The Impact of Stock Market Volatility on Corporate Bond Credit Spreads

P. Berg[a], R. Bewley[b] and D. Rees[a]

[a] *Commonwealth Bank of Australia, Sydney 1155 (bergpa@cba.com.au) and (reesda@cba.com.au)*

[b] *School of Economics, The University of New South Wales, Sydney 2052, (r.bewley@unsw.edu.au)*

Abstract: There has been a rapid increase in the number of corporate bonds issued in Australia since the middle of 1998. This increase has stimulated interest in characterising the yield curves and the factors that determine changes in these spreads. The focus of this paper is on measuring any impact of stock market volatility on spreads using two different measures. One measure is based on volatility implied from options prices while the other is derived from a conditional heteroscedastic volatility model of changes in a stock market index. It is found that the former has no significant impact on spreads but the latter is both significant and stable over time. This impact is estimated to be negative implying that an increase in volatility cause a decrease in corporate bond spreads.

Keywords: Forecasting; GARCH; Implied volatility; Time series analysis; VAR

1. INTRODUCTION

It is only since the middle of 1998 that the number of corporate bonds on issue has been sufficiently large to enable an econometric characterisation of credit spreads to be undertaken across a broad range of credit ratings and industries. In a previous line of research, Berg et al. [2000] fitted yield curves to daily data on spreads to Commonwealth Government Securities for nine (Standard and Poor's®) credit ratings categories (AAA to BBB) and five industry categories.

The model constructed by Berg et al. [2000], and published under the name CBASpectrum®, is estimated separately for each day in an ongoing basis so that the time series of estimated spreads that flow from this model can reasonably be used as data input for this research. Appropriate time series methods can then be used to gain insight into how spreads evolve over time and, in particular, into how they react to market forces external to the corporate bond market.

Since the family of estimated yield curves is characterised by twelve parameters, one strategy is to model the evolution of the estimated parameters. Indeed, such an approach implies that the full system could be modelled by a finite number of equations whereas the direct modelling of spreads, as is done in this paper, requires that an arbitrary finite subset of an infinite number of ratings/duration combinations be considered. In the first instance, the direct modelling approach has been followed as it has far greater intuitive appeal and can reasonably limited to modelling a smaller number of ratings/duration combinations than the number of parameters in the system.

While credit spreads can be constructed from actual data, such an approach would suffer from the consequent evolution in duration of the selected bonds over time and any particular changes in the characteristics of these bonds. These problems are removed by implying a credit spread for a generic bond of fixed duration and rating over time. The 'bonds' chosen for this study are AAA, AA, A and BBB, each of a constant five year duration, and the spreads are measured as deviations from the estimated Commonwealth Government Securities yield curves using CBASpectrum®.

Two possible external influences on the corporate bond market are considered in this paper: the stock market, as measured by the ASX Large Capitalisation Index, and implied volatility. While 425 daily observations are used in the analysis of bond spreads: 1st July 1998 to 16th March 2001, 800 daily observations are used to measure volatility so as to remove any end-point effects due to the estimation procedure.

2. STOCK MARKET VOLATILITY

It is now commonplace to measure volatility in financial time series using Engle's [1982] ARCH or Bollerslev's [1986] GARCH models. These models are based on the notion that the innovations of a time series process unconditionally have a fixed variance, but that volatility clustering occurs in the sense that the conditional variance of a process varies over time.

A GARCH(p,q) model can be expressed as

$$h_t = \alpha_0 + \Sigma_1^p \alpha_i \varepsilon_t^2 + \Sigma_1^q \beta_j h_{t-j} + \Sigma_1^m \psi_l g_{lt} \qquad (1)$$

where ε_t are the innovations in the levels model, h_t is the conditional variance, and g_{lt} factors that determine changes in conditional volatility.

Autoregressions were considered for the log first-difference of the index but no such terms were significant. Indeed, a constant term was not significant and, therefore, was excluded. Thus, ε_t is the change in the log of the index. The preferred model is a GARCH(1,1). The estimated parameters, with approximate t-ratios, is presented in Table 1.

Table 1. Estimated GARCH(1,1) Model.

Parameter	Estimate
α_0	0.0588 (2.78)
α_1	0.1034 (3.46)
β_1	0.7999 (14.76)

The implied conditional standard deviations $h_t^{1/2}$, from this model are depicted in Figure 1. It shows some possible initial value problem due to having to estimate $h_0^{1/2}$. The impact of an outlier at 18th April 2000 (Day 573) can also be noted. Owing to the scale of this figure it is difficult to note that the conditional standard deviation more than doubles going from day 572 to 573 but it takes approximately 10 days to return to the level of day 572 owing to the persistence in the estimated GARCH model.

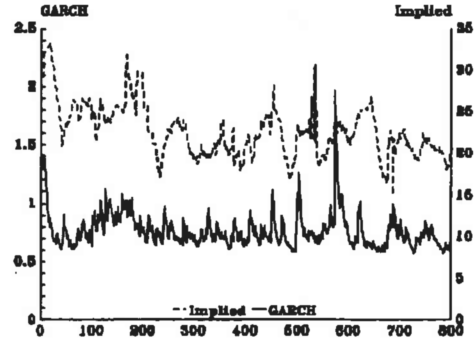

Figure 1. GARCH and Implied Volatility.

Implied volatility is an alternative measure of uncertainty in the stock market based on options prices and it can also be noted from Figure 1 that it has very different time series characteristics from the GARCH variant. Implied volatility takes a longer term view of uncertainty than a GARCH measure, and it is forward looking, rather than being an instantaneous measure of volatility, as is the GARCH estimate.

The question posed in this paper is whether stock market volatility, as represented by either the GARCH model, or implied volatility, impact on the estimated credit spreads generated by CBASpectrum®.

3. CREDIT RATINGS SPREAD MODELS

The estimated spreads for the four generic bonds of five years to maturity are presented in Figure 2. It can be noted that each spread has undergone a number of changes in the sign of the local trend and the relative spreads have also varied substantially during that period.

Unit root tests were performed on each series using the Augmented Dickey-Fuller [1979] test, including a constant term and a time trend, on the last 425 observations. The lag length was chosen by choosing using the AIC. In results not reported, the null hypothesis of a unit root could not be rejected at any reasonable level of significance for any time series. Moreover, the estimate of the largest root is numerically close to one in each case. In the absence of other information, it is worthwhile proceeding as if each series is integrated to order 1, I(1).

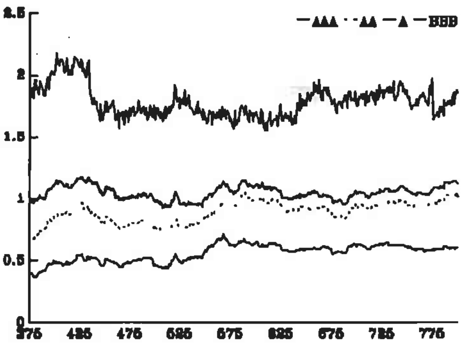

Figure 2. Estimated Spreads.

It is of some interest to investigate whether or not the four time series of credit spreads are cointegrated. However, it should be noted that the outcome of such tests depend on the ratings and durations of the generic bonds chosen for inclusion in the system. The closer are the choices of ratings bands and/or the durations, the more likely would two series be expected to move closely together.

In results also not reported, Johansen's [1988] tests provide borderline support for cointegration at the 5% level. However, the single candidate cointegrating vector had no obvious economic interpretation. It was resolved that this characteristic was more due to the generic bonds being of the same duration than any long term economic relationship. As a result a Vector Autoregression (VAR) in the first differences of the spreads was selected.

4. EQUITY VOLATILITY AND SPREADS

The main problem of performing a joint time series analysis of the four changes in credit spreads and one of the conditional standard deviation series, or the implied volatility, in a five equation VAR is that this would imply that the lagged credit spreads affected the conditional standard deviations in contradiction of the model reported in Table 1. The GARCH process was re-estimated incorporating some such lagged effects as g_{it} variables in the notation of (1). However, it was not possible to allow for more than one lag in (1) without numerical problems arising. With one lag, there were no significant ψ_i so that the GARCH model of Table 1 need not be modified. Moreover, this implied that the appropriate VAR is a four equation model of the differences of the spreads but with each equation augmented by the lags of volatility time series.

The volatility measures were allowed to enter contemporaneously since the GARCH estimates depend only upon the lagged log-differences of the stock market index through ε_{t-i} in (1). For equivalence, implied volatility also enters from lag 0. The AIC selected a lag of three in the levels for the four equation VAR including either volatility measure. Although the GARCH series is not I(1), the parameter estimates strongly suggested that differencing of this variables is appropriate so that, in the final model, two lags of each first-differenced spread are used with the current first-difference of the volatility measure entering exogenously.

Since volatility is a critical variable in this analysis, the maximum lag for this variable was allowed to be different from that for the other four variables. Maximum lags of 0 to 2 in the differences were used in turn for this variable, along with two lags of the other (differenced) variables. The AIC was used to select this maximum lag length and zero was chosen for both volatility measures.

Table 2. t-ratios on Volatility Variables.

	Spread			
Variable	AAA	AA	A	BBB
GARCH	-3.42	-2.40	-1.92	-0.73
Implied	1.02	0.87	0.39	0.56

When each volatility measure enters separately, the GARCH-variant t-ratios, shown in Table 2, are negative and progressively less significant as the rating level rises. On the other hand, the implied volatility measure is everywhere positive and insignificant. When both measures are simultaneously included, there was no change in these conclusions about the two measures.

Since the two volatility measures provide different conclusions it is appropriate to further analyse the models to look for possible structural instability or influential data. In order to gain some insight into the stability of the model, Bewley and Yang's [2000] long-run CUSUM tests were applied to each equation. This test differs from the standard CUSUM test in that it is applied to the Bewley [1979] transform of the VAR and, as a result, breaks are analysed in the means of the time series rather than in the constants of a VAR which depend on all of the means and all of the lagged dependent variable coefficients. Bewley and Yang show that the long-run test is more powerful that its short-run counterpart in most areas of interest.

Since no long-run CUSUM trace breaks its critical value, there is no evidence of any structural instability from these tests. The long-run recursive coefficients, using the Bewley transformation of the VAR, on the volatility measure were also considered.

The largest estimate of volatility occurred on day 573 so that it is of some interest to examine the recursive traces for changes around that day. While there is some mild evidence of sensitivity to this observation in the BBB spreads, there is no evidence elsewhere. Indeed, the traces are remarkably flat over the last three-quarters of the sample. Perhaps what is more important is that the coefficients do not change in the neighbourhood of day 573 but the confidence band narrows considerably. This behaviour is consistent with the largest volatility observation being consistent with the rest of the data.

Table 3. One-Step Forecast Accuracy.

Model	AAA	AA	A	BBB
With Vol	0.783	1.416	0.977	27.972
No Vol.	0.810	1.441	1.012	27.370

As a further diagnostic test, one-step ahead forecasts were generated for the last 300 days, updating the estimates of the coefficients at each origin. Although from Table 2, the volatility measure is significant at the 5% level in two of the equations, and at the 10% level in the A equation, there was no real gain in mean square error as can be noted from Table 3 which compares the forecasting ability of each credit spread in the preferred VAR, including GARCH volatility, against a benchmark model that excludes this variable.

The lack of a major improvement in forecast accuracy naturally raises the question of why none was found. It is not uncommon in econometrics for variables to be significant within sample but not to contribute to forecasting performance. However, the stability and significance of the volatility coefficients over time is not consistent with this phenomenon. There being no obvious route to follow on further analysing this question, the symmetry of the impact of changes in volatility (in terms of increases and decreases) was generalised by allowing for an interactive dummy on the change in the volatility. Although the symmetric model did not include a constant, the option of including one in the asymmetric model was considered. The results with and without a constant are given in Table 4.

Table 4. Comparison of Impact of Increases and Decreases in Volatility.

Variable	With	Without
AAA		
Inc.	-0.0215 (-3.14)	-0.0193 (-3.00)
Dec.	-0.0126 (-0.63)	-0.0248 (-1.65)
AA		
Inc.	-0.0210 (-2.30)	-0.0174 (-2.03)
Dec.	-0.0073 (-0.28)	-0.0268 (-1.33)
A		
Inc.	-0.0236 (-2.26)	-0.0201 (-2.05)
Dec.	0.0167 (0.55)	-0.0020 (-0.09)
BBB		
Inc.	-0.0275 (-0.70)	-0.0292 (-0.79)
Dec.	-0.0098 (-0.09)	-0.0009 (-0.01)

Clearly, the increases are much better determined than the falls and this could point to an asymmetric effect. However, there is much similarity between the estimated coefficients and the lack of significance of decreases in volatility could be due to lack of variability. While there are some large shocks to the stock market producing apparently sharp peaks, the persistence in the GARCH model causes the falls to be much slower than the increases. Since a return to a period of low volatility in a symmetric model would force forecasts to necessarily return to previous levels, there is some merit in further considering the one-sided model that only includes the upside term. To allow for a non-zero mean in this one-sided term and no constant in the regressions, the mean was removed from the positive increases in volatility.

Using AAA spreads as an example, forecasts from the different models are considered for the most volatile period in the sample. Assuming that future changes in the stock market are known, estimated coefficients based only on prior data are given in Figure 3 using the symmetric model. It can be noted that a downturn is predicted but that the return to previous levels is too quick compared to the actual.

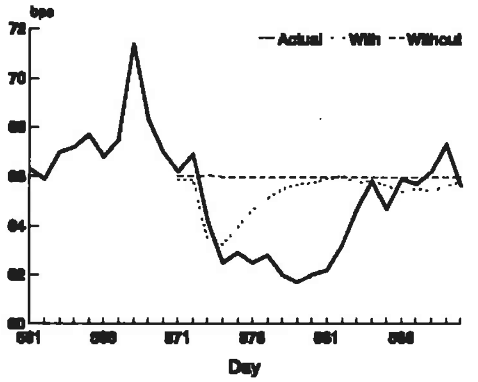

Figure 3. Forecasts from Symmetric Models With and Without Volatility.

A number of forecasts of AAA spreads from successive origins in the same period is given in Figure 4 from the model with no volatility variables included and in Figure 5 for the model that only includes increases in volatility.

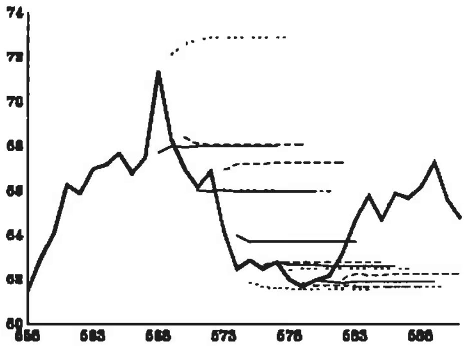

Figure 4. Forecasts from Different Origins and Actuals Using Models that Exclude Volatility.

Forecasts from the model without volatility, although having two lags, behave similarly to a naive no-change model; the forecast for all future time is close to the last observation. On the other hand, there is a richness in the one-sided model forecasts that seem to give quite plausible results in AAA spreads during this most volatile of periods. In times of low volatility, or falling volatility, the forecasts from this model will be similar to those from the model that excludes volatility. Similar forecast performance is noted for the other three spreads.

The main problem with such a model is that the results may be too dependent on the one large observation on volatility (Day 573). In an alternative approach to diagnostic testing, the observations were re-ordered by the size of the volatility measure, but, of course, keeping the appropriate lags in sequence. In this way, the estimates using all of the data are necessarily the same as before but recursive techniques now focus on the size of the changes rather than the chronological ordering. Since there a number of zeroes corresponding to the falling volatility measure, this ordering is not unique.

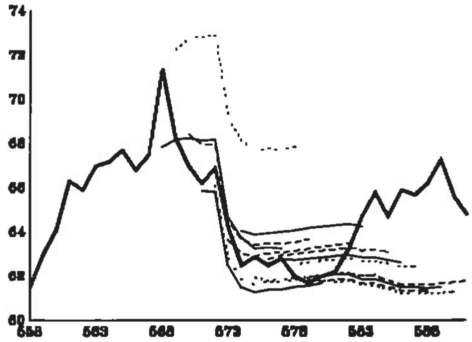

Figure 5. Forecasts from Different Origins Using the Asymmetric Volatility Model.

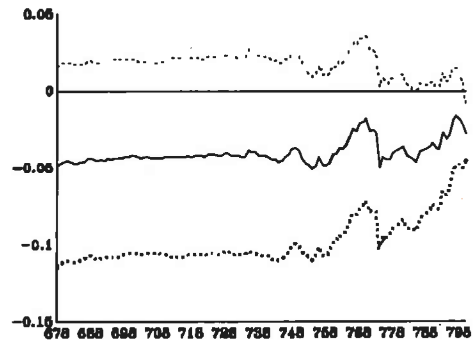

Figure 6. Recursive Coefficients with 95% Confidence Intervals Estimated from Data Ranked by Volatility, AAA Bonds.

One of these recursive traces is given in Figure 6 for the coefficient on volatility in the AAA equation. It is deceptively flat in the early part of the re-ordered sample because the volatility variable is essentially a constant term until the bigger shocks appear. Nevertheless, the impact is quite constant over the re-ordered sample. If the mean is not removed from the one-sided variable and a constant is included, the process takes longer to settle down but still exhibits a similar pattern in the last 50 or so observations. Thus, it is argued that the significance of the GARCH-based volatility variable in the model is not due just to the highly volatile period but there is a reasonably similar effect for the less volatile periods.

As a final piece of analysis, the model with a constant and a symmetric volatility measure was augmented by the lagged change in the log of the Commonwealth Government Securities yield, $\Delta\ln(CGS)$. This variable was included to consider whether overall shifts of all of the yield curves impact on the magnitude of the spreads. The coefficients and t-ratios on the GARCH-based volatility variable and $\Delta\ln(CGS)$ are presented in Table 5.

There is a reasonably consistent positive impact on the yield curves from $\Delta\ln(CGS)$ and it is significant at the 5% level in two equations, that of AA and A. However, there is no major impact on the size or significance of the volatility estimates.

5. CONCLUSIONS

The credit spread is the additional yield that is priced into a bond as a compensation for various risks (default, credit migration) and for illiquidity. It has been found in this paper that implied stock market volatility, derived from options prices, has no significant impact on these spreads on a day-to-day basis.

Table 5. Impact of Changes in CGS Yields on Spreads.

Equation	Volatility	$\Delta\ln(CGS_{-1})$
AAA	-0.0187	0.0151
	(-3.16)	(1.91)
AA	-0.0167	0.0233
	(-2.11)	(2.21)
A	-0.0150	0.0246
	(-1.66)	(2.04)
BBB	-0.0190	0.0609
	(-0.56)	(1.33)

On the other hand, an instantaneous measure of market volatility, derived from a GARCH model, has a significant negative impact on these spreads. That is, an increase in market volatility causes a narrowing of spreads to Commonwealth Government Securities.

A number of diagnostic tests were performed on the model which established the robustness of the model to outliers and structural change over time. However, it was found that a one-sided model that only passed on increases in volatility to spreads provides more realistic forecasts.

6. REFERENCES

Berg, P., R. Bewley and D. Rees, Yield Curves for Australian Corporate Bonds, mimeo, 2000

Bewley, R.A., The Direct Estimation of the Equilibrium Response in a Linear Dynamic Model, *Economics Letters*, 3, 357-361, 1979.

Bewley, R. and M. Yang, Testing for Structural Breaks in the Long-Run Means of a VAR, paper presented to the New Zealand Econometrics Study Group, July 2000.

Bollerslev, T., Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics*, 31, 307-327.

Dickey, D.A. and W.A. Fuller, Distribution of the Estimator for Autoregressive Time Series with a Unit Root, *Journal of the American Statistical Association*, 74, 427-31, 1979.

Engle, R.F., Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of UK Inflation, *Econometrica*, 50, 987-1007, 1982.

Johansen, S., Statistical Analysis of Cointegrating Vectors, *Journal of Economic Dynamics and Control*, 12, 231-254, 1988.