



ELSEVIER

Journal of Banking & Finance 24 (2000) 1131–1154

Journal of
BANKING &
FINANCE

www.elsevier.com/locate/econbase

Value at risk models for Dutch bond portfolios

Peter J.G. Vlaar *

*Econometric Research and Special Studies Department, De Nederlandsche Bank NV, P.O. Box 98,
1000 AB Amsterdam, The Netherlands*

Received 23 July 1998; accepted 27 May 1999

Abstract

This study investigates the consequences of dynamics in the term structure of Dutch interest rates for the accurateness of value-at-risk models. Therefore, value-at-risk measures are calculated using both historical simulation, variance–covariance and Monte Carlo simulation methods. For a ten days holding period, the best results were obtained for a combined variance–covariance Monte Carlo method using a term structure model with a normal distribution and GARCH specification. Term structure models with a *t*-distribution or with cointegration performed much worse. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: E43; G28

Keywords: Value-at-risk; Bond portfolio; Historical simulation; Variance–covariance method; Monte Carlo simulation

1. Introduction

Risk management by banks is currently receiving a great deal of attention. Recent bank failures, those of BCCI and Barings Bank, for instance and the difficult situation in which the Japanese banking system and a number of French banks find themselves, have undoubtedly contributed to this attention.

* Tel.: +31-20-524-3189; fax: +31-20-524-2529.

E-mail address: p.j.g.vlaar@dnb.nl (P.J.G. Vlaar).

In addition, the increased international cross-linking between financial conglomerates has impressed on supervisory authorities the possible serious consequences a bank failure may have for international payments and thus for the economic system as a whole. This situation is generally referred to as the systemic risk.

Where traditionally bank failures were primarily the consequence of an excess of non-performing loans (the so-termed credit risk), the failure of Baring's Bank has shown that changes in financial asset prices may also lead to bankruptcy. This price risk, also known as market risk, will probably increasingly determine banks' profitability as well as their risks of failure. The reasons for this include the growing importance of portfolio investment by banks and the increasing use of derivatives such as options and futures. After all, these products are characterised by greater leverage, large-value commitments being entered into with only a limited amount of invested capital.

In response to the increased market risks, the banking supervisory authorities have drawn up rules, in international consultation both within the Basle Committee on Banking Supervision (1994, 1996) and the European Union, laying down capital adequacy requirements to cover these risks. Initially, in the determination of these requirements relatively little allowance was made for hedging. More recent proposals, however, view the risks of investment portfolios in a more integrated manner, with a balanced portfolio of assets and liabilities leading to less strict requirements. An important new concept in this context is the so-termed value-at-risk (VaR). The VaR is an expression of the market risk of a financial institution's total portfolio, representing the maximum amount which may be lost within a particular period (the so-termed holding period) on the portfolio, if it is not modified in the meantime, in all but exceptional (say 1% of) cases.

Using over 17 years of daily data, this article will compare the out-of-sample performance of VaR standards calculated on the basis of both the three permissible methods and a combination thereof, for 25 hypothetical portfolios consisting of Dutch government bonds of eight different maturities. An analysis will be made of which assumptions (regarding the expectation, the variance as well as the distribution of interest rate changes) are important in determining an adequate VaR standard for these bond portfolios. As the VaR is supposed to give the amount of money that can be lost with a (say) 1% probability, the VaR is considered adequate if the frequency of actual losses in excess of the calculated VaR is approximately 1%.

2. Value at risk

The Basle directives for the determination of capital adequacy requirements for market risk (BIS, 1996), which also form the basis for European and Dutch

legislation, allow banks and investment institutions to determine their capital requirements for market risk on the basis of VaR models. In doing so, they must make allowance for a holding period of 10 days and a 99% confidence level. In order to determine the capital adequacy requirement, this VaR must be multiplied by an add-on factor which may vary from three to four, depending on the actual percentage of exceedances from the past. Partly for this purpose, VaR models should be evaluated at least once a year.

Despite the fact that the holding period is two weeks, the models must be based on daily data. In this context, three different methods to determine the VaR are permitted, namely historical simulation, Monte Carlo simulation or variance-covariance techniques. In the case of *historical simulation*, a calculation is made for a past period of the hypothetical returns of the current portfolio over an assumed two-week holding period. The return exceeded in 99% of cases is taken as the VaR standard. There is no consensus on the preferred length of the simulation period. In the case of a short period, the result will be very sensitive to accidental outcomes from the recent past. A long simulation period, on the other hand, has the disadvantage that data may be included which are no longer relevant to the current situation. This method is characterised by the fact that only a few observations determine the entire result. A problem of this method for multi-day VaRs is that, due to overlapping data, the various yields are no longer independent.

In order to apply the Monte Carlo method and the variance-covariance method, it is necessary to adopt an assumption about the multivariate statistical distribution of the asset returns, including assumptions regarding the average yield and the degree of uncertainty. In a Monte Carlo simulation, a large number of random samplings is taken from this distribution after which the change in value is calculated for each of the samplings. In that case, the VaR, similarly to that in the historical simulation, is defined as the result that is exceeded in 99% of cases. In the variance-covariance method, the VaR is calculated exactly on the basis of the distributional assumptions. The market risks of products which are nonlinearly dependent on the price level (such as options), can only be included as a linear approximation in this method. In applying this method, for the non-linear risks an additional capital requirement must be calculated. The Monte Carlo method does, in principle, permit nonlinearities to be taken into account.

3. The models considered

3.1. Composition of portfolios

This article will compare VaR standards for 25 portfolios comprising Dutch fixed-interest securities. The choice of fixed-interest securities rather than

shares, for instance, is based on the dominant influence of interest-driven business items on banks' balance sheets. With a view to the calculation of risks, we distinguish the following eight different maturities: money market rates (Eurocurrency) with maturities of one, three, six and twelve months, respectively, and capital market rates (public authority bonds) with remaining maturities of one to three, three to five, five to seven and seven to ten years, respectively. As regards these eight maturities, this article has used the daily interest rate levels over the period January 1980 through March 1997.

Ideally, the weights of the different interest rate items in the 25 portfolios should correspond as closely as possible to real positions. Dimson and March (1995), for instance, use data from trading portfolios to determine the portfolio composition in evaluating capital requirements. As no such data were available for the present article, the weights were determined as the outcome of a probability experiment. For each portfolio, a random draw was taken from an eight-dimensional uniform distribution from -1 to 1 . For each maturity, the probability of a short position is, therefore, equal to that of a long position, and the weights for the different maturities are independent. The weights should be interpreted as the portfolio value's sensitivity with respect to changes in the various interest rates. In order to avoid dominance of the portfolios with the largest weights, for each portfolio the weights are subsequently scaled such that the absolute values of the weights sum to 100. Consequently, if all interest rates move exactly one percentage point in an unfavourable direction, the amount of money lost will be 100. As long-term bonds are more interest-sensitive than short-term bonds, the adopted approach implies unbalanced portfolios with a relatively large number of short-term certificates. In addition to bonds, portfolios may also comprise other products whose values linearly depend on interest rate changes. We refer to Estrella et al. (1994) for the evaluation of VaR standards regarding non-linear products such as options. The same 25 portfolios are used in all calculations. A comparison with results for 100 simulated portfolios showed that 25 portfolios was a sufficient number for a comparison of methods. The differences over time are more important than the differences in portfolio composition.

3.2. Historical simulation

The first method to determine the VaR is historical simulation. This method takes the empirical distribution of the investment results in the past as being constant and representative for the coming ten days. Table 1 lists the results for this method where the empirical distribution is determined from respectively the 250, 550, 750, 1250 and 2550 hypothetical ten-day investment results, directly preceding the day for which the current VaR is calculated.

Table 1 shows that the average VaR (represented as $\overline{\text{VaR}}$), averaged both over time and portfolios, increases as the period over which it is determined

Table 1
Evaluation of VaR measures based on historical simulation^a

History	$\overline{\text{VaR}}$	Surplus	% exceedances	Needed	# years with exceed >1%
250	18.76	-4.41	2.181	23.17	10
550	21.49	-1.39	1.298	22.88	7
750	22.66	-0.45	1.096	23.11	5
1250	24.66	1.71	0.708	22.95	3
2550	25.49	5.27	0.338	20.22	1

^a *Explanation:* The VaR measures have been calculated as the average of the two ten-day investment results which are just respectively just not exceeded in 99% of all cases. The following items are listed in succession: the history taken into account when calculating the VaR, the average size of the VaR, the amount that may be deducted from each VaR to arrive at exactly 1% of exceedances, the realized percentage of exceedances, the amount needed to arrive at a maximum of 1% exceedances (VaR minus surplus) and the number of years (of 13) in which the exceedance percentage is higher than 1%.




Fig. 1. Standard deviation interest rate changes.

becomes longer. A similar pattern has been demonstrated by Hendricks (1996) for currency portfolios over the 1978–1995 period. The explanation can probably be found in the gradual decrease in the volatility of both (European) exchange rates and interest rates. Fig. 1 shows the volatility, calculated over a moving period of five years, for the eight interest rates considered in the present article. The four money market rates all experienced a decrease in volatility over the entire period, while the volatility of the capital market rates also decreased almost consistently up to 1994, when it started to rise again slightly.

Given the fact that the volatility of the individual interest rates has decreased so much, it is surprising that the VaR standards that were calculated from empirical distributions over at most 750 days, are on average insufficient. Should the volatility increase again in the future, these standards would almost certainly underestimate the real risks. The cause for this underestimation of risks can probably be found in the dominant influence which only a few accidental results have on the size of the calculated VaR, particularly as a result of the use of overlapping data. Even if the underlying stochastic process which determines interest rate movements does not fluctuate over time, the VaR calculated on the basis of simulations will fluctuate due to accidental results. These accidental fluctuations will lead to too high a percentage of exceedances of the VaR as the probability of additional exceedances at too low a calculated value is greater than the probability of an exceedance less in the event of too high a calculated value.

As regards the differences between portfolios, no systematic patterns are discernible. The only factor that seems to have some influence on the exceedance rate is the relative weight of capital market interest rates in the portfolio. The VaR of portfolios with relatively large capital market positions, either long or short, are exceeded somewhat more often. For none of the models this influence was statistically significant however. Differences in long or short positions, either in money market or capital market interest rates had no systematic impact on the exceedance rates.

3.3. Monte Carlo simulation

The second method considered is the Monte Carlo simulation technique. This technique requires an assumption regarding the multivariate statistical distribution of price changes of the assets considered – in this case interest rates. This model assumption can be conceived of as comprising three sub-elements, namely the expected change in value, the degree of uncertainty in the form of the variance–covariance matrix, and the type of distribution. In this study, the parameter values of the return distribution used in the VaR calculations for a particular calendar year are estimated with regard to data of the five calendar years preceding that year, thereby assuming a yearly update of the models.

3.3.1. The expectation

As regards many assets, there is no real need for a model for the expected price change as the latter can be set at zero. The danger of modelling accidental ‘noise’ in the data seems greater than the potential benefits to be gained from a more accurate prediction of the average (Figlewski, 1994). It is, however, open to question whether this also applies to interest rate changes. Interest rates for

different maturities will, after all, show some interrelation. According to the theoretical expectations hypothesis of the term structure, interest rates for long maturities should be equal to a geometric mean of expected future short-term interest rates. In accordance with these models, interest rates for long maturities should thus have predictive value for future short-term interest rates, which in turn should result in a long-term equilibrium relationship between interest rates for different maturities. Within a system of interest rates for eight different maturities, this would in principle have to lead to seven equilibrium relationships.¹ This pattern may be disturbed, however, by possible time-varying term premiums in interest rates for longer maturities. The expectations hypothesis of the term structure has been frequently studied, see for instance Campbell and Shiller (1984, 1987), Taylor (1992), Anderson et al. (1992) and Shea (1992). Although generally the existence of equilibrium relationships is proven, the specific assumptions underlying the expectations hypothesis must be rejected. Furthermore, the equilibrium relationships prove to be sensitive to the monetary policy pursued. Bradley and Lumpkin (1992) model only one single equilibrium relationship between seven interest rates and subsequently show that more accurate forecasts can be made if this relationship is included in the forecasts. Almost all of these studies, however, are based on monthly data. It is not at all certain that the information contained in the term structure also has predictive value for daily data. However, given the fact that the models will be used to forecast two weeks ahead, it seems appropriate to make allowance for the interrelationship between interest rate levels. The present article will consider both models with and without equilibrium relationships. In this context, the number of relationships and the coefficients within these relationships will be determined on statistical grounds on the basis of the so-called cointegration technique (Johansen, 1988).

In addition to interrelationships between the levels, information about lagged changes in interest rates may also have predictive value for future interest rates. Unexpected changes in official interest rates, for instance, may have a more immediate impact on the money market than on the capital market. Furthermore, an analysis of interest rate differences per maturity (see Appendix A) indicates that capital market rates in particular show significant positive autocorrelation.² The model will make allowance for autocorrelation by including lagged interest rate changes, and will analyse both the results for

¹ Seen over very long periods, it seems logical that the short-term interest rate itself, too, tends towards long-term equilibrium (in statistical terms, the interest rate should be stationary). Within a five-year period, such a pattern cannot be identified, however, as is shown in the ADF test in Appendix A.

² This seems to be contradictory to the efficient market hypothesis. It is however questionable whether the predictable gains exceed transaction costs. In addition, the phenomenon may have to do with infrequent trade in some of the bonds from which the interest rates have been derived.

one and for two lags.³ In addition, the random walk model (an expected change of zero) will be reviewed.

3.3.2. The variance–covariance matrix

As regards the variance–covariance matrix of interest rate changes, too, several specifications are possible. The fact that the volatility of financial assets fluctuates over time has been known for a long time (see, for instance, Fama, 1965). Yet, the modelling of this time-varying uncertainty only started properly with the so-termed Auto-Regressive Conditional Heteroskedasticity model (ARCH) (Engle, 1982). In the ARCH model, the variance is modelled as a linear function of lagged squared prediction errors. The success of this specification has led to very many applications and extensions (see Bollerslev et al., 1992, for an overview). The most popular adaptation is the so-termed generalized ARCH model (GARCH) of Bollerslev (1986). In this specification, the variance at time t depends on both lagged squared disturbance terms and lagged conditional variances:

$$h_t = \alpha_0 + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j h_{t-j},$$

where h_t and ε_t are the conditional variance and the prediction error at time t , respectively, and α_i and β_j are parameters to be estimated. In practice, it almost always proves possible to limit the number of lagged squared disturbance terms and conditional variances to one (GARCH(1,1)). Given the clear GARCH effects in the data (see Appendix A), this specification will also be tested in the present article, on the assumption that the correlations between the disturbance terms of the various model comparisons are constant (Bollerslev, 1990). This assumption, however, does not prove to be entirely tenable for the differences in interest rates in the present article (see Appendix A).

It must be said that a potential drawback of the GARCH specification is that the parameters must be estimated, which may cause problems especially if the number of variables becomes large. This is why simplifications of the GARCH model have been introduced, of which the RiskMetrics method by JP Morgan (1996) is by far the most popular. In the RiskMetrics method, the GARCH parameters α_0 , α_1 and β_1 are not estimated, but set at 0, 0.06 and 0.94, respectively. In this specification, the current conditional variance, therefore, equals a weighted average of the one-day-lagged squared disturbance term (with a 0.06 weight) and the one-day-lagged conditional variance (with a

³ The problem is that different information criteria to determine the optimal number of lags, lead to different conclusions. According to the Schwarz criterion, only one lag would suffice, according to Hannan–Quinn two, while Akaike deems nine lags necessary.

weight of 0.94). Substituting the lagged conditional variance results in the following expression:

$$h_t = \sum_{i=1}^{\infty} 0.94^{i-1} 0.06 \varepsilon_{t-i}^2.$$

The conditional variance, therefore, proves to be an exponentially weighted average of lagged squared disturbance terms. The present article will test both the aforementioned specification and a parameterization of α_1 and β_1 of 0.03 and 0.97, respectively.

The GARCH-type methods will be compared with a naive variance method in which the conditional variance equals an unweighted average of squared prediction errors over the five years immediately preceding the ten-day period to be forecast.

3.3.3. The distribution

As regards the assumed distribution, the present article will consider two alternatives. The first is the normal distribution. The most important reason underlying this choice is that VaR calculations almost always assume a normal distribution, mainly because of its practical advantages. After all, if the normal distribution is assumed it is possible to determine an analytic solution for VaR standards with a holding period of more than one day.⁴ Another practical advantage of the normal distribution is that the model parameters are usually easier to determine. Particularly for models with cointegration, it is a great advantage that an analytic solution for the model parameters for the expected change is available (Johansen, 1988). In order to benefit from this feature, the interest rate models with GARCH and normal distribution will be estimated in a two-step procedure, in which the first step will be the estimation of the expected changes in interest rates and the correlation matrix of the prediction errors,⁵ and the second the estimation of the variances. This procedure provides consistent estimations, although it would be more efficient to estimate all parameters at the same time.

The second distribution that will be reviewed is the Student- t distribution with five degrees of freedom.⁶ Despite the practical advantages of normal distribution, its assumptions are rarely fulfilled in daily financial data. Also for our data, the real distribution proves to be more ‘thick-tailed’ than the normal distribution

⁴ The normal distribution is, after all, the only stable (that is to say, capable of being aggregated) distribution with a finite variance.

⁵ If the correlation matrix is estimated in the second step, together with the GARCH parameters, the average VaR turns out to be slightly lower and the number of exceedances slightly higher.

⁶ Free estimation of the t -distribution’s number of degrees of freedom has also been considered, but in several cases the estimated degrees of freedom was below two, which would mean that the variance did not exist. This result is probably due to a limited number of outliers (Lucas, 1997b).

in all cases (see Appendix A). The t -distribution is more ‘thick-tailed’ than the normal distribution and thus seems to correspond more closely to reality.

As there are no analytic solutions for any parameters in the Student- t model, the entire model must be optimized by means of an iterative procedure in which the maximum likelihood estimator is sought. Especially for the models with cointegration this proved quite a difficult problem. In contrast to the models with normal distribution, the models with GARCH specification and t -distribution were optimized as a whole (i.e. with the expectation and variance covariance matrix taken together). In this context, the determination of the greatest possible number of equilibrium relationships was based on the same asymptotic theory as that used for the models with a constant variance (Lucas, 1997a), although formal proof that this is indeed correct in the case of GARCH has never been provided. More research would be needed.

3.3.4. The simulation procedure

The Monte Carlo VaR standards are based on 10 050 calculations with regard to random samplings from the assumed distribution. As the portfolios comprise eight different maturities and given the ten-day prediction horizon, this means 804 000 random samplings per model per day. The variances of two to ten days ahead per individual maturity have been determined for the GARCH and weighted average methods on the basis of the sampled disturbance terms. For the naive method, a constant variance was assumed for the prediction horizon. The naive variance method in the Student- t distribution is based on the estimation over the preceding five calendar years as there is no analytic expression for the maximum likelihood estimator in the case of the t -distribution. A daily update of the variance would take too much time, therefore. The ten-days-ahead predictions for the eight maturities are calculated, for each of the 25 portfolios, the 10 050 investment results are subsequently ranked, after which the VaR standard is calculated as the average of the 100th and 101st results.

3.3.5. Results

Table 2 shows the results for the VaR standards on the basis of this method. At least six aspects draw our attention. First, on average none of the VaR standards is large enough, in the sense that the realized number of exceedances of VaR is always higher than the predicted 1%. Consequently, portfolio decisions based on the assumption that these VaR measures are unbiased will be too risky. The high exceedance percentages are all the more striking given the fact that the VaR based on historical simulation over approximately five years only saw an exceedance percentage of 0.7% (see Table 1). In addition, the volatility over the sample period decreased (see Fig. 1), which should, in fact, lead to VaR calculations that are systematically too high. To what extent this result is due to too limited a number of simulations will have to be studied in

Table 2
Evaluation of VaR standards based on 10,050 Monte Carlo simulations^a

Expectation	Distribution	$\bar{\varepsilon}^2$	\bar{h}	$\overline{\text{VaR}}$	Surplus	%	Needed	Years >1%
Random walk	A	68.9	86.7	20.19	-0.96	1.314	21.15	6
	B	68.9	59.3	16.37	-3.91	2.798	20.28	13
	C	68.9	57.8	16.42	-3.49	2.457	19.91	13
	D	68.9	80.0	19.72	-0.58	1.326	20.30	6
	E	68.9	51.3	14.90	-7.00	4.878	21.90	11
	F	68.9	59.2	15.96	-5.89	4.378	21.85	11
$p = 1$ No cointegration	A	67.7	97.5	20.97	-1.26	1.245	22.23	6
	B	67.7	70.6	17.56	-3.73	2.431	21.29	13
	C	67.7	69.2	17.64	-3.26	2.195	20.90	13
	D	67.7	85.4	20.03	-0.89	1.211	20.92	6
	E	67.5	55.8	15.47	-6.68	4.311	22.15	11
	F	67.6	60.5	16.23	-5.32	3.749	21.55	11
$p = 2$ No cointegration	A	67.2	100.0	21.18	-1.40	1.260	22.58	6
	B	67.2	73.6	17.88	-3.63	2.456	21.51	13
	C	67.2	72.2	17.96	-3.34	2.161	21.30	13
	D	67.2	88.3	20.30	-1.05	1.215	21.35	6
	E	67.5	54.4	15.30	-6.80	4.452	22.10	11
	F	67.6	57.4	15.97	-5.54	3.832	21.51	11
$p = 1$ With cointegration	A	90.5	90.8	20.12	-3.22	1.967	23.34	9
	B	90.5	70.3	17.43	-5.70	3.919	23.13	13
	C	90.5	68.6	17.44	-5.54	3.527	22.98	13
	D	90.5	81.4	19.41	-3.06	1.966	22.47	10
	E	77.5	52.9	14.93	-7.54	5.301	22.47	12
	F	79.7	55.9	15.57	-6.64	4.537	22.21	12
$p = 2$ With cointegration	A	95.9	94.6	20.73	-2.68	1.917	23.41	10
	B	95.9	73.8	18.06	-5.39	3.680	23.45	13
	C	95.9	72.1	18.07	-4.96	3.328	23.03	13
	D	95.9	85.5	20.08	-2.61	1.927	22.69	10
	E	77.9	51.7	14.81	-7.60	5.410	22.41	12
	F	81.0	55.7	15.43	-6.54	4.522	21.97	12

^a *Explanation:* The following items are listed: the model for the expected change in interest rates, in which p is the number of lagged interest rate changes included, an indication of the stochastic process, the average squared prediction error, the average value for the conditional variance, the average VaR, the surplus if exactly 1% of the results would have to exceed the VaR, the percentage of exceedances, the minimum amount needed to arrive at exactly 1% of exceedances, and the number of years (of a total of 13) in which the exceedance percentage is higher than 1%. The following stochastic processes have been distinguished: (A) normal distribution with naive variances; (B) normal distribution with exponentially weighted variances, with parameters 0.94 and 0.06; (C) like B, with parameters 0.97 and 0.03; (D) normal distribution with GARCH; (E) t -distribution with naive variances; (F) t -distribution with GARCH.

greater detail. In practice, however, simulation experiments are rarely conducted with more than 10 000 samplings.

Second, the large differences in average size of VaR standards calculated by different methods (also in comparison also with calculations on the basis of the historical simulations, see Table 1) indicate that the small fluctuations in the add-on factor (between 3 and 4) to determine the capital adequacy requirement, provide little incentive to banks to achieve a correct determination of the VaR. After all, a calculated VaR of 15, for instance, multiplied by an add-on factor of 4 still results in a lower capital requirement than a VaR of 25 multiplied by an add-on factor of 3. Greater fluctuations in the multiplier seem desirable as the most accurate VaR calculation might otherwise not lead to the on average lowest capital requirement.

A third striking aspect is that the models based on the t -distribution in all cases lead to a much higher percentage of exceedances than the models based on the normal distribution.⁷ This is all the more striking as the 1% critical value of the t -distribution with five degrees of freedom (2.6065 times the standard deviation) is 12% higher than that for the normal distribution (2.3263). The cause for the higher percentage of exceedances is that the maximum likelihood estimator for the variance–covariance matrix under the t -distribution assigns less weight to extreme observations. After all, because the chance of ending up in the tail of the distribution is greater with the t -distribution than with the normal distribution, the parameter values under the t -distribution will be forced less in a direction to prevent extreme observations. If the underlying stochastic process is indeed t -distributed, this provides the best result as accidental outliers have a smaller impact on the total result. If, however, the extreme observations are not representative of the distribution at other moments, the normal distribution is more satisfactory because it gives great weight to the extreme observations which are so crucial in the calculation of VaR. Given the many exceedances for the t -distribution, we must conclude that this distribution does not adequately describe reality.⁸

⁷ The differences between the normal and the distribution for the naive variance models cannot be explained by the fact that for the t -distribution the variances are not modified daily. If for the normal distribution constant variances per calendar year are assumed, slightly higher VaR standards and consequently lower exceedance percentages are the result.

⁸ The fact that the normal distribution provides better results than the t -distribution does not detract from the fact that the normal distribution does not give a good description of reality either. A mix of two normal distributions might be a better candidate. Such a mix may be seen as comprising a distribution for normal circumstances plus the probability of an external shock (due to, for instance, an unexpected change in the discount rate or unexpected macro-economic news, see Vlaar and Palm (1993)). Another solution may be provided by the use of so-termed tail estimators, which are not based on the total statistical distribution, but only on the extreme observations. Danielson (1997) has shown that VaR standards based on these estimators with a holding period of one day work well for share portfolios.

Table 3
Optimal number of cointegrating relationships in accordance with 95% critical value for models with restricted constant^a

Lags	$p = 1$			$p = 2$			
	Distribution	Normal	t	$t + \text{GARCH}$	Normal	t	$t + \text{GARCH}$
80–84		5	5	5	4	5	5
81–85		5	5	5	4	5	6
82–86		4	4	5	4	4	4
83–87		3	4	5	2	4	5
84–88		2	3	4	2	3	4
85–89		3	4	5	2	4	5
86–90		4	4	4	3	4	4
87–91		4	3	3	3	4	3
88–92		4	4	3	4	4	3
89–93		5	5	5	5	5	4
90–94		6	5	6	6	5	5
91–95		6	5	5	5	5	4
92–96		6	4	6	5	5	6

^a *Explanation:* The results for the normal distribution are based on Johansen (1988), those for the t -distribution on Lucas (1997a).

A fourth striking result is that the models with cointegration generate considerably worse results than those without. The average squared prediction error is considerably greater if cointegration is allowed for (especially on the assumption of normality), while the average conditional variances are smaller. As a consequence, the percentage of exceedances for models with cointegration is consistently higher than that for models without. Also, the minimum amount needed on the basis of the models with cointegration to arrive at 1% of exceedances, is always higher than that in the comparable models without cointegration. Apparently the long-term ‘equilibrium relationships’ are insufficiently stable over time to be useful in forecasting. An indication can also be found in Table 3, which specifies the estimated number of equilibrium relationships per sub-period for each of the distribution assumptions. This number turns out to fluctuate considerably over time.⁹ The differences between the distributions are relatively small. In any case, no systematic patterns can be found.

A fifth point of special interest is the VaR’s high percentages of exceedances for the models with variances on the basis of exponentially weighted squared prediction errors (RiskMetrics). In all cases, the percentage of exceedances is higher than that in the naive or GARCH predictions. In addition, only for these models does the VaR standard turn out to be exceeded every year more

⁹ Even if the number of equilibrium relationships was restricted to one, the average squared prediction error, as well as the percentage of exceedances and the minimum amount needed, proved higher than without any equilibrium relationship.

often than in 1% of all cases. The reason for these results is that the implicit restrictions on the GARCH parameters do not apply to the changes in interest rates most of the time (see Table 4). Although the GARCH parameters are close to being integrated, over one third of the α_0 estimates are significantly larger than zero at the one percent level. According to a Wald test statistic the RiskMetrics parameterisation with β_1 equal to 0.94, has to be rejected at the 1% level for the random walk model in 59 out of 101 cases. For the specification with β_1 equal to 0.97 the 1% rejection rate is still 31 out of 101. The results for other mean specifications were similar. As a consequence of the neglect of the constant, possibly combined with too great an influence from the lagged disturbance terms, the conditional variance decreases too quickly after a few low prediction errors. Even if no great interest rate changes have occurred for some time, there is still a considerable probability that these may occur again some time in the future.

Finally, the differences between the random walk model and the models that make allowance for lagged interest rate changes seem minimal. Although the average squared prediction error is marginally higher for the random walk model, the minimum amount needed is nevertheless lower. The exceedance percentages are virtually identical, and the differences between the models with one or two lagged changes are negligibly small.

3.4. The variance–covariance method

The third method to determine VaR is the variance–covariance (VC) method. This method is the most widely applied in practice. As said above, this method can be only be applied to a prediction horizon of more than one day on the assumption of normality. Given the disappointing results for the t -distribution, there is no need to consider this a great disadvantage. In this method, the predicted interest rates over ten days are expressed, by means of substitution of forecasts of one to nine days ahead, in current and lagged interest rates and the disturbance terms for the following ten days. Given the variance–covariance matrix of the disturbance terms for the following ten days and the 1% critical value for the normal distribution, the VaR can subsequently be calculated. This calculation method is based on the independence of the disturbance terms over the various days, as a result of which the variance calculated over ten days equals the sum of the variances over the individual days. This assumption, however, is not met in the case of GARCH or the exponentially weighted average of squared disturbance terms as in these models the variance two or more days ahead depends on the size of the disturbance terms one day ahead. Independence only applies to the naive variance models. In practice, however, independence is assumed nevertheless, and the expected variances are calculated conditionally on the current information. This method has also been applied in Table 5. Given the consistently worse results for the

Table 4
 GARCH-parameters of the random walk model with normal distribution^a

sample	1 month	3 months	6 months	12 months	1–3 years	3–5 years	5–7 years	7–10 years
	$\alpha_0 (\times 1000)$							
80–84	1.998	0.251	0.152	<i>0.041</i>	0.001	<i>0.268</i>	0.147	0.134
81–85	0.971	0.244	0.157	0.054	0.028	0.214	0.184	0.101
82–86	0.860	<i>0.304</i>	0.389	0.061	0.060	0.335	0.187	0.077
83–87	0.279	<i>0.245</i>	0.350	0.082	0.035	<i>0.122</i>	0.161	0.073
84–88	<i>0.405</i>	0.307	1.226	0.885	0.001	0.041	<i>0.102</i>	0.034
85–89	0.209	0.260	0.499	<i>0.747</i>	0.000	0.040	<i>0.081</i>	0.033
86–90	<i>0.303</i>	0.217	0.662	0.605	0.000	0.023	0.054	0.031
87–91	0.222	0.185	<i>0.418</i>	0.583	0.013	<i>0.017</i>	0.032	0.026
88–92	0.198	0.261	<i>0.237</i>	<i>0.254</i>	<i>0.049</i>	0.031	0.039	0.022
89–93	0.179	<i>0.225</i>	0.287	<i>0.170</i>	0.039	0.028	0.336	0.026
90–94	0.148	<i>0.130</i>	0.218	0.221	0.037	0.035	0.042	0.037
91–95	<i>0.127</i>	0.101	0.141	0.249	0.027	0.039	0.057	0.057
92–96	0.028	0.005	0.000	0.135	0.044	0.054	0.078	0.074
	α_1							
80–84	0.152	0.049	0.026	0.024	0.035	0.074	0.071	0.095
81–85	0.125	0.068	0.041	0.027	0.019	0.070	0.098	0.087
82–86	0.110	0.091	0.083	0.034	0.011	0.062	0.081	0.067
83–87	0.077	0.114	0.095	0.039	0.001	0.096	0.118	0.093
84–88	0.097	0.135	0.199	0.209	0.011	0.068	0.089	0.063
85–89	0.092	0.116	0.099	0.158	0.003	0.070	0.109	0.072
86–90	0.053	0.083	0.093	0.109	0.011	0.048	0.061	0.040
87–91	0.040	0.065	0.074	0.087	0.031	0.049	0.069	0.044
88–92	0.017	0.045	0.027	0.023	0.078	0.050	0.057	0.034
89–93	0.010	0.024	0.013	0.013	0.070	0.056	0.061	0.039
90–94	0.015	0.028	0.018	0.023	0.079	0.064	0.063	0.045
91–95	0.024	0.030	0.023	0.026	0.059	0.066	0.065	0.053
92–96	0.011	0.008	0.005	0.018	0.053	0.063	0.061	0.056
	β_1							
80–84	0.763	0.937	0.965	0.974	0.967	0.872	0.895	0.869
81–85	0.825	0.915	0.948	0.970	0.978	0.878	0.851	0.882
82–86	0.825	0.884	0.881	0.960	0.979	0.823	0.839	0.896
83–87	0.895	0.866	0.859	0.950	0.989	0.852	0.799	0.867
84–88	0.843	0.825	0.550	0.606	0.990	0.910	0.836	0.913
85–89	0.882	0.852	0.796	0.701	0.996	0.906	0.820	0.901
86–90	0.903	0.886	0.756	0.757	0.988	0.933	0.889	0.932
87–91	0.918	0.899	0.819	0.773	0.956	0.933	0.896	0.930
88–92	0.947	0.904	0.920	0.918	0.883	0.924	0.900	0.944
89–93	0.953	0.928	0.923	0.945	0.895	0.917	0.900	0.936
90–94	0.953	0.940	0.927	0.915	0.897	0.910	0.908	0.932
91–95	0.935	0.938	0.935	0.897	0.927	0.908	0.898	0.911
92–96	0.979	0.990	0.994	0.935	0.923	0.907	0.895	0.902

^a Explanation: α_0 -values that are significantly different from zero, using heteroskedasticity consistent standard errors, at the 1% or 5% level are in bold face respectively italics.

Table 5
Evaluation of VaR standards on the basis of the variance–covariance method^a

Expectation	Variance	$\bar{\varepsilon}^2$	\bar{h}	$\overline{\text{VaR}}$	Surplus	%	Needed	Years >1%
Random walk	Naive	68.9	86.7	20.18	−0.57	1.200	20.75	5
	Expon. 0.94	68.9	55.8	15.66	−4.62	3.069	20.28	13
	Expon. 0.97	68.9	56.1	16.06	−3.90	2.524	19.96	12
	GARCH	68.9	73.9	18.69	−1.13	1.396	19.82	7
$p = 1$	Naive	67.7	97.5	20.96	0.04	0.968	20.92	4
	Expon. 0.94	67.7	67.2	16.91	−3.41	2.347	20.32	12
	Expon. 0.97	67.7	67.4	17.31	−2.88	1.960	20.19	10
	GARCH	67.7	81.3	19.33	−0.36	1.109	19.69	4
$p = 2$	Naive	67.2	100.0	21.18	0.46	0.890	20.72	4
	Expon. 0.94	67.2	70.2	17.25	−3.00	2.084	20.25	11
	Expon. 0.97	67.2	70.5	17.64	−2.54	1.738	20.18	10
	GARCH	67.2	84.4	19.65	−0.03	0.982	19.68	4

^a *Explanation:* Only models without cointegration have been included; see also Table 2.

models with cointegration (with respect to, among others things, the average squared prediction error, see Table 2), only models that have no equilibrium relationships have been included.

A comparison of the results of the Monte Carlo method and the VC method for naive variance models shows that the average size of the VaR standards and that of the conditional variances are virtually equal. The average number of exceedances of the VaR is nevertheless considerably smaller for the VC method. This is true especially for the models in which lagged interest rate movements are included. The calculated VaRs for these models prove sufficiently large in the VC method. Given the fact that the VC method calculates the theoretically correct VaR (assuming correctness of the estimated interest rate models), while the Monte Carlo method only provides an approximation, the difference in exceedances clearly indicates that 10 050 random samplings in the Monte Carlo method are still not sufficient to determine the real 1% exceedance probabilities. The random scatter of the simulated exceedance levels around the real levels (given the assumed distribution) causes a higher percentage of exceedances afterwards as the probability of an additional exceedance at too low a simulated value is greater than the probability of an exceedance less at too high a simulated value.¹⁰ The

¹⁰ This result underlines once more the problems in the historical simulation method (see Table 1). After all, in the historical simulation, the number of samplings is equal to the length of the history included. Moreover, if a ten-day VaR is computed, the results are even worsened by the fact that overlapping data are used.

consequences of incorrectly assuming independence of disturbance terms on different days can be seen in the difference in conditional variances between the VC and Monte Carlo methods. A comparison of Tables 2 and 5 makes clear that the VC method underestimates the real variance by up to approximately 8%. Partly for this reason, the exceedance percentages for the GARCH and exponentially weighted average methods are almost always too high.

3.5. Combined Monte Carlo variance–covariance method

One reason why the number of simulations for the Monte Carlo method must be so high is that in order to determine the VaR, only the 1% level is important. As a consequence, only a very limited number of simulations is really important for the calculation of the VaR. If the normal distribution is used as a basis, it is, however, possible to make more efficient use of the simulations. In that case, the distribution of the final result will also be normal. As the normal distribution is described entirely by the average and the variance, the variance which has been calculated from all simulations may be used. Given this variance and the normal distribution's 1% critical value, the VaR can subsequently be calculated. The results of this procedure are specified in Table 6.

A comparison of the results for the naive variance models in Tables 5 and 6 makes clear that 10050 Monte Carlo simulations are indeed sufficient if the

Table 6
Evaluation of VaR standards on the basis of the combined method^a

Expectation	Variance	VaR	Surplus	%	Needed	Years >1%
Random walk	Naive	20.18	-0.61	1.191	20.79	5
	Expon. 0.94	16.16	-4.21	2.815	20.37	13
	Expon. 0.97	16.32	-3.67	2.412	19.99	12
	GARCH	19.44	-0.50	1.239	19.94	7
$p = 1$	Naive	20.95	0.02	0.962	20.93	4
	Expon. 0.94	17.36	-3.04	2.144	20.40	12
	Expon. 0.97	17.54	-2.67	1.883	20.21	10
	GARCH	19.84	0.11	0.999	19.73	4
$p = 2$	Naive	21.16	0.45	0.890	20.71	4
	Expon. 0.94	17.68	-2.63	1.905	20.31	11
	Expon. 0.97	17.86	-2.31	1.665	20.17	11
	GARCH	20.12	0.41	0.880	19.71	4

^a *Explanation:* The variances were calculated on the basis of 10050 Monte Carlo simulations, after which the VaR standards were determined on the basis of the VC method.

VaR is calculated from the simulated variance. The average VaR size and the percentage of exceedances are virtually identical. If, therefore, the portfolio only comprises products which depend linearly on interest rates, this combined method is preferable to a pure Monte Carlo simulation. If, on the other hand, non-linear products comprise an important part of the portfolio, the error that is made in the exclusively linear inclusion of these products must be weighed against the error made by the Monte Carlo VaR standards' random fluctuations around the real value.

Using this calculation method, the VaR standards based on the exponent-weighted squared prediction errors are also insufficient. The use of this method is, therefore, not be recommended for interest rate portfolios.

A comparison of the naive models on the one hand, and the models with GARCH on the other, shows that the percentage of exceedances for the naive model is virtually equal to that for the GARCH model. The exceedance percentage for the naive method, however, has probably been more strongly influenced by the virtually continuous decrease in volatility over the sample period (see Fig. 1) as the GARCH models assign greater weight to more recent (and thus, within the sample, on average less volatile) observations. Thus, the GARCH models seem to inspire greater confidence than naive models for periods with increasing volatility, for instance in connection with uncertainty surrounding the euro. The minimum amount needed to prevent too high a percentage of exceedances is lower for the two GARCH models than for the naive model. This is another argument in favour of the use of GARCH models.

A comparison of the random walk models and the models with lagged interest rate changes leads to a slight preference for the latter category. The slightly lower average squared prediction error, combined with the higher conditional variance (see Table 2), leads to a lower percentage of exceedances. The error made by assuming the normal distribution, with too small a probability of extreme observations, is thus compensated for more effectively by a greater overestimation of the conditional variance. The amount needed to arrive at exactly 1% of exceedances is also slightly lower for the models including lagged changes, at least if a GARCH specification is modelled as well.

Table 7 provides a more detailed breakdown of the exceedance percentages per year. These results are no cause for great concern. The percentages, for instance, are not only low in the years in which the volatility decreased most (see Fig. 1), and there is also no series of consecutive years with too high a percentage of exceedances. The differences between the models are small. Not much can be said about the significance of the various percentages because both the results across the various portfolios and those over time are not completely independent due to overlapping observations.

Table 7
Percentage of exceedances per year of VaR standards on the basis of the combined method^a

<i>p</i>	0		1		2	
	Naive	GARCH	Naive	GARCH	Naive	GARCH
1985	0.74	1.17	0.37	1.27	0.29	1.27
1986	0.05	0.20	0.00	0.14	0.02	0.12
1987	1.21	0.92	0.98	0.77	0.92	0.75
1988	0.89	0.61	0.84	0.55	0.75	0.40
1989	1.54	1.99	1.32	1.74	1.17	1.31
1990	2.16	2.61	2.16	2.56	2.21	2.50
1991	0.31	0.35	0.31	0.29	0.15	0.25
1992	2.95	2.78	3.08	2.84	3.02	2.66
1993	0.74	1.03	0.34	0.49	0.08	0.25
1994	2.42	1.74	1.46	0.74	1.51	0.62
1995	0.60	0.62	0.32	0.26	0.28	0.22
1996	0.89	0.81	0.49	0.43	0.41	0.31
1997	0.09	1.58	0.09	0.47	0.09	0.37

^a Explanation: See Table 6.

4. Conclusions

This article has sought to indicate which elements are important in determining VaR standards, with a 99% confidence level and a holding period of ten trading days, for Dutch interest rate-related items. To this end, a great number of VaR standards for the period 1985–1997 were studied. On the basis of this analysis, the following conclusions can be drawn:

1. The differences in average size of VaR standards on the basis of different methods may vary to such an extent that the possible fluctuations in the add-on factor to determine the capital adequacy requirement are too small to actively encourage banks to accurately calculate the VaR. Accurate VaR calculations might very well lead to higher capital requirements than simple but inaccurate measures.
2. As regards the three permissible methods to determine the VaR, the following general conclusions may be drawn:
 - The historical simulation method regarding the period studied is only satisfactory if a long history is included. The fact that the VaR does provide sufficient cover for a long period is, however, mainly due to the decrease in volatility over the random sample. The average size of the VaR on the basis of historical simulation must be relatively large.
 - The Monte Carlo method requires a very great number of samplings (10 050 was too small) to arrive at the theoretically correct size of the 1% exceedance level.
 - The variance–covariance method (which can only be applied on the assumption of normality) works well for models with a naive variance.

For GARCH-like variance specifications, however, the method leads to an underestimation of the real variance, resulting in too large a percentage of exceedances.

3. For models with a time-varying variance and a normal distribution, a combined Monte Carlo variance–covariance method, in which the variance is derived from the Monte Carlo simulation and the VaR is subsequently calculated on the basis of the variance–covariance methods, provides good results.
4. As regards model assumptions that underlie the Monte Carlo and variance–covariance methods, the following can be said:
 - models which make allowance for equilibrium relationships between interest rates in the determination of the expected interest rate conditions provide considerably worse VaR standards;
 - modelling the expected interest rate changes as a function of lagged interest rate changes provides slightly better forecasts and slightly lower exceedance percentages. The differences between one and two lags are minimal;
 - VaR standards based on models with variance specifications computed as exponentially weighted averages of squared prediction errors (RiskMetrics) are structurally exceeded too often;
 - models with naive or GARCH variance specifications do provide correct VaR standards. For the naive models in particular, however, this result is due in part to the decreased volatility over the analysed period;
 - the *t*-distribution provides much worse VaR standards than the normal distribution. This is probably the consequence of the fact that the extreme observations are not representative of the total stochastic process. This does not, however, detract from the fact that the number of extreme observations is higher than predicted by the normal distribution.

Of course, one should keep in mind that these results are only investigated for Dutch interest rate related items, so they need not hold for international portfolios or share portfolios. Given the close relationship between Dutch and German interest rates, it is most likely that similar results hold for German interest rates. It would be interesting to investigate other items as well to see whether these results can be generalised.

Acknowledgements

I would like to thank André Lucas for kindly providing me with his GAUSS-programs to compute the optimal number of cointegrating relationships under non-normal distributions. Furthermore, the useful comments by an anonymous referee, seminar participants at De Nederlandsche Bank, Peter van Els and André Lucas are gratefully acknowledged.

Table 8
Summary statistics of interest rates by maturity for three sub-samples^a

	1 month	3 months	6 months	12 months	1–3 years	3–5 years	5–7 years	7–10 years
<i>1980 up until 1984</i>								
ADF	-1.43	-1.29	-1.21	-0.97	-2.54	-0.36	-0.49	-0.49
s.d.	14.65	13.10	12.79	12.36	12.27	6.74	6.13	5.88
m3	-0.92	-0.38	-0.27	-0.36	-1.20	0.12	-0.26	-0.22
m4	16.07	8.56	5.90	9.66	16.22	3.52	4.27	4.21
ρ	0.05	0.00	0.00	-0.04	0.05	0.24	0.30	0.30
Qe(25)	27.64	14.92	24.24	38.60	34.79	119.02	135.81	140.35
Qee(25)	170.47	156.28	139.97	62.78	85.08	236.60	253.23	365.74
Correlation matrix	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	0.72	0.76	0.66	0.14	0.28	0.87	0.82	
	0.65	0.59	0.13	0.32	0.27	0.79		
	0.56	0.12	0.33	0.31	0.25			
	0.15	0.28	0.33	0.32				
	0.26	0.31	0.33	0.31				
	0.26	0.29	0.31					
	0.25							
<i>1986 up until 1990</i>								
ADF	0.07	0.12	0.19	0.11	-0.11	-0.44	-0.41	-0.44
s.d.	7.90	7.93	6.64	6.51	4.21	3.65	3.75	3.68
m3	0.48	0.03	-0.06	-0.33	4.98	0.71	-1.49	-0.35
m4	18.69	16.51	5.69	7.25	121.70	23.20	35.63	13.24
ρ	-0.11	-0.20	-0.14	-0.14	0.12	0.19	0.17	0.17
Qe(25)	27.09	39.08	35.31	35.16	60.68	82.62	33.31	30.88
Qee(25)	108.06	524.32	107.49	98.59	0.69	89.73	226.74	359.17

Table 8 (Continued)

	1 month	3 months	6 months	12 months	1–3 years	3–5 years	5–7 years	7–10 years
Correlation matrix	1.00							
Δr	0.64	1.00						
s.d.	0.55	0.67	1.00					
Δr	0.47	0.56	0.76	1.00				
m3	0.12	0.20	0.20	0.23	1.00			
m4	0.18	0.25	0.26	0.26	0.44	1.00		
ρ	0.17	0.27	0.28	0.30	0.50	0.78	1.00	
Qe(25)	0.19	0.28	0.28	0.30	0.49	0.77	0.91	1.00
Qee(25)								
1992 up until 1996								
ADF	-1.47	-1.54	-1.71	-1.70	-1.53	-1.44	-1.38	-1.39
r	5.20	5.20	5.35	5.42	3.96	4.22	4.24	4.20
Δr	-0.79	-0.88	-1.67	-1.68	0.01	0.33	0.36	0.42
m3	30.64	16.65	19.54	19.73	11.39	4.09	3.48	4.01
m4	0.00	-0.02	-0.03	0.02	0.23	0.22	0.17	0.13
ρ	19.89	24.51	18.16	20.54	86.61	75.01	60.95	47.70
Qe(25)	131.34	88.27	70.39	110.05	108.97	368.90	477.61	482.16
Qee(25)								
Correlation matrix	1.00							
Δr	0.80	1.00						
s.d.	0.72	0.76	1.00					
Δr	0.63	0.67	0.79	1.00				
m3	0.28	0.27	0.29	0.35	1.00			
m4	0.16	0.17	0.18	0.23	0.81	1.00		
ρ	0.12	0.12	0.14	0.18	0.74	0.93	1.00	
Qe(25)	0.09	0.09	0.10	0.14	0.67	0.86	0.95	1.00

^a Explanation: For every sub-sample the following statistics are given: an Augmented Dickey–Fuller test statistic for stationarity of the interest rate levels (ADF), the standard deviation (s.d.), the skewness parameter (m3), the excess kurtosis (m4), the first order autocorrelation coefficient (ρ), a Ljung–Box test statistic for 25th order autocorrelation (Qe(25)), both adjusted for GARCH-like heteroskedasticity (see Diebold, 1987), a Ljung–Box test statistic for 25th order autocorrelation in de squared interest rate changes (to detect GARCH-effects) (Qee(25)) and the correlation matrix of the interest rate changes of different maturities. Values that are significantly different from zero, assuming normality, at the 1% or 5% level are in bold face respectively italics.

Appendix A. Summary statistics of interest rates by maturity for three sub-samples

Table 8 presents the summary statistics of interest rates by maturity for three sub-samples.

References

- Anderson, H.M., Granger, C.J.W., Hall, A.D., 1992. A cointegration analysis of treasury bill yields. *Review of Economics and Statistics* 74, 116–126.
- Basle Committee on Banking Supervision, 1994. Risk Management Guidelines for Derivates. Bank for International Settlements, Basle.
- Basle Committee on Banking Supervision, 1996. Supplement to the Capital Accord to Incorporate Market Risks. Bank for International Settlements, Basle.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Bollerslev, T., 1990. Modeling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *The Review of Economics and Statistics* 72, 498–505.
- Bollerslev, T., Chou, R.Y., Kroner, K., 1992. ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics* 52, 5–59.
- Bradley, M.G., Lumpkin, S.A., 1992. The treasury yield curve as a cointegrated system. *Journal of Financial and Quantitative Analysis* 27, 449–463.
- Campbell, J.Y., Shiller, R.J., 1984. A simple account of the behavior of long-term interest rates. *American Economic Review* 74, 44–48.
- Campbell, J.Y., Shiller, R.J., 1987. Cointegration and tests of present value models. *Journal of Political Economy* 95, 1063–1088.
- Danielson, J., 1997. Extreme returns, tail estimation and Value-at-Risk. In: The 52th ESEM conference, Toulouse.
- Diebold, F.X., 1987. Testing for serial correlation in the presence of ARCH. In: Proceedings from the American Statistical Association, Business & Economics Section, pp. 323–328.
- Dimson, E., Marsh, P., 1995. Capital requirements for securities firms. *Journal of Finance* 50, 821–851.
- Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987–1007.
- Estrella, A., Hendricks, D., Kambhu, J., Shin, S., Walter, S., 1994. The price risk of options positions: Measurement and capital requirements. Federal Reserve Bank of New York Quarterly Review 19, 27–43.
- Fama, E.F., 1965. The behavior of stock market prices. *Journal of Business* 38, 34–105.
- Figlewski, S., 1994. Forecasting volatility using historical data. Working Paper 13. New York University, New York.
- Hendricks, D., 1996. Evaluation of Value-at-Risk models using historical data. Federal Reserve Bank of New York, Economic Policy Review, pp. 39–69.
- Johansen, S., 1988. Statistical analysis of cointegrating vectors. *Journal of Economic Dynamics and Control* 12, 231–254.
- JP Morgan, 1996. RiskMetrics Technical Document, 4th ed. New York.
- Lucas, A., 1997a. Cointegration testing using pseudo likelihood ratio tests. *Econometric Theory* 13, 149–169.

- Lucas, A., 1997b. Robustness of the Student t based M-estimator. *Communications in Statistics; Theory and Methods* 26.
- Shea, G.S., 1992. Benchmarking the expectations hypothesis of the interest-rate term structure: an analysis of cointegration vectors. *Journal of Business & Economic Statistics* 10, 347–366.
- Taylor, M.P., 1992. Modelling the yield curve. *Economic Journal* 102, 524–537.
- Vlaar, P.J.G., Palm, F.C., 1993. The message in weekly exchange rates in the European monetary system: Mean reversion, conditional heteroskedasticity and jumps. *Journal of Business & Economic Statistics* 11, 351–360.