Modeling the Evolution of Interest Rates: The Key to DFA Asset Models

by Gary G. Venter, FCAS

Modeling the Evolution of Interest Rates: The Key to DFA Asset Models

Gary G. Venter

INSTRAT

MODELING THE EVOLUTION OF INTEREST RATES: THE KEY TO DFA ASSET MODELS

Fluctuations in short term and long term interest rates can have significant impact on insurer financial results. Hence projecting the probabilities of these possible fluctuations is an important step towards a credible dynamic financial analysis. Changes in the level of interest rates as well as shifts in the shape of the yield curve both need to be modeled. Such shape-shifting is not an unconstrained random process - there are relationships among the yields of different terms - yet a good deal of flexibility is required to be able to reproduce historical curves.

Financial theory suggests that the yield curve at any point in time is a function of the probabilities of the future values of the short-term interest rate. Thus a process that produces probabilities for the evolution of the short-term rate will also have implications for the entire term structure. Simulating probabilities for future yield curves can proceed by first simulating short-term rate probabilities over an extended horizon, and then using those to simulate yield curve probabilities for a shorter horizon.

To illustrate that procedure, this paper has four sections: first models for short-term interest rate changes will be discussed, followed by a discussion of how to produce yield curves from those models. Then estimation issues are addressed, and the final topic is adding other correlated economic variables.

1 MODELING SHORT-TERM INTEREST RATES

Most models of short-term rates are expressed as stochastic differential equations involving Brownian motion. At the time of this writing that is not a topic on the CAS Syllabus, so a short deviation will be taken to explain the notation to be used.

1.1 Brownian Motion

Standard Brownian motion is a sequence of random variables X_t indexed by time t, where X_t is normally distributed with $E(X_t)=0$ and $Var(X_t)=t$. Thus the variance grows over time. A *Brownian motion with drift μ and variance σ^2* has X_t normally distributed with mean μt and variance $\sigma^2 t$.

There are other technical requirements: a Brownian motion must be a continuous process with stationary independent increments. That means that the increments $X_s - X_t$ are independent for different choices of s and t, and are stationary in the sense that the distribution of $X_s - X_t$ is the same for any s and t with a common value of s-t. One reason for the popularity of Brownian motion is that the converse is also true: a continuous stochastic process with stationary independent increments must be a Brownian motion. (See L. Brieman *Probability* Addison-Wesley 1968 ch. 12.) This can be related to the Central Limit Theorem. The sum of a lot of independent increments would tend to normality.

1.2 Stochastic Differential Equations

Methods for solving stochastic differential equations will not be addressed in this paper, but the notation will be used as a recipe for simulation. For example, let z be a standard Brownian motion, and consider the following equation for the short-term interest rate r:

$$dr = \mu dt + \sigma dz \tag{1}$$

This can be interpreted as a way to simulate changes in r over short intervals. Say the short interval has length Δt. Then simulate the change in r (i.e., Δr) as a draw from a normal distribution with mean $\mu\Delta t$ and variance $\sigma^2\Delta t$.

1.3 Models of Short-Term Interest

It is fairly common in financial mathematics to express interest rates as continuously compounding - what actuaries call the force of interest. Thus the usual interest rate i becomes the force of interest r where $(1+i)^t = e^{rt}$. The short-term rate is the instantaneous continuously compounding rate, which can be thought of as the limit of shorter and shorter terms. Sometimes this is estimated as the one-month rate, or even the three-month rate, or as a projection backwards from a few of the rates for shorter terms.

One possible model for r is expressed in (1) above: r is a Brownian motion. Another possibility would be to let y = ln r be a Brownian motion. This is called *geometric Brownian motion*, and excludes any possibility that r could become negative. *Single factor models* are those that can be expressed using only one Brownian motion process. A number of these models are of the form:

$$dr = (a+br)dt + \sigma r^k dz \qquad (2)$$

where $0 \le k \le 1$. Often k is taken as ½ or 1, which would make the variance of the change in rates proportional to r or r^2. Typically a is non-negative and b is non-positive. Note that this model has four parameters to estimate, even though it is a single-factor model.

In this model it is not possible for r to become negative if a is positive, because if r gets to zero, all the terms become zero except for the positive drift, and r becomes positive in the next instant. When b is negative the process is called *mean reverting*. If $|br|$ is above a, the drift will be downward, and if below a, the drift will become upward. Thus the drift is always back towards a. A Brownian motion process with no drift that is not adjusted to be mean reverting will eventually become quite wild. The variance $\sigma^2 t$ will grow with time, so the probability of finding the process to be within a given distance of zero will diminish to the vanishing point. However non-mean-reverting processes are sometimes used in short-term forecasts of interest rates.

A mean-reverting process may display negative auto-correlation at some intervals. That is, if it is going up at some point, it is likely to be going down at some future point. Interest rates seem to display this behavior. For instance, one study found that the logs of the growth rates of short-term rates are positively correlated from one month to the next, but negatively correlated to those of six and seven months earlier (D. Becker, *Statistical Tests of the Lognormal Distribution as a Basis for Interest Rate Changes,* Transactions, Society of Actuaries, vol. XLIII).

Multi-factor models specify interest rate evolution as process involving the interaction of multiple random effects at each stage. For instance, a three factor model may specify that the interest rate evolves according to a random process in which the expected rate of change and its variance also evolve randomly. Six or seven parameters may be needed to describe these three random effects.

One apparently successful three-factor model is given by Anderson and Lund (Working Paper No. 214, Northwestern University Department of Finance):

$$dr_t = a(u_t - r_t)dt + s_t r_t^k dz_1 \qquad k>0 \tag{3}$$

$$d\ln s_t = b(p - \ln s_t)dt + v\,dz_2 \tag{4}$$

$$du_t = c(q - u_t)dt + w u_t^{1/2} dz_3 \tag{5}$$

Here there are three standard Brownian motion processes, z_1, z_2, and z_3. The rate r moves subject to different processes at different times. It always follows a mean-reverting process, with the mean at time t denoted by u_t. But that mean itself changes over time, following a mean-reverting process defined by c, q, and w. The standard deviation of r_t is $r_t^k s_t$, where s_t also varies over time via a mean reverting geometric Brownian motion process. In total there are eight parameters: a, b, c, k, p, q, v, and w.

A two-factor model by M. Tenney (The Double Mean Reverting Process™, Society of Actuaries Technical Report, 1996) takes a somewhat different approach toward keeping the interest rate positive. Let y = ln r be the ln of the interest rate rather than the interest rate itself. Tenney's model then can be expressed as:

$$dy_t = a(u_t - y_t)dt + v\,dz_1 \qquad (6)$$

$$du_t = c(q - u_t)dt + w\,dz_2 \qquad (7)$$

where z_1 and z_2 are correlated standard Brownian motion processes with correlation ρ. Thus there are six parameters: a, c, q, v, w, and ρ.

Neither of the above models is particularly easy to estimate from data, but once estimated, simulation is quite straightforward for either of them. Multi-factor models are used despite the estimation difficulties because of some of the weaknesses of single factor models. These include difficulty in capturing the movements of historical interest rates, and difficulty in matching the historical yield curves. The historical yield rates, for instance, display different rates of variation in different periods, and these do not necessarily correlate directly to the interest rate level. This could be evidence of stochastic movement of the variance, as in equation (4) above. It is also consistent with an infinite variance process, which would generate unstable measurements of observed variance in different periods. Modeling interest rates as such a process will not be addressed here, however. Historical yield curves occasionally have inversions, in which the short-term rates are higher than those for longer terms. This usually is not allowed by the single-factor models.

2 Yield Curves Implied by Short-Term Rate Models

One reason long-term rates are usually higher than short-term rates is that long-term investors take the risk that intervening events will render the accumulated earnings worth less. In the very long run, though, all things might average out, and so very long-term rates are not necessarily higher than long-term rates, and may even be lower. It is not unusual, for instance, for 20 year rates to be a little higher than 30 year rates. Some infinite term bonds have been issued in the UK at fairly low rates.

The standard method for producing yield curves from a stochastic generator of short-term rates is to change the parameters of the generator to make it generate higher short-term rates for time periods further into the future, and then to take the expected value of the future adjusted rates as the estimate of what the short-term rate will be at that future period. The short-term rates so estimated for each future period then can be put together to make the long-term rates. This general concept will be spelled out more precisely below.

Adjusting the future rates generated is equivalent to keeping the rates but adjusting the probabilities in a manner that increases the expected value of the future short-term rates. Adjusted probability methods seem intuitively reasonable, but they are also justified by arbitrage theory. Thus a short detour into arbitrage theory may be useful.

2.1 Arbitrage of Interest Rates

A common financial definition of an arbitrage opportunity is a possibility to make a net investment of zero, and end up with no probability of a loss and a positive probability of a gain. Arbitrage theory says that there are no arbitrage opportunities available. This is not universally accepted by casualty actuaries. Two types of comments are often heard:

1. Investment houses make arbitrage profits all the time. They borrow at the 3-month rate and lend out at the higher 6-month rate.
2. Investment houses make arbitrage profits all the time. They have sophisticated trading models that look for these opportunities continuously, and put up big bucks whenever they arise, which is often.

The first is not really an arbitrage profit, at least by the above definition. It may be a pretty good bet, but now and then the 3-month rate will jump while the money is still out at the now lower 6-month rate, and the investors will have a loss when they have to borrow at a higher rate than they are getting.

The second may indeed be true. But if it is happening, the big boys are taking out all the arbitrage profits before anyone else ever sees them. It would be highly unusual for the end-of-day published rates to have arbitrage possibilities in them. If they appear in the 20 minute delay quotes on-line, they are probably gone by the time they appear.

An example of an interest rate arbitrage is adapted from P. Boyle (*Options and the Management of Financial Risk*, Society of Actuaries, 1992). Suppose the yield curve is flat: all rates are 8%. In the next instant they will be flat again, with 50% probability of staying at 8%, but with 25% probability each of moving to 7% or 9%. Borrow 1000 due in 10 years, and use the $1000/1.08^{10}$ to make loans with single payments of $500/1.08^5$ due in 5 years and $500 \cdot 1.08^5$ due in 15 years. Each of those loans costs $500/1.08^{10}$ to make today, so this produces a net position of zero. But in the next instant if the interest rates go to 7% the net position is worth 0.550 (i.e., $500/[1.08^5 1.07^5]+500 \cdot 1.08^5/1.07^{15} - 1000/1.07^{10}$). Interestingly enough, if they go to 9% it is worth 0.449. At 8% it stays at zero. This is an arbitrage opportunity by the above definition, and so it is ruled out by arbitrage theory.

Although this is a highly artificial example, it shows that certain combinations of yield curves and interest rate movements are not possible under arbitrage theory. Boyle has more seemingly realistic examples that are likewise disallowed. This raises the issue of what yield curve / interest rate movement combinations are possible without generating arbitrage opportunities.

2.2 Pricing Consistent with Arbitrage Theory

It turns out that to rule out arbitrage possibilities, securities must be priced as the expected value of their returns under some probability distribution. The probabilities do not have to be the actual probabilities of those returns. In fact, if they were, there would be no reward for risk, which is unrealistic. Thus risk-adjusted probabilities must be used. Sometimes the risk-adjusted probabilities are called *risk-neutral* probabilities. That is because when using them you act as if risk were not important - i.e., you just take expected values. But this does not mean risk is ignored: expected value pricing based on risk-neutral probabilities is a method for building risk premium into prices.

One constraint on the risk-adjusted probabilities is that they are equivalent to the actual probabilities in the sense that they give zero probability to the same set of events. This is violated in the Boyle example above, where there is positive probability of a change in interest rates, but prices are based on expected values under the assumption of no possibility of changing rates. In this situation the adjusted probabilities give no chance to the events that can lead to the actual positive profit probabilities.

Although it is complex to prove that no-arbitrage and adjusted-probability expected-value pricing are equivalent, the following heuristic argument may help make it plausible. The key to avoiding arbitrage is to ensure that prices are additive. That is, the sum of the prices of a combination of securities that always produce the same outcomes as another given security should equal the price of that security. If not, buying the cheaper and selling the dearer set will give a profit in every case. But if prices of all securities are additive, there must be some set of event probabilities that gives those prices as expected values. The key to seeing that is to define fundamental securities that relate to the specific possible events. For instance with interest rates, those securities might pay 1 if the interest rates exceed specific targets for each term, and 0 otherwise. Such securities could be defined for any combination of term interest rates. The prices of those securities would define a joint probability distribution for the interest rates, and all other securities could be priced as combinations of those, which would be like taking their expected values under the distribution so defined.

2.3 Arbitrage-Free Pricing under Interest Rate Generators

The price at time t of a zero-coupon bond maturing at time T is the discounted value of the payment. With constant interest this is no problem, but with stochastic interest this would require the expected average discount, with the discount taken with respect to the risk-adjusted probabilities. The price of a bond that pays 1 at maturity can be expressed as:

$$P(t,T) = E_t^*[\exp(-\int r_s ds)] \tag{8}$$

where the integral goes from t to T, and E^* is the mean using the risk-adjusted probabilities. From the price of the bond, the implied interest rate for that term can then be backed out. In practice, the integral is evaluated as a sum over the small intervals used in the generation of short-term rates.

Thus the term structure is tied to the future paths of the short-term rate. Once an interest rate generator is available, what needs to be specified is how the risk-adjusted probabilities are to be defined. What is usually done is to change the generator so that it produces higher interest rates over time. Strictly speaking this gives a higher rate at each probability, but this then produces higher probabilities for the higher rates, which is what arbitrage theory is looking for.

A typical adjustment is to add something to the drift terms. For instance, Anderson and Lund to (3) - (5) above add $\lambda s_t r_t dt$ to the r diffusion and $\kappa u_t dt$ to the u diffusion. This gives a new process for generating risk-adjusted short-term rates, as below:

$$dr_t = a(\lambda s_t r_t + u_t - r_t)dt + s_t r_t^k dz_1 \qquad \lambda,k>0 \tag{9}$$

$$d\ln s_t = b(p - \ln s_t)dt + v dz_2 \tag{10}$$

$$du_t = c(\kappa u_t + q - u_t)dt + w u_t^{1/2} dz_3 \qquad \kappa>0 \tag{11}$$

Thus both r and u increase at higher rates, on the average, in the risk-adjusted process. The rate scenarios generated by (9) - (11) are used to evaluate the expected value of the integral in (8) to give bond prices, which are essentially the discount rates for the various terms.

Tenney similarly increases the drifts, but also changes the rates of mean reversion. The adjusted process from (6) and (7) is:

$$dy_t = \phi a(\lambda + u_t - y_t)dt + v dz_1 \quad \lambda>0,\ \phi>1 \tag{12}$$

$$du_t = \varphi c(\kappa + q - u_t)dt + w dz_2 \quad \kappa>0,\ 1>\varphi>0 \tag{13}$$

Note that in the u diffusion the mean reversion is slower than in the unadjusted process. This increases the variability of u.

3 Estimation of Parameters

Since the same parameters predict both the movements of short-term rates and the term structure, fitting can be done to either or both. If the fit is going to emphasize the term structure, equation (8) can be fit via simulation. However, this could require that a simulation be carried out a each step of a parameter search, which can be quite calculation intensive. Thus closed form or otherwise more tractable forms are usually sought for the zero-coupon bond prices at each maturity. Even if the fit is going to emphasize movements of the short-term rate, further fitting to the bond prices is needed to get the risk adjustments.

The usual approach to bond price formulation is to develop a stochastic differential equation for the price of the bond. This can then be solved explicitly or numerically. Developing such equations typically uses Ito's Lemma, which is the chain rule for stochastic calculus.

3.1 Ito's Lemma

Brownian motion is continuous, but is very jumpy at small scales, so is not differentiable. However a method of integration of these processes has been developed. This allows the use of differential notation, but the usual rules of derivative calculus, such as the chain rule, do not apply. However an analogue of the chain rule has been developed, and is known as Ito's Lemma. Suppose a process x can be expressed by $dx = y\,dt + s\,dz$, where z is a standard Brownian motion. If f is a twice differentiable real-valued function, then:

$$df(x) = \tfrac{1}{2}f''(x)s^2dt + f'(x)dx \qquad (14)$$

The second term is the usual chain rule, while the first is sometimes called the *convexity* term.

For example, suppose the change in r is proportional to the current level of r:

$$dr = \mu_t r\,dt + \sigma r\,dz \qquad (15)$$

Let $y = f(r) = \ln r$. Then $f'(r)=1/r$ and $f''(r)=-1/r^2$, so:

$$dy = -\tfrac{1}{2}\sigma^2dt + \mu_t dt + \sigma dz \qquad (16)$$

Thus the two methods illustrated above for keeping r positive - namely Brownian motion proportional to a power of r and geometric Brownian motion - are closely related. When converting a lognormal mean to the normal mean you add $\tfrac{1}{2}\sigma^2$ before exponentiating, which corresponds to the extra term in (16).

3.2 Solving for Bond Price

The price of a bond is some function f of the interest rate r. Thus Ito's Lemma can be used to derive a differential equation for the bond price. The usual approach in the single factor setting is to specify a risk-adjusted diffusion for r like:

$$dr_t = u_t(r_t)dt + \sigma_t(r_t)dz \qquad (17)$$

Then the price at time t of a bond maturing at time T, expressed as $P(r_t,t,T)$, can be shown to follow:

$$rP = P_t + uP_r + \tfrac{1}{2}\sigma^2 P_{rr} \qquad (18)$$

which is a differential equation with boundary condition $P(r,T,T) = 1$.
For example, see Vetzal *A Survey of Stochastic Continuous Time Models of the Term Structure of Interest Rates,* Insurance Mathematics and Economics (14), 1994.

In the multi-factor setting a similar differential equation can be derived, with partial derivatives of the price entering from all factors. In general if Y is the vector of factors, u is the vector of risk-adjusted drifts for the factors, and σ is the vector of variances, then the bond price P satisfies:

$$rP = P_t + u^T P_Y + \tfrac{1}{2}\mathrm{tr}[\sigma\sigma^T P_{YY}] \qquad (19)$$

It turns out that the discount formula (8) above is a solution to (19), sometimes called the *Feynman-Kac* solution. Thus using simulation to solve (8) does solve (19). The advantage of going to (19) directly is that it can sometimes be solved in closed form, as in the example below, or by numerical methods that are less intensive than simulation. This is the approach taken by Tenney, for example, to estimate the parameters in the geometric Brownian motion two-factor model (6)-(7) above. However, Anderson and Lund solve for the diffusion parameters directly, which requires using (8) and a lot of computation or (19) and some derivation and numerical methods to get the risk terms.

3.3 A Simplified Model

Since the model parameters affect both the evolution over time of the short-term rate and the term structure at each point, both effects can be used to evaluate the goodness of fit. Thus both will influence the choice of parameters - i.e., parameters are needed that fulfill both roles. In the models above it is difficult to illustrate this interaction, as the generation of the term structure is complex.

Thus to illustrate these general concepts, parameter estimation will be discussed for a somewhat simplified three-factor model of the short-term interest rate - one with a closed form solution for the yield curve. This is the model of Kraus and Smith (*A Simple Multifactor Term Structure Model*, The Journal of Fixed Income, March 1993).

K&S postulate that the term structure at any point in time can be described as a function of three factors r, μ, and α. These factors evolve over time through Brownian motion, according to the equations below.

$$dr = \mu dt + \sigma dz_1 \quad (20)$$

$$d\mu = m dt + s dz_2 \quad (21)$$

$$d\alpha = dm - d(\sigma^2) = b dt + v dz_3 \quad (22)$$

So r is just a Brownian motion process, but the drift and variance both change over time. Here the time subscripts on the variables that change over time are omitted, but are implied. The drift μ is itself subject to a Brownian motion. The variance of the r diffusion, σ^2, is linked with the drift of the drift, m, as a single process $\alpha = m - \sigma^2$. This turns out to simplify the term structure formulation. However the variance itself could follow some unspecified process, like mean-reverting geometric Brownian motion.

On empirical and practical grounds it seems reasonable to set b to zero. Doing so leaves only two parameters - s and v. It is problematic that the model is not mean reverting, and interest rates are allowed to become negative. However, this may be a reasonable model to use for short-term projections. The advantage of its formulation is in the simplicity and flexibility of the resulting term structure. With some additional assumptions, the term structure turns out to be a simple polynomial form in the term T, with coefficients that are linear functions of the parameters and factors of the model. Thus long-range simulations are not needed to generate a term structure distribution for the near future.

K&S assume that the yield rates are linear functions of the factors (but not of the term). They also introduce three risk-adjustment coefficients, one for each factor, that are similar to the λ and κ risk terms in the models above. They then derive, using a no-arbitrage argument, a polynomial form for the term structure. This proceeds by setting up a differential equation for bond prices, from (19):

$$rP = P_t + (\mu_j+\lambda_r)P_r + (m+\lambda_\mu)P_\mu + \lambda_\alpha P_\alpha + \tfrac{1}{2}[\sigma^2 P_{rr} + s^2 P_{\mu\mu} + v^2 P_{\alpha\alpha}] \qquad (23)$$

where λ_Φ is the risk factors for factor Φ. Here an additive constant risk factor has been added to each drift coefficient. This equation has a closed form solution. Let $y_j(T)$ denote the yield for term T at time j. The assumptions then give:

$$y_j(T) = r_j + T(\mu_j+\lambda_r)/2 + T^2(\alpha_j+\lambda_\mu)/6 + T^3\lambda_\sigma/24 - T^4 s^2/40 - T^6 v^2/504 \qquad (24)$$

Note that the equation is a sixth degree polynomial in the term, and the higher order coefficients are negative. This implies that for long enough terms the yield rate will decrease. For the parameters estimated below, the thirty year rate is often less than the twenty-five year rates, which is often the case in the data as well. Of course this also implies that very long term rates will be negative, which is not realistic. The model clearly should not be used for very long terms.

The first three (quadratic) terms of the polynomial vary with time. An interpretation is then that there is a fixed sixth degree polynomial for the standard yield curve, and this gets shifted up and down by a quadratic over time. The graph to the left shows the polynomial defined by the estimates of the fixed terms of (17) - i.e., all the elements without subscripts, namely $T\lambda_r/2 + T^2\lambda_\mu/6 + T^3\lambda_\sigma/24 - T^4 s^2/40 - T^6 v^2/504$.

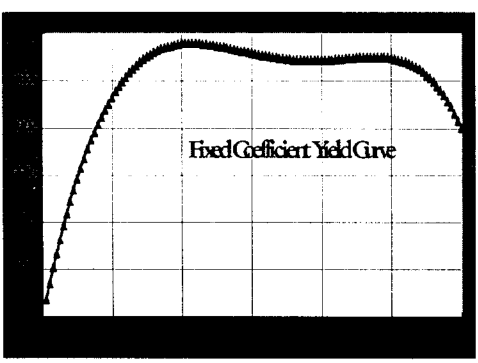

The quadratic shift gives quite a bit of flexibility to the shape of the yield curve. Reversals and other unusual shapes are readily produced. This is in contrast to single factor models, which typically allow only parallel shifts. The quadratic terms for a selected quarter and the resulting yield curve are shown below.

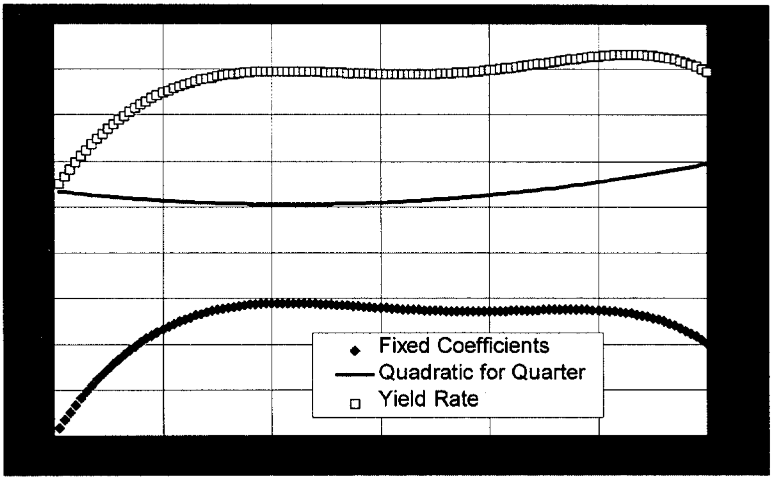

3.4 Results and Discussion

The fitting approach described below gave the following estimates based on quarterly data from 82:4 to 95:4.

λ_r	λ_μ	λ_σ	s^2	v^2
.0146	-3.75E-3	4.82E-4	6E-6	3.8E-8

The fit was reasonably good, which to some extent justifies the assumptions. However it seems unusual that λ_μ is negative, in that the risk adjustments are supposed to push rates up.

3.5 Fitting Parameters

Parameters were fit using term structure data for 1982:Q4 - 1995:Q4. For quarterly observations the evolution equations become:

$$\Delta r = \mu/4 + \varepsilon_1\sigma/2 \tag{25}$$

$$\Delta\mu = m/4 + \varepsilon_2 s/2 \tag{26}$$

$$\Delta\alpha = \varepsilon_3 v/2 \tag{27}$$

Here the ε's are random draws from the standard normal distribution. The term structure is what is observable, so the fitting is based on fitting equation (24) to the data for each quarter. The evolution equations (25) - (27), however, put constraints on the parameters. There are at least these constraints:

1. The T^4 and T^6 coefficients must be negative.
2. v^2 in the T^6 coefficient is the variance of the changes in the T^2 coefficients α.
3. s^2 in the T^4 coefficient is the variance of the changes in the T coefficients μ.
4. The average change in the T coefficient is m which is imbedded in the T^2 coefficient.
5. The constant term r_i changes by an average of μ_I from the T term and by a variance of σ_I, which is imbedded in the T^2 term.

The basic approach is to get coefficients of T^i by regression, subject to the constraints. Suppose we have a preliminary estimate of the coefficients of the T^i. Then some of the constraints can be used to separate regression coefficients into components, and then these can be used to check other constraints. For instance, the constant term r_i should change by an average of μ_i and by a variance of σ_i. These relationships can be used to estimate λ_r and σ_i. To see this in greater detail, to estimate λ_r, by adding up the changes in r we get:

$$r_n - r_1 \cong \Sigma\mu_i/4 \tag{28}$$

Then, since $n\lambda_r$ can be expressed as $\Sigma(\mu_i+\lambda_r) - \Sigma\mu_i$, (28) can be used to estimate λ_r, as the first sum can be calculated from the T coefficients.

Having estimated λ_r, the T coefficients then give the μ_i's. To estimate σ_i, use that:

$$E(r_{i+1}-r_i-\mu_i/4)^2 = \sigma_i^2/4 \tag{29}$$

The expression inside the expectation on the left-hand side of this equation can now be calculated for each i. Call it B_i. The following ad hoc method is one way to estimate the expectation at each i. Take a seven term centered moving average of the B_i at each point, with the middle three points getting double weight. To reduce the effect that extreme observations have on this average, trim each point to a maximum of twice its own centered seven point average. Then re-average the points to estimate the expected B_i at each point, and use that as the estimate of $\sigma_i^2/4$ from (22). There are clearly other ways to estimate the σ_i^2. Since these do not directly impact the term structure, their estimation is not critical to the overall fit. The smoothing of the B_i to estimate $\sigma_i^2/4$ is shown in the graph below.

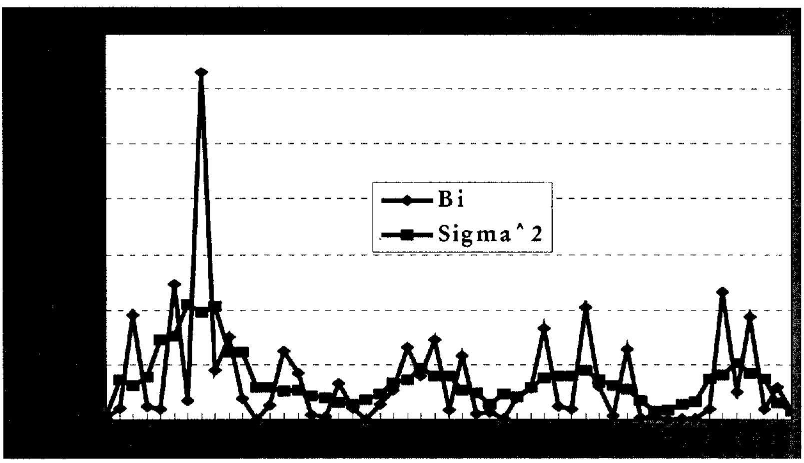

A similar procedure can be used to split out m_i and λ_μ from the T^2 coefficient. Since m_i is just $\alpha_i - \sigma_i$, subtracting σ_i from the coefficient just leaves $m_i + \lambda_\mu$. Using the fact that μ_i changes by an average of m_i and with variance s^2 (known from T^4) gives a way to estimate λ_μ. From the constraint (26) the following should hold:

$$s^2/4 = E(\mu_{i+1}-\mu_i-m_i/4)^2 \tag{30}$$

Each value of λ_μ implies values for the m_i's, so λ_μ can be estimated as the value that would give the resulting m_i's that satisfy (30). The m's that result from matching this variance are graphed along with the change in μ that they are meant to average in the graph below. The variance is quite large, so the fit may be reasonable.

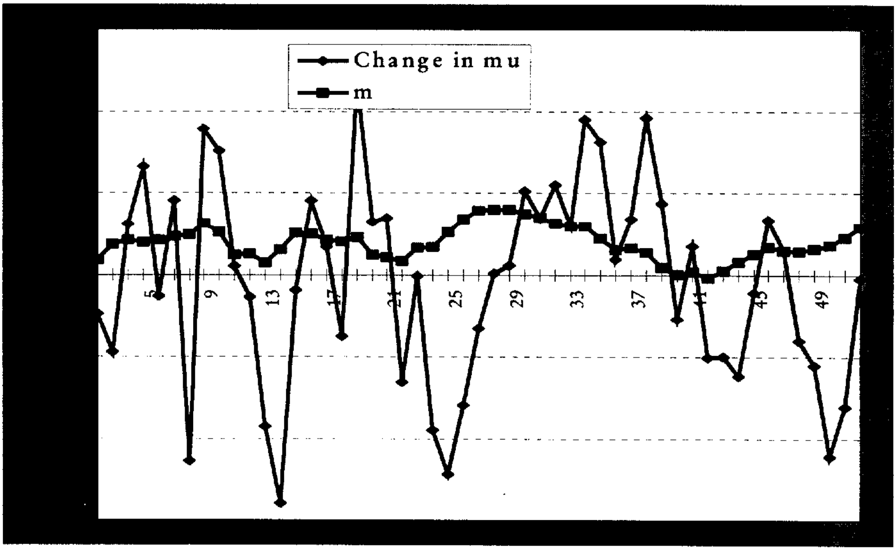

Another requirement is $\mu_n-\mu_1 \cong \Sigma m_i/4$, seen by summing (26) over all periods. This has not been used to estimate anything, so it was used as a test of the fit. The relationship $E(\alpha_{i+1}-\alpha_i)^2 \cong v^2$, from (27) was used as a constraint. The right side of this is the T^6 coefficient, and the left side is called the *implied* v^2.

A search procedure (simplex based) was used to fit the parameters. The search is looking for the three higher order coefficients defined by λ_σ, s, and v. What is minimized is the sum of squared errors between the actual and fitted yield rates at each period plus a weighting constant times the difference between v^2 and the implied v^2. In the minimization, for each trial λ_σ,s,v triplet, the sum of the higher terms, i.e., $T^3\lambda_\sigma/24-T^4s^2/40-T^6v^2/504$, is subtracted from the yield rates. This leaves a quadratic expression for the yields so adjusted for each quarter. The three quadratic coefficients are fit by regression for each time period to the observed interest rates for the terms used. These were 3 month, 6 month, 1 year, 2 year, 3 year, 5 year, 7 year, 10 year, and 30 years. This gives estimates for the coefficients r_j ,$(\mu_j+\lambda_r)/2$, and $(\alpha_j+\lambda_\mu)/6$ for each period j. Then the above approach is used to split these out into μ_j, λ_r, α_j, λ_μ, m_j, and σ_j. The sum of squared errors and the difference between the trial and implied v^2's are computed. New triplets are tried until the sum is minimized. That gave the parameters above. The resulting λ_μ of -0.00375 is roughly in the ballpark of the value of -0.00303 need to equalize $\mu_n-\mu_1$ and $\sum m_i/4$.

3.6 Goodness of Fit

The graph below shows the actual and fitted interest rates by term for a bad fitting quarter, a good fitting quarter, and the average of all the periods. The bad fit was actually somewhat exceptional, as the good fits were more typical.

The Best of Fits and the Worst of Fits (and the Average)

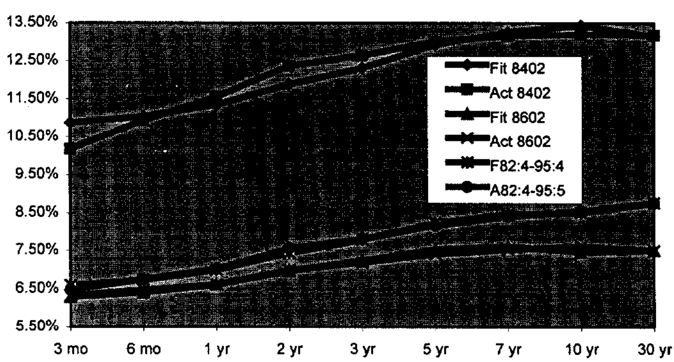

Although the fitting was simplified due to the closed form of the yield curve, the relationships between the yields and the movement of interest rates would hold in the more complex cases as well. Thus this example illustrates the interrelationships to be preserved in interest rate fits.

4 OTHER ECONOMIC VARIABLES

The term structure of interest rates incorporates investors' anticipations of future rates and thus implicitly of future levels of prices and economic activity. In recent years a number of articles have been published which attempt to forecast economic variables based on the term structure. For instance, see Eugene F. Fama *Term Structure Forecasts of Interest Rates, Inflation, and Real Returns,* Journal of Monetary Economics 25 (1990) pp. 59-76. It appears from this research that aspects of the yield curve do correlate with future economic activity, and so these correlations need to be taken into account when generating economic scenarios.

Forecasts will of course not be perfect, so when forecasting economic series from interest rates the prediction distributions will need to be taken into account. In a simulation context, for a given time frame the interest rate generator will produce yield curve scenarios, and then from each of those a prediction can be made of the other economic variables. Then a random draw from the prediction distribution can be made to produce a specific simulated scenario that includes interest rates and other series. This procedure should produce scenarios that are realistic over the time period chosen and with reasonable relative probabilities.

4.1 Examples of Prediction of Economic Variables from Yield Curves

To illustrate this process, two economic series are estimated from the term structure: the Consumer Price Index (CPI) and the Wilshire 5000 Index (W5). These both could have significant impact on insurer financial results.

Measures of the term structure typically are the interest rates for different terms as well differences between interest rates for different terms, e.g., the 10 year rate minus the 3 year rate. These are used at various lags. In this exercise all rates and lags are in multiples of a calendar quarter, so for notational purposes the time periods will be expressed as quarters. Notation such as 3L40:12 will denote the third lag of the difference between the 40 quarter and 12 quarter interest rates, i.e., the 10 year rate less the 3 year rate seen 9 months ago. Without the colon 0L40 is just the 10 year rate for the current quarter. Thus the notation comes with an actuarial spirit.

4.2 Consumer Price Index

The variable estimated here, denoted qccpi, is the ratio of the CPI for a quarter to that for the previous quarter. The variables used in the fit along with indications of their significance are shown in the table below. The data used is from the fourth quarter of 1959 to first quarter 1997, as this was available from pointers within the CAS website.

Change in CPI

Variable	Estimate	T-statistic	Significance Level
1:4Lqccpi	0.9994	1649.4	<.01%
0L40:4	-0.2668	-5.3349	<.01%
2L40:20	0.8486	4.6411	<.01%
3L2:1	0.7182	3.4663	.07%

The most important indicator of inflation is recent inflation. The variable used to represent this, denoted 1:4Lqccpi, is the average of qccpi for the past four quarters. The coincident variable, 0L40:4 has a negative coefficient. This may be due to inflation influencing current interest rates, but with a greater impact on short term than long term rates, thus flattening the yield curve. At lag 2 quarters, the coefficient for 2L40:20 is positive and at lag 3 quarters that for 3L2:1 is positive. These indicate a general tendency for a steeper yield curve to anticipate future inflation. Other yield spreads also appear to have significant impact on inflation, but only a few can appropriately be included in any one regression. Interest rate series are highly correlated, and many that do not enter the formulation will still end up having significant correlations with the inflation rates produced. The coefficients suggest that over 80% of any increase in these yield spreads will be reflected in subsequent inflation.

The r-squared, adjusted for degrees of freedom, is 65%. The standard error of the estimate is 0.0051. Thus the typical predicted quarterly change is accurate to about half a percentage point. The standard error is the standard deviation of a residual normally distribution around the predicted point, which can be used to draw the scenario actually simulated. The actual vs. fit is graphed below. The series can be seen to be fairly noisy, but the model does pick up the general move-

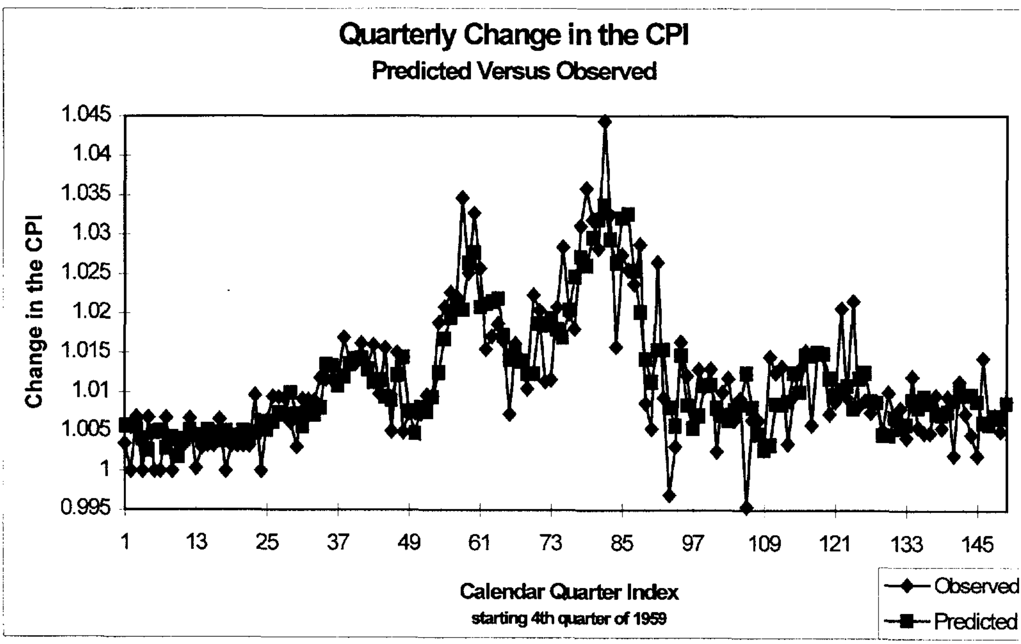

ments over time. The residuals are graphed on a normal scale below. Normality looks to be reasonably consistent with the observed residuals.

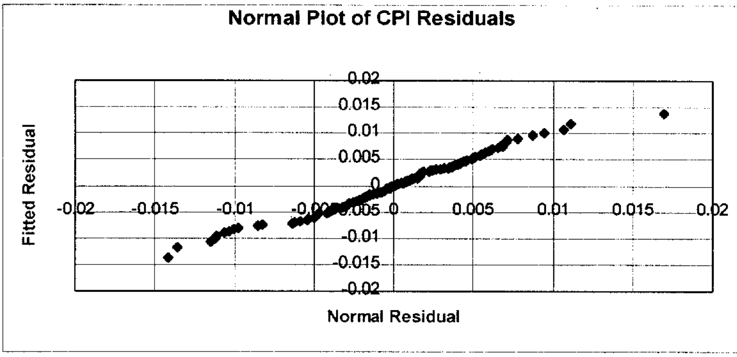

4.3 Wilshire 5000

The variable modeled, qcw5, is the ratio of the W5 at the end of a quarter to that at the previous quarter end. In this case the CPI percentage change variable qccpi was included in the regression as an explanatory variable. This allows creation of scenarios that have simulated values of W5 that are probabilistically consistent with the CPI value for the scenario.

The fitted equation for quarter ending data 1971 through first quarter 1997 is shown in the table below. In this regression only two variables were used, but they are composite series. The first, denoted 0-4Lqccpi, is the increase in qccpi over the last year, i.e., the current rate less the rate a year earlier. This variable has a negative coefficient, indicating that an increase in inflation is bad for equity returns. The other variable is denoted qcrelsprd. It represents the previous quarter's increase in the long-term spread less this quarter's increase in the short-term spread. Here the long-term spread is the difference between 10-year and 5-year rates, and the short-term spread is the difference between 6-month and 3-month rates. The increases noted are the quarter-to-quarter arithmetic increases in these spreads.

The coefficient on qcrelsprd is positive. This variable is positive if the increase in the short-term spread is less than the previous increase in the long-term spread, or if its decrease is greater. Either could suggest moderating inflation and interest rates, and thus be positive for equity returns.

Quarterly Change in Wilshire 5000

Variable	Estimate	T-statistic	Significance Level
0-4Lqccpi	-2.7113	-3.1936	0.2%
qcrelsprd	11.869	4.5273	<.01%
constant	1.02316	145.311	<.01%

The adjusted-r-squared is only 24% for this regression, indicating that the fit is not particularly good. The residual standard deviation is .0721, which allows a

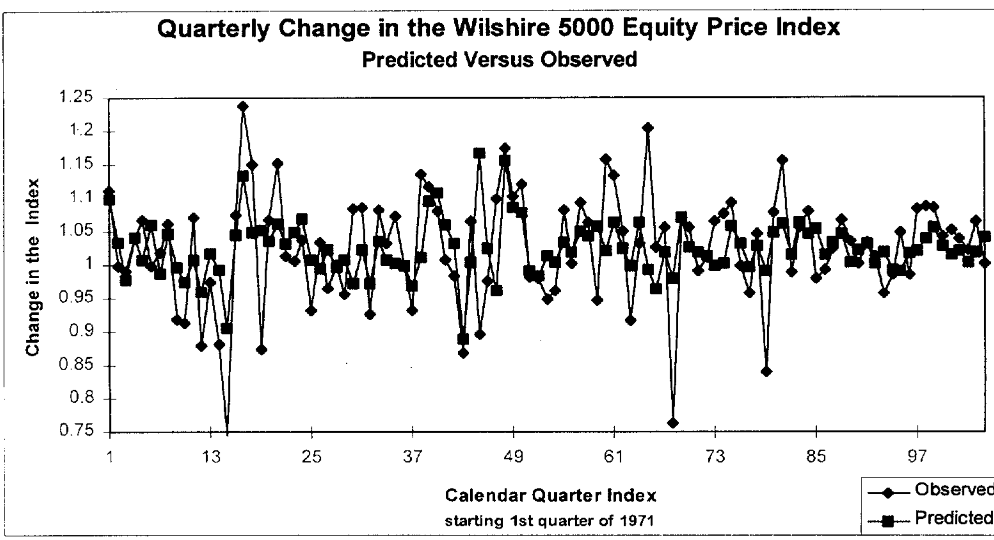

fairly wide deviation from the model. The actual vs. fit is graphed below. The residuals also appear to be more heavy-tailed than for a normal distribution. They are graphed on a normal scale below.

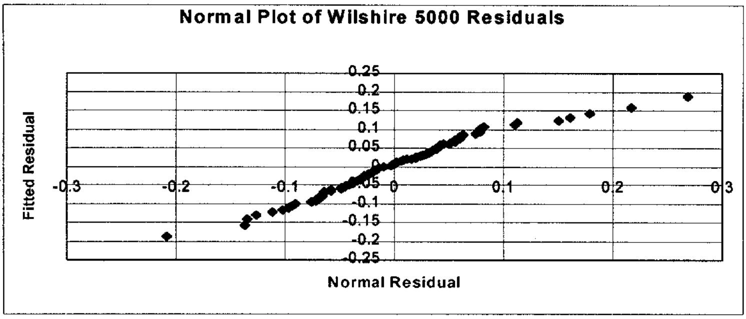

An alternative would be to simulate equity returns independently of interest rates and inflation. However, even the weak relationship found here would reflect the correlations that are likely among these variables, and would thus be preferable to assuming independence. More research into appropriate models for equity returns would be worthwhile.

Fitting percentage changes typically gives low r-squareds. The fit is usually better when translated to cumulative. The graph below shows cumulative products of the actual and fitted changes since 1974. The fit appears better on this basis.

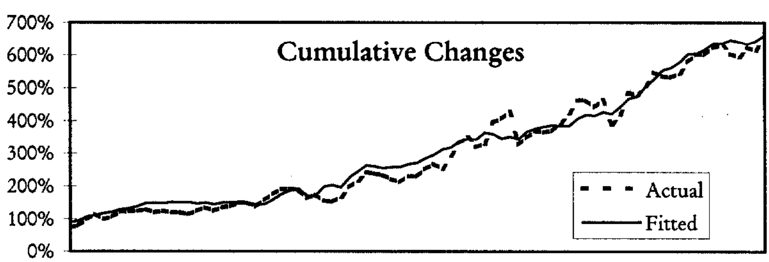

The fit is not unbiased for products of factors over the entire horizon, so the above graph begins at a point selected to give a horizon where it is unbiased. A fit in the logs of the change ratios is unbiased for products when exponentiated. The graph below shows this fit for the same variables used in the original regression. The fit is actually not quite as good as the original. Both are fairly close cumulatively for the past ten years, however.

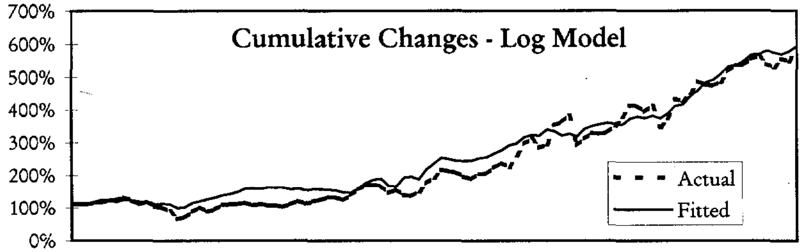

5 APPLICATIONS

Duration matching, which seems to work well for life insurers, is problematic for P&C carriers, who have shorter duration liabilities, and so would have to give up expected return to match. Simulation studies have suggested that going longer on assets provides a margin to P&C insurers, which can compensate for duration mismatch. Realistic stochastic asset generators may help quantify this trade-off. Even duration matching and its refinement to convexity matching do not provide complete hedges against interest rate movements. As an alternative, the robustness of investment strategies can be tested against the whole range of possible outcomes, by probability level, by measuring against simulated assets.

Asset simulations can be tied to liability simulations as well, e.g., by linking inflation movements to loss trends. The total risk of assets and liabilities can thus be quantified simultaneously by such dynamic financial analysis.

Acknowledgment: The valuable assistance of John Gradwell in gathering data and fitting models is gratefully acknowledged.

Appendix - Summary Evaluation of Variables

Variable Wilshire 5000 Equity Price Index

Rationale Broad-based indicator of value of equity investments

Source Downloaded from the website "Wilshire Index History", address wilshire.com/home/products. Could not find a pointer on the CAS website.

Method of Analysis Multiple regression based on series already modeled, which in this case were Treasury yields and the CPI.

Variable Consumer Price Index

Rationale Inflation measure that covarys with interest rates, trend factors, and other economic series.

Source Downloaded from the CAS DFA website, following the pointers "Data Access", "financial and economics databases", "Consumer Price Index", at the address http://205.230.252.34. However this site ends with data from first quarter 1995.

Method of Analysis Multiple regression based on series already modeled, which in this case were Treasury yields.

Variable US Treasury yields for 3 months, 6 months, 1 year, 2 years, 3 years, 5 years, 7 years, 10 years, and 30 years

Rationale Many insurers invest in Treasury securities, and when these are carried at market value their price will depend specifically on the interest rates. Other economic variables are correlated to interest rate movements.

Source Downloaded from the website maintained by the Saint Louis Federal Reserve Bank's "FRED Database" , address www.stls.frb.org/fred/, which has a pointer from the CAS DFA website.

Method of Analysis Simulation based on multi-factor arbitrage-free diffusion processes fit to historical interest rate movements and yield curves.