm{USD}}$ over the evaluation sample. The local trending of the exchange rate seen in the first panel of figure 1 is modelled by the varying mean of the exchange rate returns. We note that this local trending behaviour is not apparent from autocorrelation of the exchange rate returns; day-to-day jumps are large relative to a possibly varying mean.

In table 1 the modes of the posterior distributions of the parameters are presented. For the **WN** model, a very small mean return μ is found. The mode of the standard deviation of the observation equation (2), σ_ϵ, is 0.68; for the other models a similar value is found. The distribution of the standard deviation σ_η of the disturbance in transition equation (3) for the model **LL** has a mode which is 30-fold smaller than that of σ_ϵ, indicating that the signal is weak. For the other

[3]Instead of resampling the posterior distribution of the parameters for each sample size t, we use the same posterior distribution $\theta|\mathcal{I}_T$ for evaluating all predictive densities $P(s_{T+j}|\mathcal{I}_T), j > 0$ in the evaluation period. As the estimation sample is large compared to the evaluation sample, the approximation error appears to be small.

[4]All calculations are performed in Ox version 2.20 (see Doornik 1999), using SsfPack 2.3 (Koopman, Shephard and Doornik 1999)

[5]We assume the investor hedges the currency risk using 1-month forward contracts. If the hedge position is changed, the overlay manager may need to reverse contracts.

[6]Source: Datastream, series `DMARKER/USDOLLR`, `ECWGM1M`, `ECUSD1M`.

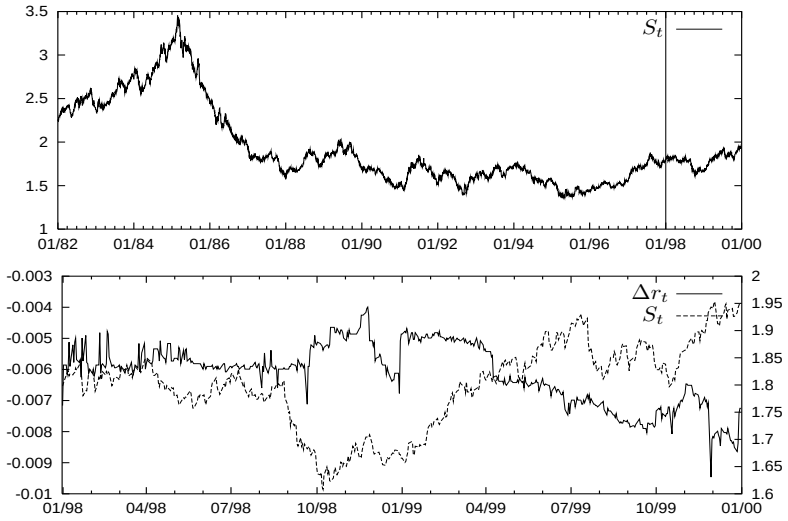

Figure 1: DMark/US Dollar exchange rate and the interest rate differential

Table 1: Summary statistics of posterior distributions

Parameter	**WN**	**LL**	**GLL**	**GLL-Student t**	**GLL-GARCH**	**GLL-GARCH-Student t**	**GLL-SV**
$\mu \times 100$	-0.39						
ρ	1	1	0.75	0.75	0.83	0.86	0.77
$\sigma_\eta \times 10$	0	0.23	0.59	0.53	0.60	0.49	0.54
σ_ϵ	0.68	0.67	0.67	0.67	0.65	0.76	0.67
δ					0.90	0.92	
$\alpha \times 10$					0.65	0.64	
ν				4.48		4.82	
μ_h							-1.07
ϕ							0.93
σ_ξ							0.28
S/N $\times 100$			2.36	1.66	2.92	1.46	2.06
Ln PO	-6.5	-54.6	0	143.8	163.2	256.0	270.9

The table reports the modes of the posterior density of the parameters, together with the signal-to-noise ratio and the logarithm of the posterior odds.

models, an AR coefficient ρ for the signal around 0.8 is found. Together with σ_η in the range 0.05-0.06, this results in an average signal-to-noise ratio

$$\text{S/N} = \frac{\sigma_\mu^2}{\sigma_\epsilon^2} = \frac{\sigma_\eta^2/(1-\rho^2)}{\sigma_\epsilon^2} \tag{6}$$

between 0.0146 and 0.0292 as reported in the bottom of the table. The low signal-to-noise ratio also results in a wide HPD region (not reported here) of the parameter ρ. The GARCH parameters δ and α and also the SV parameters ϕ and σ_ξ are estimated with high precision, indicating the importance of catering for varying variances in the model. The degrees of freedom parameter ν is estimated between 4 and 5, corresponding to tails in the disturbance density which are thicker than the tails of the normal density.

The last row of the table indicates the logarithm of the posterior odds (see Kass and Raftery 1995, Chib 1995), comparing to the **GLL** model. The models are given equal prior probabilities (therefore the posterior odds equal the Bayes factors). As the number of observations is large, only a small evidence per observation in favour of a certain model already causes huge (log) posterior odds. Even so, the results calculated here are stable between different samples from the posterior distribution. The importance of modelling the varying mean μ_t is clear from the results. Only the **LL** model has a lower marginal likelihood than the **WN** or **GLL** models. Including varying variances improves the marginal likelihood, with highest posterior odds/marginal likelihood for the **GLL-SV** model.

3 The variability of hedging decisions

Given the posterior distributions from the previously described models, we evaluate the optimal hedging decision for a risk-averse investor, with a risk tolerance parameter γ of -10. Table 2 reports statistics on the hedging decisions, table 3 sheds light on the resulting returns and utilities both in the first year and over both years in the evaluation period 1/1/1998-31/12/1999 jointly. Figure 2 presents the evolution of the optimal hedge ratio through time.

Table 2: Variability of hedging decisions

Model	$\overline{H}$	$H=0$	$H=1$	$\overline{\lvert\Delta H\rvert}$	3M-L	3M-G
Full hedge	1.00	0	523	0.000	−0.49	−0.33
No hedge	0.00	523	0	0.000	−12.86	9.57
RW	0.46	281	242	0.471	−7.49	8.16
WN	0.91	0	0	0.004	−1.33	0.32
LL	0.47	206	191	0.074	−5.51	8.55
GLL	0.78	0	119	0.105	−1.58	2.09
GLL-Student t	0.75	0	99	0.122	−1.77	2.39
GLL-GARCH	0.61	65	162	0.179	−3.35	4.38
GLL-GARCH-Student t	0.58	66	135	0.176	−3.60	5.65
GLL-SV	0.62	46	135	0.187	−2.16	3.95

Columns report the average hedge ratio, number of occurrences of a no-hedge or fully hedged position, the average absolute change in position and the maximum loss and gain over a period of three months over the period 1/1/1998-31/12/1999, for a risk averse ($\gamma = -10$) investor.

The first panel in tables 2 and 3 corresponds to special hedging cases where no use of a model was made. An infinitely risk-averse investor would choose never to run any exchange rate risk and will have $H = 1$ over the complete evaluation sample, 1/1/1998-31/12/1999. No changes in the hedging position occur, and the total return C over the two years is -3.20%, the cumulative

Table 3: Return and utility of hedging decisions

	First year		Both years	
Model	C	U	C	U
Full hedge	−1.48	−1.48	−3.20	−3.20
No hedge	−7.65	−11.52	8.18	0.24
RW	−3.61	−5.34	7.56	3.35
WN	−1.99	−2.01	−2.18	−2.24
LL	−8.02	−8.93	3.51	−0.16
GLL	−2.56	−2.73	−1.05	−1.59
GLL-Student t	−2.59	−2.78	−0.50	−1.17
GLL-GARCH	−3.35	−3.96	−0.56	−2.62
GLL-GARCH-Student t	−3.28	−3.95	2.31	−0.01
GLL-SV	−1.54	−1.97	4.29	3.10

Columns report cumulative returns and utilities over the first year and over both years jointly, for a risk averse ($\gamma = -10$) investor.

interest rate differential over the period. A risk-seeking investor, not hedging at all during the 523 days, obtains a return equal to the return on the exchange rate. Note that this includes periods with large losses and others with large gains. Reported in the columns labeled by 3M-L and 3M-G are the maximum losses and gains encountered over a 3-month period. The power utility function with risk tolerance parameter $\gamma = -10$ is skewed, penalizing losses more than equally sized gains can counterbalance, such that the utility of the no-hedge case is -11.52 for a return of -7.65 in the first year, and only 0.24 for a total return of +8.18 over the two years. The third row reports the results from a Random Walk (**RW**) strategy, where tomorrow's risk is fully hedged whenever today a loss is led, and vice versa. The hedging position fluctuates strongly, judging from the average absolute change $\overline{|\Delta H|}$. The first panel in figure 2 displays the hedging decision for this case. The maximum loss over a three month period is no more than 7.49%, substantially less than the 12.86% loss of the exchange rate itself. The total return is slightly less than the return on the exchange rate itself, even though on almost half of the days the risk was hedged. This results in a positive utility after two years of 3.35, as most of the exchange rate return is obtained at a lower risk. It turns out that improvements on this simple strategy during this evaluation period of a strongly appreciating currency are not easy to find.

The second panel in table 2 reports results on the hedging decisions based on the predictive density $P(s_{t+1}|\mathcal{I}_t)$ resulting from the models, after integrating out the parameters. The **WN** model allows neither the mean nor the variance to vary over time. Therefore, the only fluctuating element in the return equation (1) is the interest rate differential Δr_t. The hedge ratio is high on average, with a final return and utility, in table 3, close to (but slightly higher than) the fully hedged return of -3.20%.

The **LL** model takes an extreme position: The exchange rate is modelled as an I(2) process. When the hedging decision changes, it changes strongly. In 397 out of 523 days a border solution $H = 0$ or $H = 1$ is chosen. The final hedging result is halfway the no-hedge and full-hedge return. Even though the return is lower than the no-hedge return, utilities are of similar size, as the return is gained incurring less large losses. The maximum 3-month loss and gain are similar to the results of the **RW** strategy.

The **GLL** model is more conservative, with a slower adaptation to shocks. It results in a higher average hedge ratio (see also panel 4 in figure 2). This model points out that the risk resulting from the varying exchange rate is most of the time too high, resulting in a negative return (though it is still higher both in return and utility than the **WN** or fully hedged return). The downside risk, judging from the value of 3M-L, is covered to a large extent.

The **GLL-Student t** model includes a heavier tailed disturbance distribution. The extreme position $H = 1$ occurs less often, leading to a higher final return than for the **GLL** model. As the

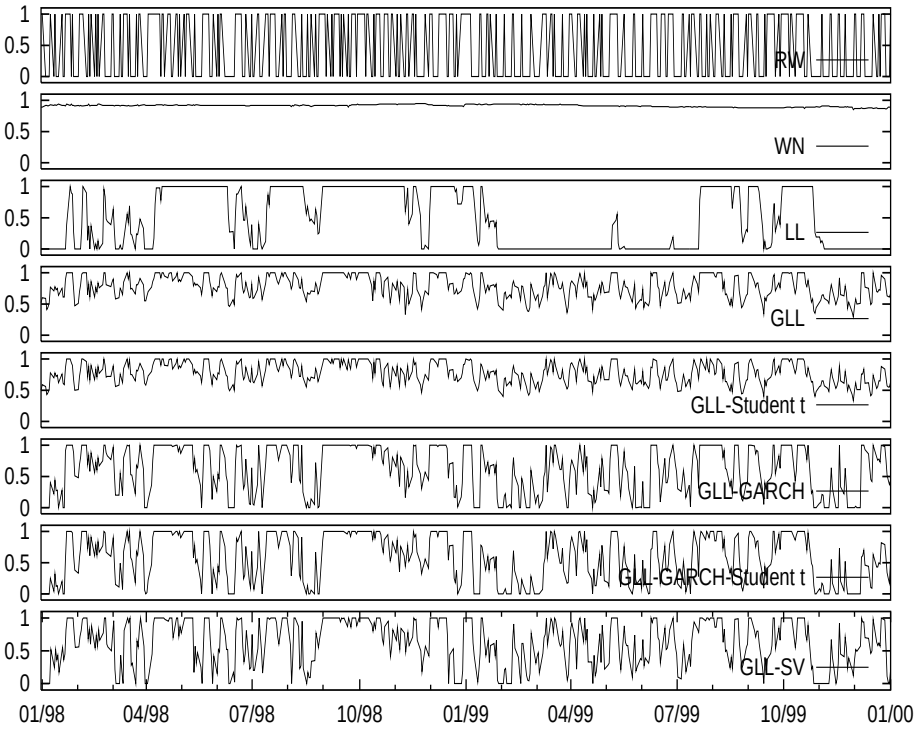

Figure 2: Hedging decisions during the evaluation sample

influence of the heavy tails on the hedging decision is not large, returns and utilities do not differ much with the **GLL** model.

The next model introduces changing variances through a GARCH component. Periods of lower risk can now be recognized, with a lower average hedge ratio as a result. Both downside risk and upside risk appear to be larger compared with the **GLL** and **GLL-Student t** models; the final return is slightly higher. The flexibility of the heavier tails and the GARCH effect are combined in the **GLL-GARCH-Student t** model. The downside risk hardly changes, but larger gains are made in periods of appreciation. The **GLL-SV** model is indicated by the posterior odds as the model with the best fit. It leads to less extreme hedging positions (both $H = 0$ and $H = 1$ are chosen less often), but the hedging position changes considerably ($\overline{|\Delta H|}$ is largest). The downside risk is lower compared with the other models with varying variance. In the first year, with strong depreciation in the foreign currency, the **GLL-SV** model manages to incur no stronger loss than the cumulative interest rate differential, leading to the best utility of the model-based strategies. In the second year, highest positive gains are made, leading to a utility which is of a size similar to the cumulative result of the **RW** case.

Table 4 displays the correlation between the hedging decisions. The similarity between the results for the **GLL** and **GLL-Student t** models results in a correlation coefficient of 0.96. Likewise, inclusion of the Student t disturbance distribution only marginally alters decisions from the **GLL-GARCH** model: Even though the final cumulative return is 2.9% higher, correlation between the H_t's is 0.96 as well. The difference in returns results from a small number of days with large appreciations on which the **GLL-GARCH-Student t** model leads to a lower hedge ratio. It is interesting to note that the **RW** strategy, seemingly so random, has a positive correlation of 0.5 with most model based strategies.

The results for the **RW** hedging strategy indicate a higher utility than for the model based strategies, at least for the two-year evaluation period with the strong appreciation of the US

Table 4: Correlation between hedging decisions

Model	**RW**	**WN**	**LL**	**GLL**	**GLL-Student t**	**GLL-GARCH**	**GLL-GARCH-Student t**
WN	−0.02						
LL	0.20	0.11					
GLL	0.55	0.07	0.57				
GLL-Student t	0.60	0.13	0.57	0.96			
GLL-GARCH	0.48	−0.01	0.58	0.91	0.88		
GLL-GARCH-Student t	0.49	0.04	0.63	0.88	0.91	0.96	
GLL-SV	0.54	0.02	0.50	0.88	0.89	0.92	0.91

Dollar vis-a-vis the German DMark, and with a risk tolerance of $\gamma = -10$. Figure 3 displays the cumulative utilities over the evaluation period for the set of models and the **RW**-based strategy, over a range of risk tolerance parameters γ. For $\gamma \in [-14, -6]$, the utility of the **RW** strategy is (slightly) higher than for the best model strategy. For lower values of γ, the losses of the **RW** strategy are penalized such that the final utility is lower; for values of $\gamma > -6$ the average hedge ratio of the model based strategies becomes smaller, such that the appreciation of the exchange rate is picked up in the returns and utility. In unreported calculations we switched focus to an investor based in the US, seeking to hedge the currency risk connected to investments in the German market. In this case the **RW** strategy loses out compared to the other strategies, for almost all values of γ. The **RW** strategy is not able to effectively cover downside risk.

In figure 3 it is also seen how the ordering of the models indicated by the posterior odds in table 1 does not correspond to a clear ordering of models by their cumulative utility over the two years considered. Models with varying variance tend to deliver better hedging results, though differences are small. Taking an incorrect decision on just one or two days can lead to a lower final utility, which can switch the ordering in utility between models. Switching the focus to a US investor in the German market alters the ordering, then the **GLL-GARCH** model comes out first.

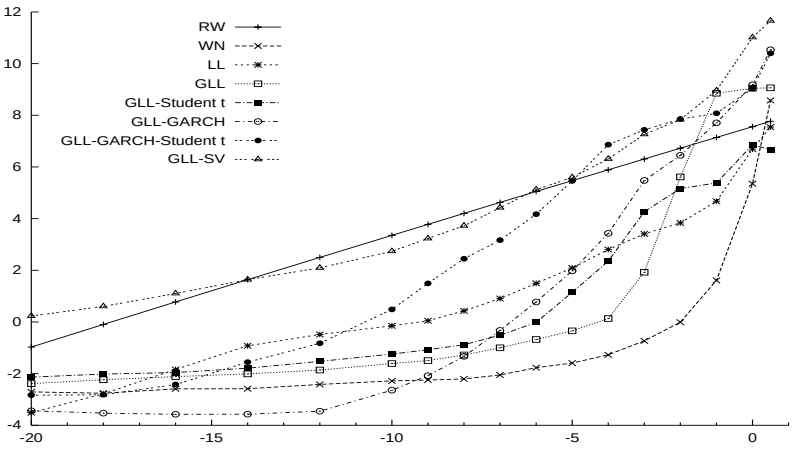

Figure 3: Cumulative utility attained using different strategies over the period 1/1/1998-31/12/1999, for a range of risk tolerance parameters

4 Summary of the results

For 7 models, differing in the method of modelling varying means, varying variances, and the distribution of the disturbances of the exchange rate return, we evaluated the optimal strategy of hedging the currency risk for a risk-averse investor. The models were compared on the basis of the optimal hedging decisions, and of the financial return, downside risk and the upside risk.

It was found that even though the exchange rate returns exhibit very little correlation, it does make a large difference in the hedging decision whether local trending behaviour in the exchange rate is modelled. An even larger effect was found when the volatility was allowed to change over time. In periods of lower variance, an investor could adapt his hedging position accordingly, while still hedging against depreciations in periods of higher volatility. The difference in hedging decisions between the **GLL-GARCH**, **-GARCH-Student t** and **-SV** models is small. Even though the correlation in decisions is high, a small difference in the hedging ratio on only a few days with stronger movements in the exchange rate can alter the (ordering of the) final cumulative returns and utilities. Furthermore, it was found that in periods of strong appreciation simple strategies that hedge little can appear to be best, and likewise during strong depreciation a strategy always hedging a high percentage can perform well. The added value of model-based hedging decisions is better appreciated in periods with switching high and low downside risk, which e.g. the **GLL-SV** model covers well. The framework described here can be of use for an overlay manager, as a decision-supporting tool.

An evaluation of the performance of the hedging strategies over a longer period, and on exchange rates between other currencies, can shed more light on the robustness of the results presented in this paper, and is left for later research.

References

Bauwens, L., Bos, C. S. and van Dijk, H. K. (1999), Adaptive polar sampling with an application to a Bayes measure of Value-at-Risk, Tinbergen Discussion Paper TI 99-082/4, Tinbergen Institute.

Bos, C. S., Mahieu, R. J. and van Dijk, H. K. (2000), 'Daily exchange rate behaviour and hedging of currency risk', *Journal of Applied Econometrics* **15**(6), 671–696.

Carter, C. K. and Kohn, R. (1994), 'On Gibbs sampling for state space models', *Biometrika* **81**(3), 541–553.

Chib, S. (1995), 'Marginal likelihood from the Gibbs output', *Journal of the American Statistical Association* **90**(432), 1313–1321.

Chib, S. and Greenberg, E. (1995), 'Understanding the Metropolis-Hastings algorithm', *The American Statistician* **49**(4), 327–335.

Doornik, J. A. (1999), *Object-Oriented Matrix Programming using Ox*, 3rd edn, London: Timberlake Consultants Ltd. See `http://www.nuff.ox.ac.uk/Users/Doornik`.

Geweke, J. (1989), 'Exact predictive densities for linear models with ARCH disturbances', *Journal of Econometrics* **40**, 63–86.

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge.

Jorion, P. (1985), 'International portfolio diversification with estimation risk', *Journal of Business* **58**(3), 259–278.

Kass, R. E. and Raftery, A. E. (1995), 'Bayes factors', *Journal of the American Statistical Association* **90**(430), 773–795.

Kim, S., Shephard, N. and Chib, S. (1998), 'Stochastic volatility: Likelihood inference and comparison with ARCH models', *Review of Economic Studies* **64**, 361–393.

Koop, G. and van Dijk, H. K. (2000), 'Testing for integration using evolving trend and seasonal models: A Bayesian approach', *Journal of Econometrics* **97**(2), 261–291.

Koopman, S. J., Shephard, N. and Doornik, J. A. (1999), 'Statistical algorithms for models in state space using SsfPack 2.2', *Econometrics Journal* **2**, 107–160.