

*Testing Stochastic Interest Rate Generators for
Insurer Risk and Capital Models*

Gary G. Venter, FCAS, MAAA

Testing Stochastic Interest Rate Generators for Insurer Risk and Capital Models

Gary G Venter, Guy Carpenter InStrat

Stochastic models for interest rates are reviewed and fitting methods are discussed. Tests for the dynamics of short-term rates are based on model fits. A method of testing yield curve distributions for use in insurer asset scenario generators is introduced. This uses historical relationships in the conditional distributions of yield spreads given the short-term rate. As an illustration, this method is used to test a few selected models.

Acknowledgement

I would like to acknowledge the invaluable assistance of Andrei Salomatov in both preparing the exhibits and helping to develop the methodology for this paper.

Testing Stochastic Interest Rate Generators for Insurer Risk and Capital Models

P&C insurers are looking to financial modeling to address how risk is diversified among assets, liabilities and current underwriting results. Before fast computer models, actuaries measured asset risk by a few simple constants, like duration and convexity. Asset managers have their own collection of risk measurement constants for hedging issues, identified by Greek letters, and so often referred to as “the Greeks.” Appendix 1 summarizes these measures.

These asset risk scalars typically measure the sensitivity of the asset portfolio to changes in some particular risk event, such as a change in the average interest rate, or a change in the volatility of interest rates. With stochastic generators, however, two degrees of specificity are added. First the dimension of probability of risk events is incorporated. Risk scalars show the sensitivity to a change but not the probability of that change. Second is the response to a much broader range of possible risks. Complex combinations of risk situations can occur, and stochastic modeling can quantify the combined risk picture.

These added dimensions come from representing the distribution of possible outcomes for an asset portfolio. Models can then combine asset outcomes with liability development and underwriting return outcomes to give a more comprehensive risk profile. Asset models generate a large variety of asset scenarios, ideally each showing up by the probability of its occurrence, and apply them to the asset portfolio to measure the distribution of asset risk.

Although useful and general in theory, the possible weakness of this approach is that in practice the model might not capture the full range of economic outcomes, or it could over-weight the chances of some occurrences that are in fact not all that likely to happen. Thus a significant risk to this methodology is generating the wrong distribution of financial events.

This paper looks at evaluating interest rate generators by testing the distribution of yield curves. Empirical research on the dynamics of the short-term rate is reviewed, then tests of the generated distribution of yield curves are introduced and applied to a few models. Interest rate models in other areas of finance tend to be used to price options, so they are evaluated on how well they can match option prices. Insurer models are more focused on the risks inherent in holding vari-

ous asset mixtures for a period of time, and sometimes on trading strategies, and so realistic distributions of ending yield curves and probabilities for movements in yields are of more direct concerns than are option prices.

1. Models of Interest Rates

The primary focus here is on arbitrage-free models of interest rates. There is still some debate among actuaries on whether this is the best approach, and some of this debate is summarized in Appendix 2, but it is such models which will be emphasized here. The tests on the yield curve distributions introduced below, however, can be used on any model that generates yield curve scenarios. Interest rates are further assumed to be default free. Modeling default probabilities adds a degree of complexity that is not addressed here.

There are a few ways to generate arbitrage-free interest rate scenarios. The method illustrated below models the short-term interest rate, denoted by r , directly, and uses the implied behavior of r , along with market considerations, to infer the behavior of longer-term rates. For these models, r is usually treated as a continuously fluctuating process. This is somewhat of an approximation as actual trades occur at discrete times, but at scales longer than a few minutes it seems appropriate, at least during trading hours.

The most common financial models for continuous processes are based on Brownian motion. A Brownian motion has a simple definition in terms of the probabilities of outcomes over time: the change in r from the current position between time zero and time t is normally distributed with mean zero and variance $\sigma^2 t$ for some σ . In differential notation, the instantaneous change in r is expressed as $dr = \sigma dz$. Here z represents a Brownian motion with $\sigma=1$, and so its variance after a time period of length t is just t . If r also has a drift (i.e., a trend) of bt during time t , the process could be expressed as $dr = bdt + \sigma dz$.

Cox, Ingersoll and Ross (*A Theory of the Term Structure of Interest Rates* *Econometrica* 53 March 1985) provide a model of the motion of the short-term rate that has been widely studied. In the CIR model, r follows the following process:

$$dr = a(b - r)dt + sr^{1/2}dz.$$

Here b is the level of mean reversion. If r is above b , then the trend component is negative, and if r is below b it is positive. Thus the trend is always towards b . The speed of mean reversion is expressed by a . Note that the volatility depends on r itself, so higher short-term rates would be associated with higher volatility. Also, if $r=0$ there is no volatility, so the trend takes over. With $r=0$ the trend would be positive, so r would move to a positive value. The mean reversion combined with rate-dependent volatility thus puts a reflective barrier at $r=0$.

If this model were discretized it could be written:

$$r_t - r_{t-1} = a(b - r_{t-1}) + sr_{t-1}^{1/2}\epsilon, \text{ where } \epsilon \text{ is a standard normal residual.}$$

This is a fairly standard autoregressive model, so the CIR model can be considered a continuous analogue of an autoregressive model.

Some other models of the short rate differ from CIR only in the power of r in the dz term. The Vasicek model takes the power to be zero. Another choice is taking a power of unity.

Most of the models incorporate mean reversion, but constant mean reversion is problematic. The rates sometimes seem to gravitate towards a temporary mean for a while, then shift and revert towards some other. One way to account for this is to let the reversion mean b itself be stochastic. This can be done by adding a second stochastic equation to the model:

$$db = j(q - b)dt + wb^{1/2}dz_1$$

Here dz_1 is a second, independent standard normal variate, and so b follows a mean reverting process gravitating towards q . Again different powers can be taken for b in the stochastic term. Such two factor models are popular in actuarial literature. For instance, Hibbert, Mowbray and Turnbull in , "A Stochastic Asset Model & Calibration for Long-Term Financial Planning Purposes," Technical Report, Barrie & Hibbert Limited, use a two factor model which generalizes the Vasicek model by taking b and r both to the zero powers, so they both drop out of the stochastic terms.

The volatility can also be stochastic. For instance, Hull, J. and A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatilities," The Journal of Finance, XLII, 2, pp. 281-300 consider such a model.

Combining stochastic volatility and stochastic mean reversion, Andersen and Lund (Working Paper No. 214, Northwestern University Department of Finance) use the model:

$$dr = a(b - r)dt + sr^k dz_1 \quad k > 0$$

$$d \ln s^2 = c(p - \ln s^2)dt + v dz_2$$

$$db = j(q - b)dt + wb^{1/2} dz_3$$

This model uses three standard Brownian motion processes, z_1 , z_2 , and z_3 . The volatility parameter s^2 now also varies over time, but via a mean reverting geometric Brownian motion process (i.e., Brownian motion on the log). In total there are eight parameters: a , c , j , k , p , q , v , and w and three varying factors r , b , and s . It is thus labeled a three-factor model. The power k on r in the stochastic term is a parameter that can be estimated.

2. Dynamics of Short-Term Rates – Empirical Findings

Estimation of model parameters should be distinguished from calibration to current states. The permanent parameters of the models are estimated from historical data, whereas the variable factors are re-calibrated to current yield curves to capture the latest market conditions. Different techniques might be used for estimation vs. calibration.

Multi-factor Brownian motion models can be difficult to estimate. Some single-factor models, such as CIR, can be integrated out to form a time series, which can be estimated by maximum likelihood. In the case of CIR, the conditional distribution of the short rate at time $t+T$ given the rate at time t follows a non-central chi-squared distribution:

$$f(r_{t+T} | r_t) = ce^{-u-v(v/u)q/2} I_q(2(uv)^{1/2}), \text{ where}$$

$$c = 2as^{-2}/(1-e^{-aT}), \quad q = -1 + 2abs^{-2}, \quad u = cr_t e^{-aT}, \quad v = cr_{t+T} \text{ and } I_q \text{ is the modified Bessel function of the first kind, order } q, \quad I_q(2z) = \sum_{k=0}^{\infty} z^{2k+q} / [k!(q+k)!], \text{ where factorial of integers is defined by the gamma function}$$

This is not usually possible for multi-factor models, where the volatility and other factors can change stochastically. Further, the short-term rate is observed, or is closely related to observed rates for very short terms, but the other factors, like the reverting mean and the volatility scalar,

are not typically observed. Thus fitting techniques that match models to data will not be applicable for these factors.

A few fitting techniques have been developed for stochastic processes. The general topic of what these techniques are and how they work is beyond the scope of this paper, but one method which has been used successfully – the efficient method of moments (EMM) – is briefly discussed below. This method was introduced by Gallant, A. and Tauchen, G.: 1996, Which moments to match?, *Econometric Theory* 12, 657-681, and they provide further analysis in 1999, The relative efficiency of method of moments estimators, *Journal of Econometrics* 92 (1999) 149-172. However the optimal methodology for estimating models of this type is far from settled.

In any case, EMM is a special case of GMM, the generalized method of moments. A generalized moment is any quantity that can be averaged over a data set, such as $(3/x)\ln x$. GMM fits a model by matching the modeled and empirical generalized moments for some selection of generalized moments. EMM is a particular choice of generalized moments that has some favorable statistical properties when used to fit stochastic models.

EMM for a particular data set starts by finding the best time series model, called the auxiliary model, that can be fit to that data. If the auxiliary model is fit by maximum likelihood, then the scores of that model (i.e., the first partial derivatives of the log-likelihood function with respect to each model parameter) will be zero at the MLE estimates. These score functions can be viewed as generalized moments, which are all zero when averaged over the data. The fitted value of the scores of the auxiliary model might be hard to calculate for the stochastic model, but they can be approximated numerically by simulating a large sample from the stochastic model, and computing the scores of the auxiliary model for that sample. The parameters of the stochastic model can then be adjusted to match these moments, i.e., until all the scores approximate zero for the generated data.

The result of this technique is a parameterized stochastic model whose simulated values have all the same dynamics as the data, as far as anyone can tell by fitting time-series models to both. With this fitting done, the modeled factors then can be calibrated to current economic condi-

tions to provide a basis for simulating future possible outcomes.

Andersen and Lund (AL) did an empirical study of short-term rate dynamics by EMM-fitting their above model to four decades of US Treasury notes, incorporating data from the 1950's through the 1990's. Their results provide empirical background to evaluate other models as well.

AL estimate k as about 0.55, which supports the power of $\frac{1}{2}$ in the CIR model. In fact the AL model with this parameter is close to the CIR model at any instant of time, but the CIR parameters are subject to change over time. Other models with $k=0$ or $k=1$ appear to be discredited for US data by this result.

The period 1979-81 had high rates and high volatility, and studies that emphasize this period have found the power of $\frac{1}{2}$ on r too low. There has been some debate about whether or not to exclude this period in fitting models. These results happened, so they can happen, but it was an unusual confluence of conditions not likely to be repeated. By taking a longer period which incorporates this interval AL do not exclude it but reduce its influence.

All parameters in the AL model were statistically significant. This implies that dependence of the volatility on r is not enough to capture the changes in volatility of interest rates. There have been periods of high volatility with low interest rates, for example. Thus the one and two-factor models without stochastic volatility appear to be insufficient to capture US interest rate dynamics.

3. Generating Yield Curves

The modeled dynamics of the short-term rate can produce implied yield curves. This is done by modeling the prices of zero-coupon bonds with different maturities, from which the implied interest rates can be backed out. $P(T)$, the current price of a bond paying \$1 at maturity T , can be calculated as the risk adjusted discounted expected value of \$1 using the continuously evolving interest rate r from the short-term model. Here "expected value" indicates that the discounted mean is calculated over all possible paths for r . This can be expressed as:

$$P(T) = E^*[\exp(-\int_0^T r_t dt)],$$

where r_t is the interest rate at time t , the integral is over the time period 0 to T , and E^* is the risk-

adjusted expected value of the discounted value over all paths r can take.

If E were not risk adjusted, the expectation that gives $P(T)$ could be approximated by simulating many instances of the r process to time T over small increments and then discounting back over each increment. The risk-adjusted expected value is obtained instead by using a risk-adjusted process to simulate the r 's. This process is like the original process except that it tends to generate higher r 's over time. These higher rates usually produce an upward-sloping yield curve.

What is the risk adjusted process for r that with this procedure will generate the yield curves? If you write the price at time t for a bond maturing at time T as a Brownian motion with drift u and volatility v , i.e.,

$$dP(t,T) = u(t,T)dt + v(t,T)dz$$

then it can be shown (Vasicek 1977) that the drift u can be expressed as a function of the risk-free rate r^f , the volatility v and a quantity λ called the market price of risk, by:

$$u(t,T) = r^f P(t,T) + \lambda(t,r)v(t,T)$$

Thus the value of the bond grows by the risk-free rate plus the product of the bond's volatility with the market price of risk, plus the stochastic term $v(t,T)dz$. The market price of risk $\lambda(t,r)$ does not depend on the maturity date T , but it could depend on the interest rate r and the current time t .

The market price of risk in the bond price process is the link that specifies the risk-adjustment to the interest rate process that will generate the bond prices as the discounted expected value. As for the bond price process, only the drift of the interest rate process needs to be risk-adjusted, and the adjustment is to add the market price of risk times a function of the volatility of the interest rate process. For instance, AL suggest using the following adjusted process to simulate the interest rates in the bond price calculation:

$$dr = a(b - r + \lambda_1 rs)dt + sr^k dz_1 \quad k > 0$$

$$d \ln s^2 = c(p - \ln s^2)dt + v dz_2$$

$$db = j(q - b + \lambda_3 b)dt + wb^{1/2} dz_3$$

This adds terms to the drift of the first and third equations but not the second, as AL feel there

is little price effect of stochastic volatility. The risk-price factors λ_1 and λ_3 can be calibrated to the current yield curve along with r , s , and b . These factors do not depend on T , so are held constant throughout any simulated yield curve calculation, but they can change stochastically when a new yield curve is calculated from a new time 0.

In the AL model you have to actually simulate the dynamics of the risk-adjusted process to get the yield curves. However, in the case of the CIR model, a closed form solution exists which simplifies the calculation. The yield rate for a zero coupon bond of maturity T is given by:

$$Y(T) = A(T) + rB(T) \text{ where:}$$

$$A(T) = -2(ab/s^2T)\ln C(T) - 2aby/s^2$$

$$B(T) = [1 - C(T)]/yT$$

$$C(T) = (1 + xye^{T/x - xy})^{-1}$$

$$x = [(a - \lambda)^2 + 2s^2]^{-1/2}$$

$$y = (a - \lambda + 1/x)/2.$$

Note that the only occurrence of r is in the Y equation, so Y is a linear function of r – but not of course of T . The linearity will come into play when we look at the distribution of Y across the generated scenarios. Since all the yield rates for different maturities are linear functions of r , they will also be linear functions of each other.

4. Historical Distributions of Yield Curves

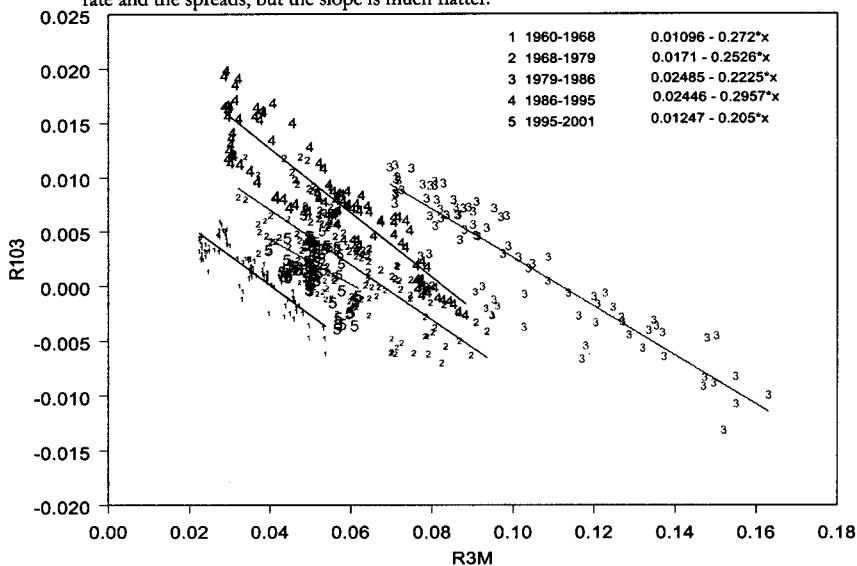
To develop tests of distributions of yield curves, it is necessary to find some properties of these distributions which remain fairly constant over time. As it is difficult to describe properties of the distribution of the entire curve, the focus will be on the distribution of yield spreads, i.e., the differences between yields.

For a property to test the models against, however, the historical distribution of a given yield spread is not necessarily all that germane. When short-term rates are high, the yield curve tends to get compressed or even inverted, so spreads get low or even negative. This is related to the mean reversion of the short-term rate. Over time it tends to move back towards its long-term

average, though with a large random deviation. Thus when it is high, a downward movement is anticipated, which produces lower long-term rates and thus negative yield spreads. If the period being projected by the model is not likely to have such high short-term rates, the yield spreads will be higher in the model than in the history.

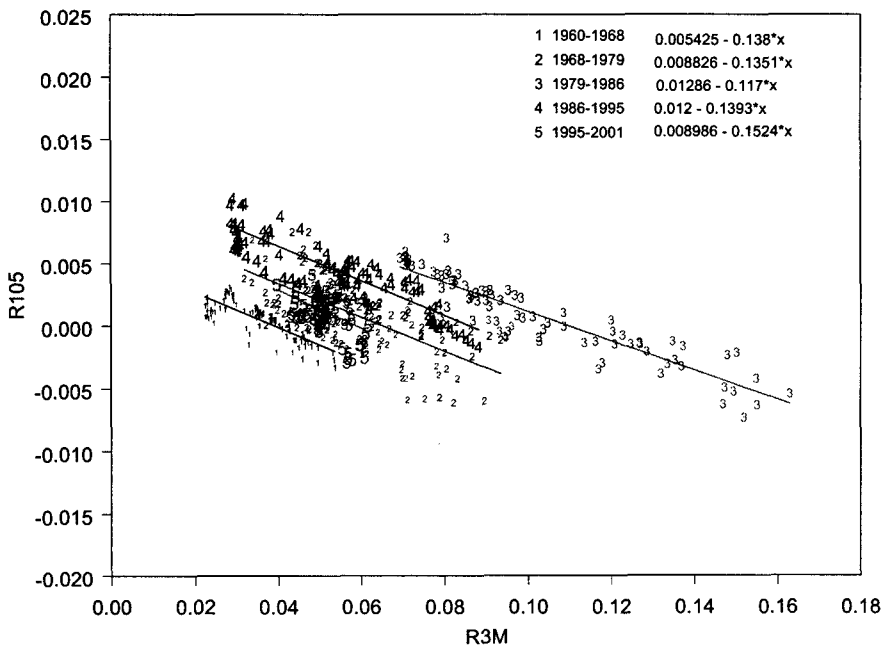
An alternative is looking at the conditional distribution of the yield spreads given the short-term rate. Over time, these conditional yield-spread distributions are more consistent than the unconditional distributions of yield spreads. The conditional distributions themselves do change in certain ways over time, however, but there are some consistencies remaining.

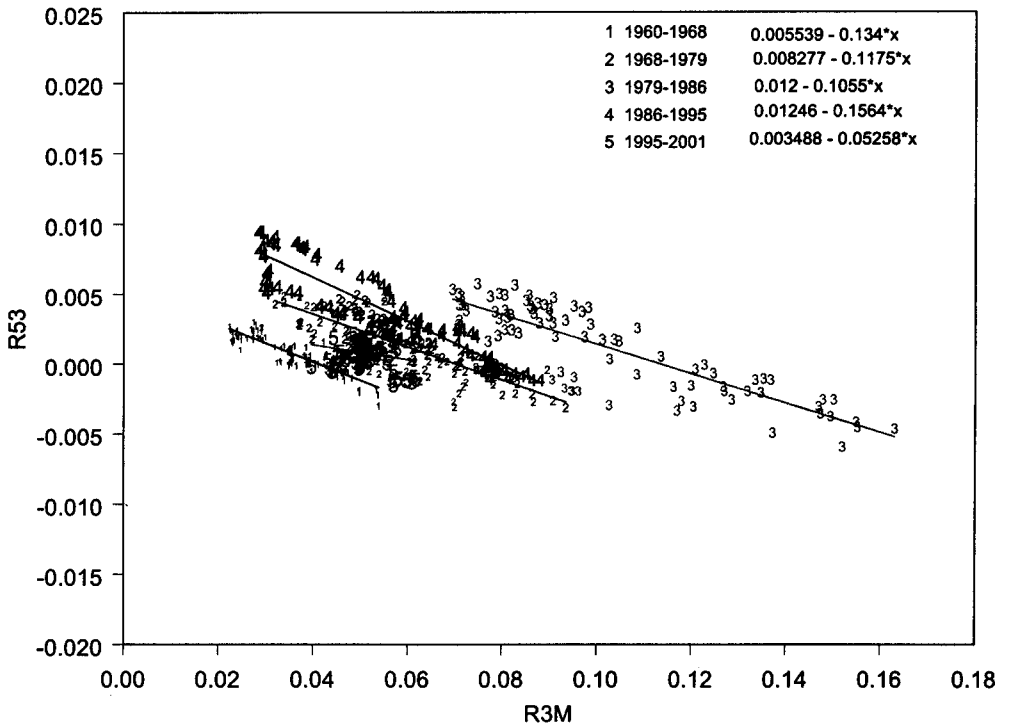
The graph below shows the US treasury three-year to ten-year yield spread as a function of the three-month rate for a 40+ year period. This period is divided up into five sub-periods, which were selected to maintain somewhat consistent relationships between the spread and the short-term rate. From the 60's to the early 80's, the short-term rates increased (sub-periods 1 – 3), then came back down after that (4 and 5). Each sub-period shows a negative slope for the spread as a function of the short-term rate, with the slopes in the range of -0.2 to -0.3 . For the entire forty year period, there still seems to be a negative relationship between the short-term rate and the spreads, but the slope is much flatter.



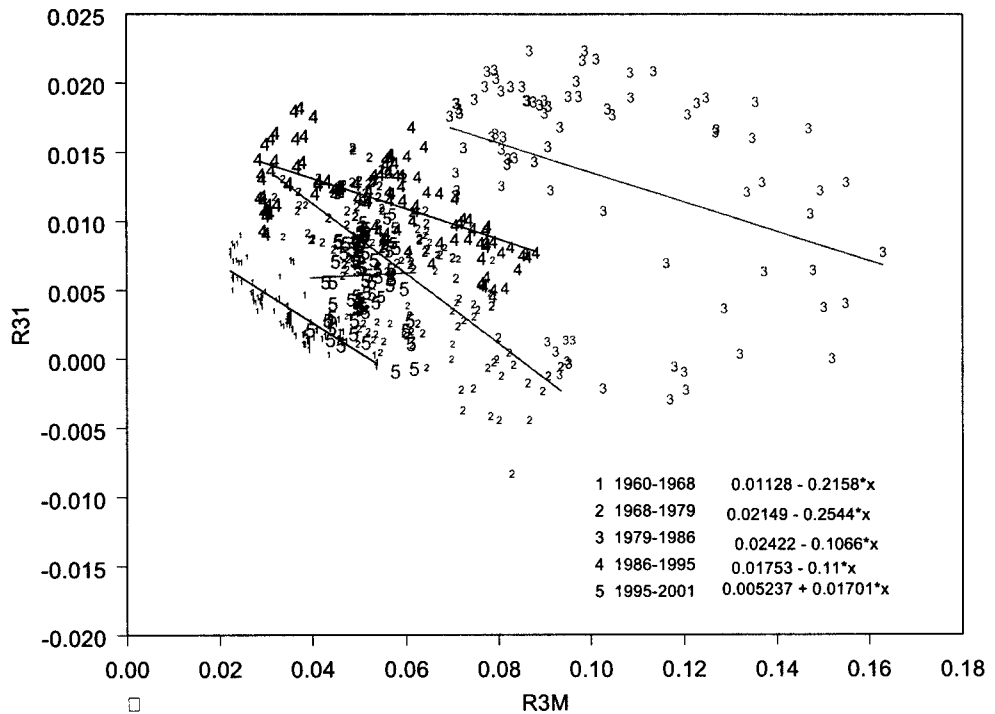
This behavior suggests that it would not be appropriate to use the conditional distribution from the entire period as a test of a scenario generator, especially if it is generating scenarios for a horizon of a few years. Over a several-year period the steeper slopes as in the historical sub-periods would be more likely to prevail. For a model projecting a few years into the future, the yield spreads would be expected to vary across scenarios, with generally lower spreads expected in those scenarios with higher short-term rates. From the historical record, it would be reasonable to expect a basically linear relationship, with a fair amount of spread around a slope in the range of -0.2 to -0.3 . This could be tested by graphing the scenarios generated by the model to see if they were generally consistent with this pattern.

The graph below shows the same thing for the five-year to ten-year spreads as a function of the three-month rate. The main difference is that the relationship of the spread to the short-term rate is less dramatic, with sub-period slopes about half what they are for the 10 – 3 spread.

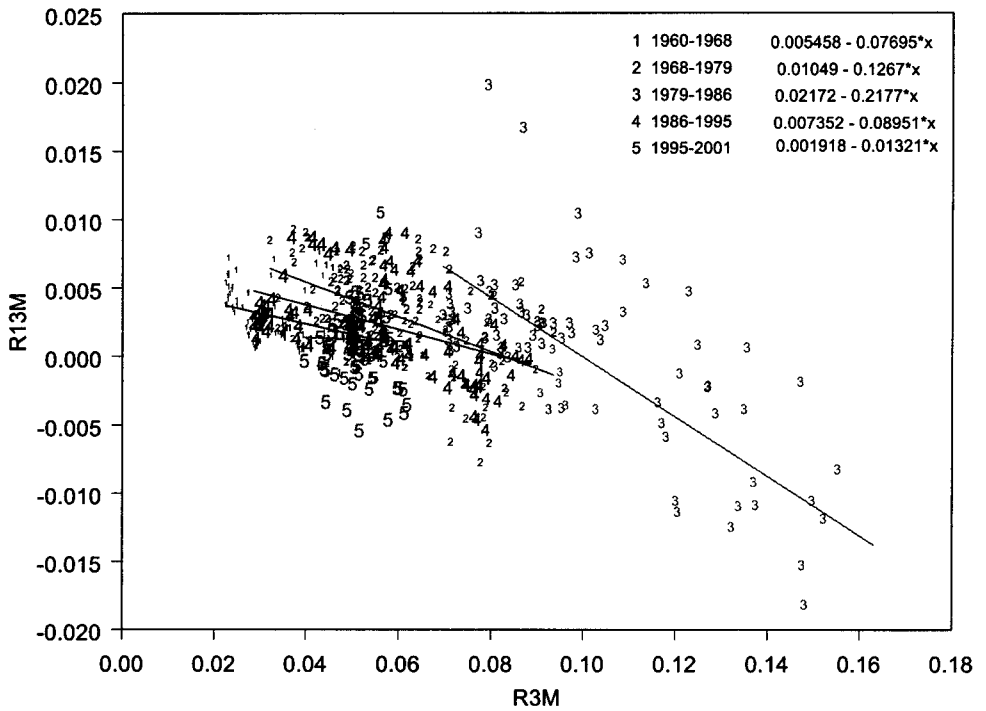




The three-year to five-year spreads show similar slopes to the 10 – 5 case, except for the latest period, which has a much flatter slope. The short-term rates in the last period have stayed in a fairly narrow range, however, making it harder to estimate the slope. In any case, relying more on the latest observations, it would seem that models producing a somewhat flatter slope in the near future should be reasonable.



The one-year to three-year spreads above show something different. Here the trend was below -0.2 in the 60's and 70's, around -0.11 in the 80's to mid-90's, and actually insignificant in the last period. Thus a flat relationship might be most appropriate in a short-horizon model.



□

The three-month to one-year spread shows even more of a break from the pattern of the longer spreads. Here the slope appears to be steeper when the short-term rates are higher, and the spreads can easily be negative. The slope is less in sub-period 5 than 1, and less in 4 than 2, suggesting that for a given short-term rate the slopes are less than they used to be. Thus a significant negative trend would not be expected for the near future, although a fair amount of randomness would still be anticipated.

5. Testing Models Against Historical Distributions

Models can be tested against the historical patterns by comparing the conditional distribution of yield spreads across the scenarios to the historical patterns to see if the patterns that have been produced historically are produced by the models. Initially two models will be compared. Both are produced by Guy Carpenter's proprietary scenario generator Global Asset Realization Processor, or GARP. They are both based on the AL specification for the short-term rate generator, but they differ in the treatment of the market price of risk. The CIR model will be included also.

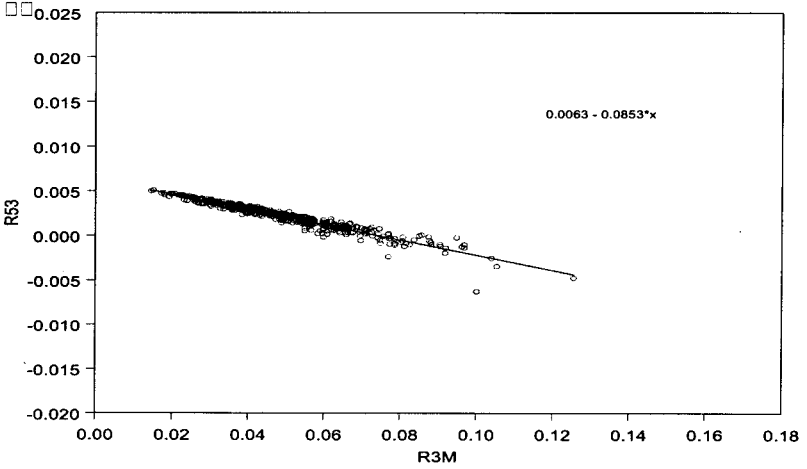
The market price of risk has to be a deterministic function across all maturities to guarantee arbitrage-free yield curves at a given time. But it can change stochastically when generating scenarios for the yield curves at another time period. Allowing the market price of risk to change stochastically produces somewhat more variability among the yield curves generated. In one model, the constant lambda model, the two AL market price of risk parameters are held constant across all simulations. In the variable lambda model, on the other hand, stochastic changes are generated from one period to the next. How best to do that is a subject of ongoing research. The variable lambda model tested here is one of many possible models of this type and has not been optimized for this test. It probably introduces a bit too much variability into the market-price of risk.

The market price of risk parameters, as well as the current values of the three factors r , b , and s are calibrated to the current yield curve to get starting values for the simulations. For this example, a yield curve from May 2001 was used for calibration. The parameters are selected that generate a current yield curve that most closely matches the selected target curve. Then yield curves are simulated at various projected periods. For periods in the near future, the curves would not be expected to be too much different from the current curves. But going out a few years produces a wider variety of yield curve scenarios. In this case the sets of curves generated for year end 2004 are used in the distributional tests. This seems like a long enough projection period to expect to see the kind of variability that exists in the sub-periods historically.

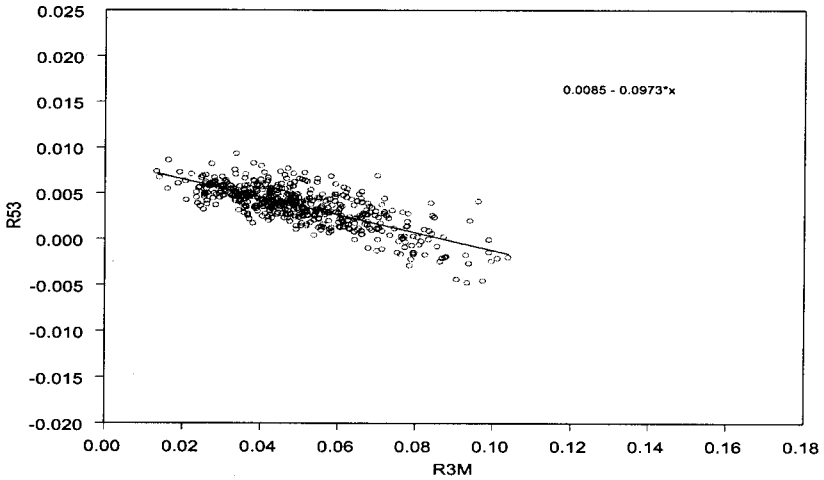
Models can be tested for the conditional distributions of all of the yield spreads. First examined is the three-year to five-year spread. Recall that the slope for this was about -0.05 in the latest sub-period, but ranged from -0.11 to -0.16 in earlier segments. The graphs below show the relationship for the simulated spreads under the two models. The constant lambda model shows a

slope of about -0.09 , vs. -0.1 for the variable lambda, which are both reasonable. There is a difference apparent in the spread around the trend line, with the constant lambda model showing little spread, and the variable lambda showing a good deal more, which is more compatible with the historical data.

GARP output 05/01 year 4 constant lambda



GARP output 05/01 year 4 variable lambda



For the CIR model it was shown above that any yield spread is a linear function of the three-month rate. Although this model does have a fair amount of flexibility in determining the slope of that relationship, there will be no variability possible around the trend line. Graphically this would look narrow like the constant lambda case, only more so. This suggests that the CIR model will necessarily produce a restricted set of yield curve scenarios, and these will not have all the variability present in historical yield curves. Thus yield curve scenarios will not be present in proportion to their probability of occurring, contrary to the criteria established above for DFA asset generators.

The table below summarizes the historical and modeled slopes and the residual standard errors from the trend lines for the sub-periods and models considered.

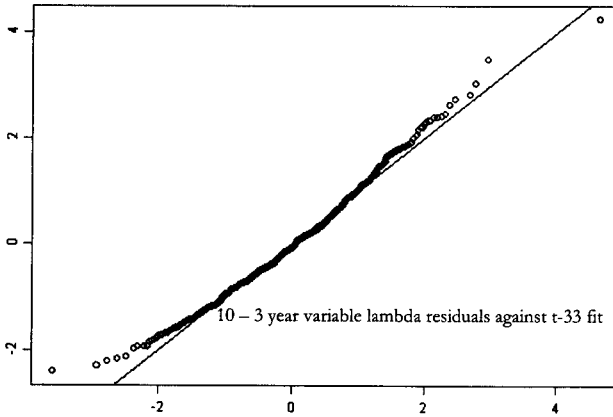
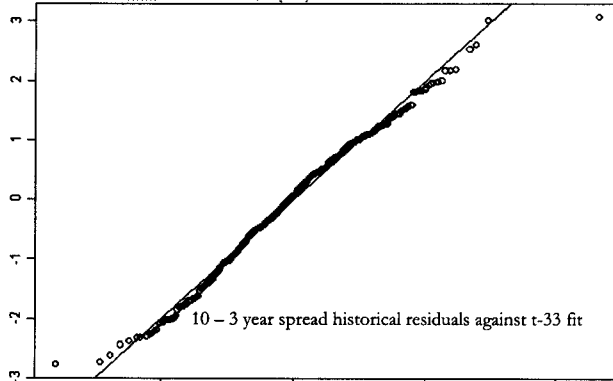
	R10 3	R10 5	R5 3	R3 1	R1 3MO
Period 1	(0.2720)	(0.1380)	(0.1340)	(0.2158)	(0.0769)
Period 2	(0.2526)	(0.1351)	(0.1175)	(0.2544)	(0.1267)
Period 3	(0.2225)	(0.1170)	(0.1055)	(0.1066)	(0.2177)
Period 4	(0.2957)	(0.1393)	(0.1564)	(0.1100)	(0.0895)
Period 5	(0.2050)	(0.1524)	(0.0526)	0.0170*	(0.0132)*
Constant λ	(0.2489)	(0.1635)	(0.0853)	(0.0721)	0.0299
Variable λ	(0.2960)	(0.1987)	(0.0973)	(0.0615)	0.0475
Period 1 se	0.0013	0.0009	0.0006	0.0013	0.0019
Period 2 se	0.0031	0.0022	0.0013	0.0037	0.0030
Period 3 se	0.0026	0.0013	0.0017	0.0070	0.0051
Period 4 se	0.0022	0.0012	0.0012	0.0024	0.0029
Period 5 se	0.0020	0.0013	0.0009	0.0028	0.0028
Constant λ se	0.0008	0.0005	0.0004	0.0013	0.0023
Variable λ se	0.0042	0.0028	0.0015	0.0021	0.0028

* Not significantly different from zero

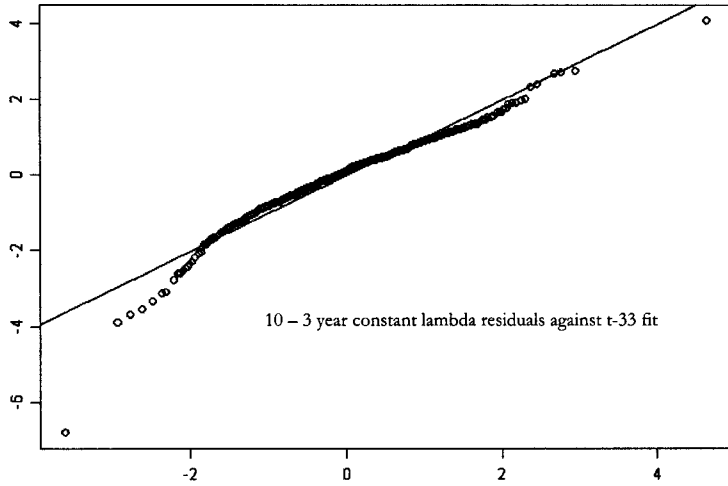
These results indicate that the constant lambda model tends to produce too little variability around the trend, whereas this formulation of the variable lambda model produces perhaps too much in the longer spreads. This suggests that allowing somewhat less variability in the stochastic processes that generates the market prices of risk could lead to still more realistic models.

6. Testing Residual Distributions

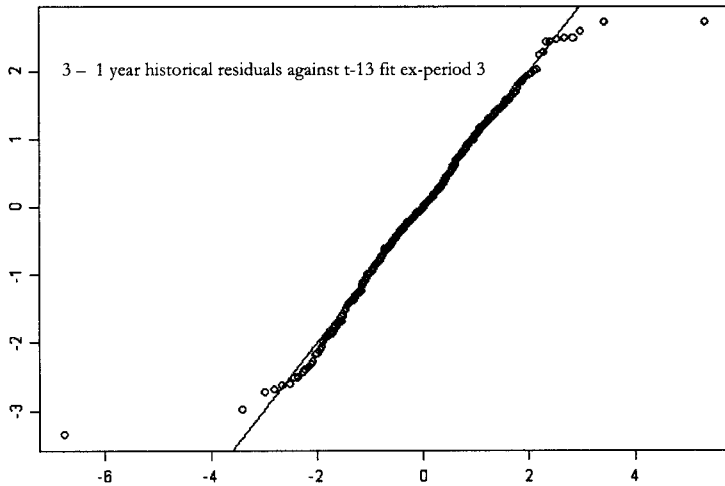
The conditional distributions of the generated yield spreads given the short-term rate have been tested against the slopes and standard errors of historical data. What about the actual distributions of the residuals around the trend lines? Are these the same historically and for the generated scenarios? This was tested by fitting t-distributions to the residuals from the model and the combined set of residuals from the historical periods. The graphs below show QQ-plots, which graph the percentiles of the residuals against the same percentiles of the fitted t 's.

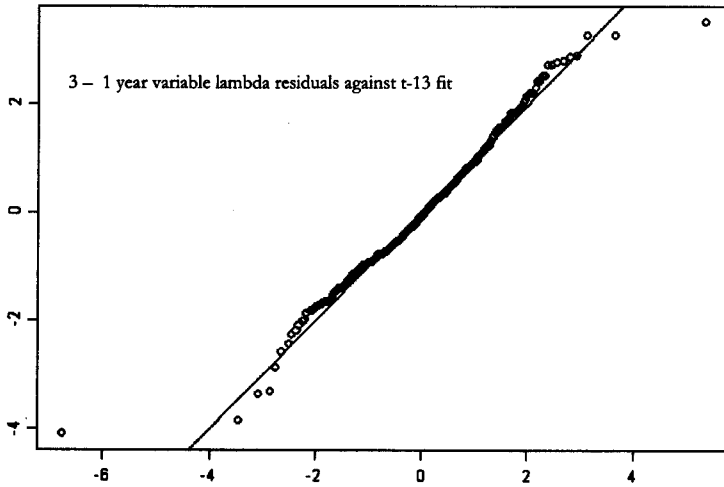


The 10 – 3 constant and variable lambda residuals look a lot like the data except in the left tail, where the constant lambda diverges. The t-distribution with 33 degrees of freedom was fit here.

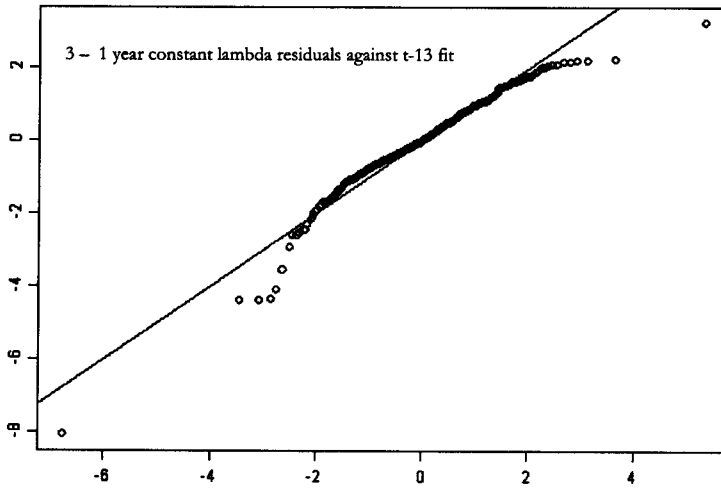


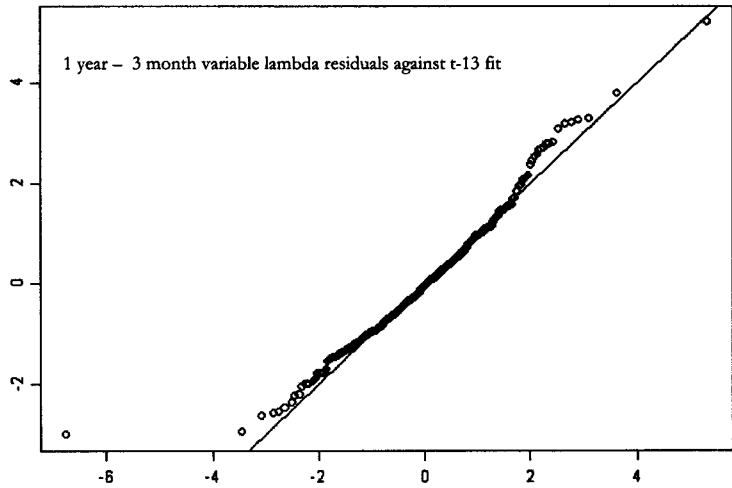
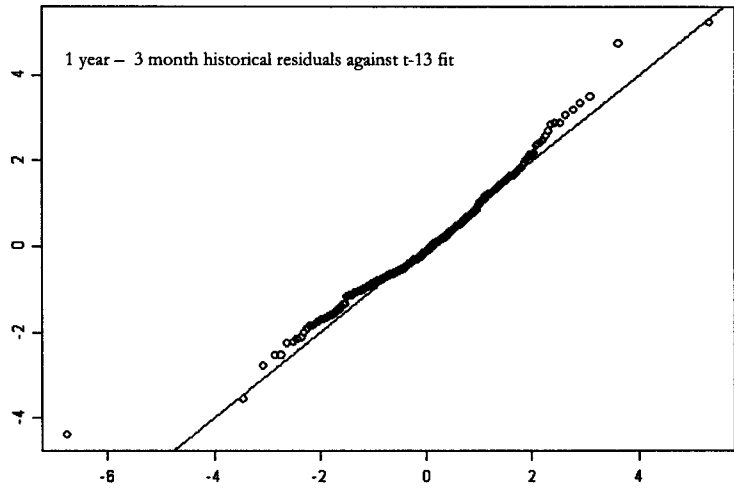
The 3 – 1 year residuals were done excluding period 3, which was unusual.

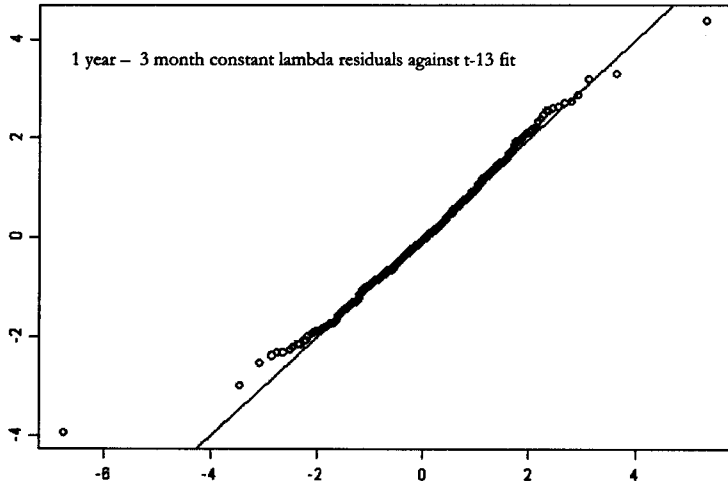




For the 3 - 1 spreads the variable lambda model provided residuals distributed similarly to those from the data, when compared to the t with 13 degrees of freedom.







Again for the 1 year - 3 month spread residuals, the data and the variable lambda model compare similarly to the t-13 fit, where the constant lambda is a little different.

Conclusion

Many models of interest rates have been proposed. For one survey, see R. Rebonato (1997) "Interest Rate Option Models," John Wiley NY. Models of the dynamics of the short-term rate apparently need to incorporate mean reversion, stochastic changes in mean reversion over time, mean sensitive volatility, proportional approximately to the square-root of the mean, and stochastic volatility as well.

Testing the conditional distribution of various yield spreads given the short-term rate appears to be a reasonable way to see if a model is generating a realistic distribution of yield curves. The unconditional distribution of generated yield spreads would not necessarily be comparable to the historical distribution, because different spreads are associated with different short-term rates, and the simulation might not be generating a distribution of short-term rates that matches the historical record, due to the particular economic conditions that prevail at the time of the simulation. The slopes of the conditional fitted lines are fairly consistent over different historical periods.

As with most tests of distributional issues, this one is not a formulaic system that gives a strict "yes/no" answer to a model's output. But it does provide a realm of reasonable results so you can give an opinion of the "probably ok/ probably not" type. For example, having no variability around the conditional trend line would seem to be too limiting. Slopes that are much steeper than historical would also seem disindicating, as would distributions of residuals around the slopes that differ substantially from the t-distributions fit. Even though these tests are not strict, better results could be sought than those of any of the models tested.

An application issue is how much variability you should have for projection periods of different lengths. When projecting out four or five years, a conditional distribution similar to those of the historical sub-periods might be appropriate. However there is some chance of entering a new realm – i.e., changing sub-periods – over that much time. In all the sub-periods graphed, changing to an adjacent sub-period would tend to flatten the conditional trend.

Appendix I – Scalar Measures of Response

A number of risk measures have been devised to look at the effect on an investment holding or portfolio of a small change in some quantity. For example, Macaulay duration measures the change in the value of a portfolio due to a change in the annualized average yield to maturity. It can be expressed as the weighted average of the times to each cash flow of the portfolio, where the weights are the cash flow amounts discounted at the average yield. Thus duration is expressed in units of time. (Duration measures value per interest rate, but as interest rate is value per time, duration is time.) One way to produce a given change in the average yield to maturity is to shift the entire yield curve by the same amount, so duration is often described as the sensitivity of the portfolio to a parallel shift in the curve.

Macaulay convexity is the weighted average of the squares of the times to the cash flows, using the same weights as for duration. It can be shown to be the square of duration less the derivative of duration with respect to the instantaneous average yield.

The analysis of derivative instruments has produced several similar measures, denoted by Greek letters, and so called “the Greeks.” These measure the change in the value of a position brought on by the change in something else that affects value. For instance, the change in the value of an option due to the change in the value of the underlying security is called delta.

For bond portfolios, each bond could be thought of as a holding of a combination of future positions in the short-term rate, which could thus be considered to be the underlying security. With the short-term rate as the underlying security, the delta risk is the change in the value of the portfolio with respect to a small change in the short-term rate. This is different than duration, as even though all the rates will change in response to a change in the short-term rate, they will not necessarily change by the same amount. This is clear in the CIR model where a change in r makes all the rates change, but each by its own $B(T)$. If the underlying security is taken to be the average yield to maturity, then delta is duration.

Gamma risk is the change in delta due to a small change in the value of the underlying security. With the short-term rate as the underlying security, in CIR gamma is zero, but for a typical asset or liability portfolio it will not be. Gamma is somewhat analogous to convexity, but as defined

here focuses on the actual short-term rate, not the average yield.

Vega measures the change in value due to a change in the volatility of the underlying instrument. The volatility of the short-term rate Brownian motion is an element in bond pricing, so vega risk is present in bond portfolios. CMO's probably have a fair degree of this risk as well, as greater interest rate volatility can increase the probability of pre-payment.

Theta is just the sensitivity of the position to a small change in the valuation date.

Rho for any portfolio measures its change in value due to a small change in the interest rate. In most asset pricing models the yield curve is assumed to be constant, so rho could be considered to be the effect of a shift in the average yield, i.e., duration.

Appendix 2 – The Arbitrage Debate

Most finance theory takes the impossibility of arbitrage as a given, but some actuaries use interest rate models that are not arbitrage-free. This may be just a matter of convenience, but two arguments are sometimes advanced for using such models:

1. Actual published yield curves are not always arbitrage-free
2. It is more important to get the statistical properties of the set of scenarios right than to avoid arbitrage.

One problem from having arbitrage possibilities in generated scenarios is that searching for optimal investment strategies would find the arbitrage strategy, and that will appear the best. It seems pretty unlikely, however, that a DFA model could identify truly risk-free high-profit investment strategies that insurers could work in practice. Even if the search disallowed the arbitrage strategies, their presence in the scenario set could have a distorting effect. However, a model that allows arbitrage only in unrealistic cases, like being able to borrow huge amounts at the risk-free rate, could be considered arbitrage-free in practice.

With this in mind, the two arguments can be reviewed separately. First, there may occasionally be some arbitrage possibilities in published yield curves. But this does not mean that these can be taken advantage of in practice. For one thing, the published curves look at trades that took place at slightly different times, so are not snapshots of one moment in time. Looking at a combination of positions in different deals that have happened recently could yield a hypothetical arbitrage, but that possibility could be gone before it could be realized. A related issue is that some of the deals might have to be scaled up significantly to get the arbitrage to work, and doing this could change the prices. In short, finding some historical published yield curves with hypothetical arbitrage possibilities in them is not reason enough to use a modeled set of scenarios that have specific arbitrage strategies built in.

The second argument is more interesting. This paper argues for the importance of getting the statistical issues right, focusing on the distribution of yield spreads across scenarios. This does not appear to be in any way inconsistent with no arbitrage. Using models like AL also emphasizes that the movement of interest rates across time should be statistically correct.. Thus both

the statistics of changes in rates over time and the distribution of yield spreads at each time are compatible with arbitrage-free scenarios. It would be interesting to see what other statistical issues there are that would require using scenarios with arbitrage built in.