Building and Testing Economic Scenario Generators for P&C ERM

Gary G. Venter, FCAS, CERA, MAAA

Abstract

This paper addresses how to test economic scenario generators (ESGs) especially in the context of P&C ERM, where the risk to portfolios from a buy-and-hold strategy is important. This calls for real-world probabilities, in contrast to much ESG work, which is often aimed at risk-neutral probabilities for options pricing. Some advances in affine models of interest rates will be discussed, as well as issues in equity and foreign exchange modeling.

1. INTRODUCTION

As part of economic capital modeling, P&C companies simulate thousands of scenarios of their income statements and balance sheets, and the scenarios are desired to be realistic and have a realistic distribution. Most asset modeling is trading focused but P&C ERM estimates the probability distribution of price changes for a fixed portfolio over various time horizons, so real-world instead of risk-neutral probabilities are needed. ESGs simulate many scenarios of economic factor drivers such as interest rates, credit spreads, equity prices, inflation and foreign exchange, which are then used in an ECM to estimate distributions of portfolio values, as well as any implications for insurance losses.

Realism of the distributions produced is thus a key evaluation criterion. For instance, arbitrage-free yield curves are important: scenarios that allow arbitrage are not realistic and can distort any analysis. If portfolio optimization is to be done, having arbitrage opportunities available will push the analysis toward taking the arbitrage. There are similar problems of realism from forcing yield curve to follow smooth curves. But it is important to note that real-world yield curves are often produced by postulating risk-neutral processes for the short rate, so risk-pricing issues are not entirely avoidable even for real-world scenarios.

Actuaries would want similar realism for economic factors that drive liability values, but again they might need both real world and risk neutral models for liability drivers like inflation if estimates of the economic value of liabilities are needed, e.g. for firm-values calculations.

There are a lot of models used for economic scenario generation. We emphasize tests on the output of the models, treating the ESG itself as a black box. But models that reproduce various aspects of historical behavior are discussed. Interest rates are the most intricate factors to model, as they come in yield curves that are integrated across maturities (tenors) in numerous ways, and so they will be the primary emphasis. The paper is organized as follows. Section 2 discusses criteria for testing ESG output. Section 3 introduces several models of risk-free interest rates and how they are used to get to bond prices. Section 4 briefly discusses expansion of these models to include risky bonds, equities and inflation. Section 5 addresses foreign exchange modeling. Section 6 reviews some calibration methodology for the interest-rate models. Section 7 illustrates the testing procedures with a few selected tests. Section 8 concludes.

2. HOW TO EVALUATE ESGS

We look at measuring how well an ESG captures the statistical properties of historical processes. For interest rates this includes times series properties such as:

- Autocorrelation
- Moments of changes in rates
- Risks to excess profit potential of longer-maturity investments

And cross-sectional properties like:

- Distribution of yield curve shapes
- Volatilities by maturity

Both types of properties can be measured by looking at the simulations from the model, viewed as a probability distribution of scenarios. But even the cross-sectional target distributions are based on the time-series history. Feldhütter (2008) details historical properties of yield curves, with references to original studies. Some of the historical facts include:

- Rates for all tenors are highly are autocorrelated
- Rates are positively skewed and show excess kurtosis, but both are fairly modest
- Rates fluctuate around temporary levels that only very slowly revert to long-term means
- Rates tend to increase with maturity while volatility decreases with maturity
- Spreads between various longer rate tend to be lower when the short-rate is higher

- Volatility slowly reverts to long-term volatility but with occasional spikes
- As zero-coupon bonds mature, their values generally increase as their rates decrease. This increase tends to be greater when the yield curve is steeper, especially for longer bonds.

Testing ESG output can proceed by measuring such properties on the simulated data, and comparing with the historical record.

To illustrate the relationship of rate spreads with the short rate, Figure 1 shows a history of the 10-year – 3-year spread and the 3-month rate for US Treasuries.

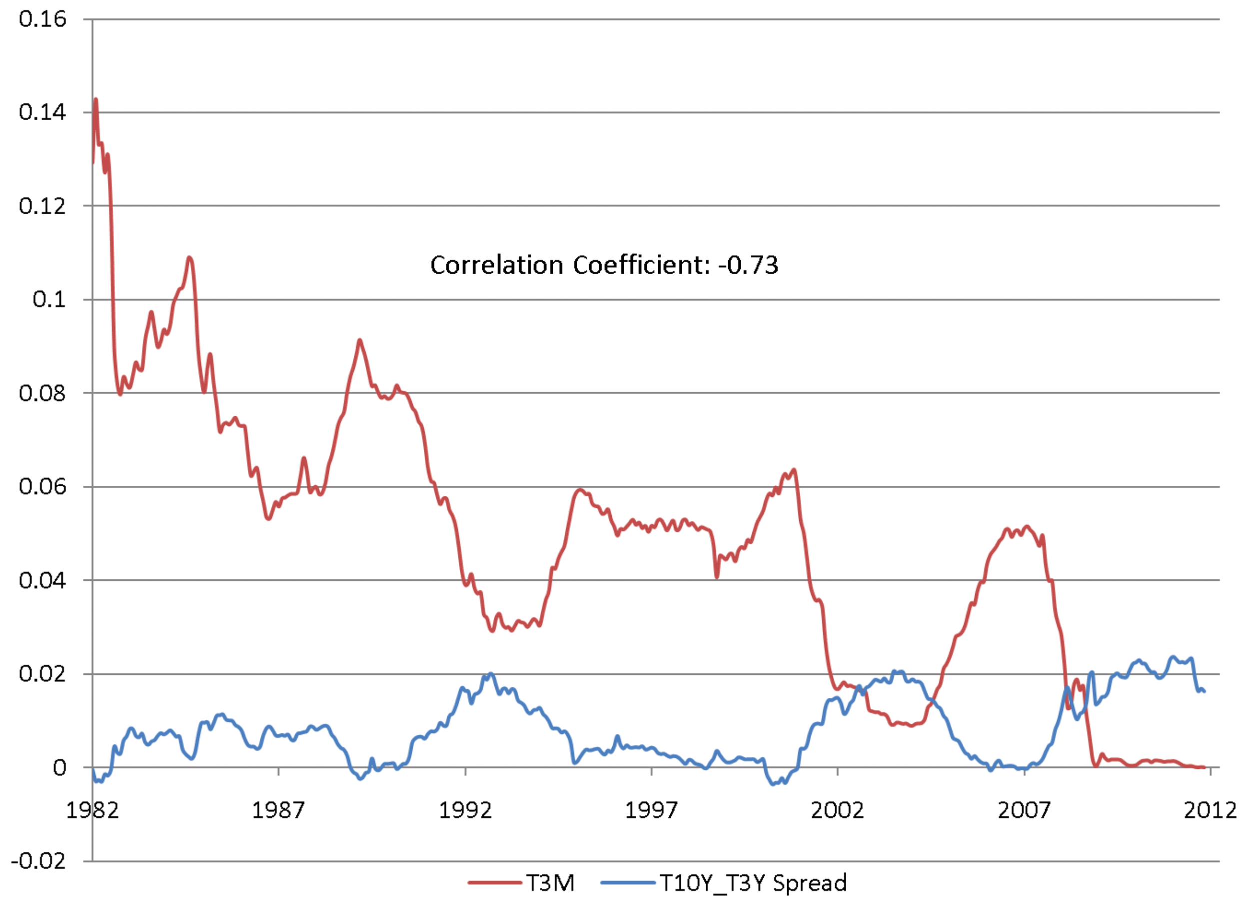

Figure 1: History of US Treasury 3-Month Rate and 1-Year – Three-Year Spread

Figure 2 shows a regression line for the spread as a function of the short rate from 1995 to 2011 with points coded for various subperiods. The spread does not get very far from this line for the entire period, and that property can be used as a test of yield curve shapes from the simulated distribution.

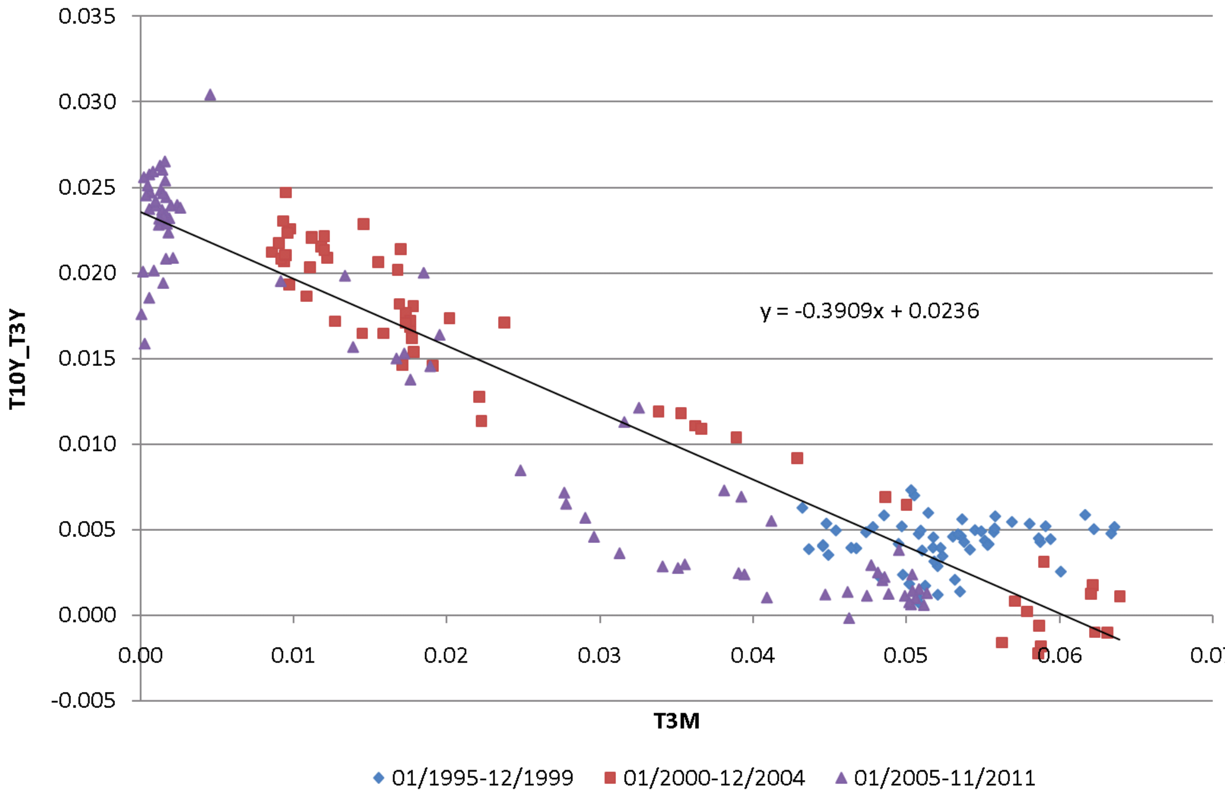

Figure 2: Spread as a Function of Short Rate 1995 - 2011

3. MODELING CHOICES

Interest rates are often modeled by short-rate models or by forward-rate models. The short-rate models postulate a process for the short rate's evolution and derive implied yield curves from that process plus a market-price-of-risk adjustment. Forward-rate models simultaneously model the evolution of all (usually risk-neutral) forward rates and back out real-world implied yield curves by arbitrage arguments. Here we will primarily focus on short-rate models that have a formulaic approach for directly calculating the implied yield curves – the affine models and related approaches.

Affine models are typically diffusion processes, usually multi-factor, for the short-rate that have solutions worked out for the implied longer-bond prices as linear exponential functions of the short-rate, with non-linear functions of time and the market price of risk included. The risk-neutral short-rate process is usually just the real-world process plus an additional deterministic drift on the short rate. This part is generalized in extensions to semi-affine and extended affine models.

Typically risk-free rates are modeled then another model is used to get to risky yield curves. This has been the most challenging step for rate modelers.

All the financial series seem to be more accurately modeled by making volatility itself stochastic. However this refinement is not needed for every problem. Modeling financial series over a single time period with no need for options pricing, which is typical for economic capital models, is not necessarily improved by considering stochastic volatility.

Inflation can be defined and quantified in various ways depending on the applications needed, and can often be modeled satisfactorily by time-series methods. Linking the evolution of inflation to that of interest rates is more problematic. Both processes display instability of correlations and auto-correlations, for example. Over long periods inflation rates and interest rates display quite similar changes, but they may be fairly uncorrelated over shorter time frames. This is typical of co-integrated processes and some models from that theory may be applicable.

Foreign exchange (FX) rates can be highly correlated for some economies and practically uncorrelated for others. Also for some pairs of FX rates, most of the correlation can come from tail events, with little correlation elsewhere in the distributions. Some advanced copulas, like the grouped-t, are capable of modeling this reasonably well. Also typical FX models, like interest-rate parity, have been found not to work in practice. But FX rates are volatile enough that modeling the mean drift has little effect on the overall distribution of outcomes.

3.1 Affine model example – the CIR (Cox, Ingersoll, Ross) model

The CIR model gives a stochastic differential equation for the evolution of the short rate. It uses Brownian motion, which is a continuous version of a time-series model, due to the fact that trades can be quite frequent and are more conveniently modeled by postulating that they are realizations of a continuous price process. A standard Brownian motion *B(t)* is defined as a continuous process whose changes over time periods of length *t* are iid normally distributed with mean zero and variance *t*. The CIR model postulates that the evolution of the short rate at *t*, *r(t)*, is given by:

$$dr(t) = k[\theta - r(t)]dt + \eta\sqrt{r(t)}dB_r(t)$$

Here $dr(t)$ can be also thought of as $r(t + dt) - r(t)$. θ is the reverting mean. If $r(t)$ is above θ, the drift (i.e., dt) term is downward, and vice versa, so the deterministic drift is always towards θ. Then k is the speed of mean reversion. The Brownian motion has variance 1 over time 1, often taken as a year. The CIR process postulates that the standard deviation of the volatility term is proportional to the square root of the rate. If the rate ever got to zero, that term would then be zero, and the drift would be upward, so the rate would then become positive. Thus zero is a reflecting barrier for a square-root process.

To get the yield curve, a market price of risk λ is first added to the drift to get a risk-neutral process with higher future short rates in general. Then the price of a bond at a future point is just the expected value of that bond payoff discounted continuously over all paths of the short rate according to the risk-neutral process. This maintains the rather confusing terminology in which the risk-neutral process is the one with additional risk, and so the one in which prices are simply means. The market price of risk is traditionally expressed as a factor times the volatility and is optionally expressed as a multiple of the short rate. Here we will write the risk-neutral process as:

$$dr(t) = [k\theta - kr(t) + \lambda\eta r(t)]dt + \eta\sqrt{r(t)}dB_r(t)$$

Denote the price at time t of a bond paying 1 at time $T+t$ as $P(t,T)$. Also let $c = k - \lambda\theta$, and $h = (c^2+2h^2)^{½}$. Then the bond price is $P(t,T) = A(t,T)e^{-B(t,T)r(t)}$, where

$$A(t,T) = \left[2h\left(2h + (c+h)(e^{hT}-1)\right)^{-1} e^{(k+h)T/2}\right]^{2k\theta/\eta^2}$$

$$B(t,T) = 2(e^{hT}-1)\left(2h + (c+h)(e^{hT}-1)\right)^{-1}$$

Because this is continuous compounding, the long rate is given by $R(t,T) = -\log P(t,T)/T$.

It turns out that single-factor models like this are too simple to capture the richness of yield-curve shapes that arise in practice. In this theory, the number of factors is the number of equations that have random draws in them. A fairly successful generalization is to postulate that the short rate is the sum of three independent unobserved CIR processes. Each process has an affine formula for the bond price, and the actual bond price turns out to be the product of the three partial bond prices, so is itself an affine formula. The partial yield rates then add to get the yield rate. Jagannathan et al. (2003) find that "with three factors the CIR model is able to fit the term structure of LIBOR and swap rates rather well. The model is able to match the hump shaped unconditional term structure of volatility in the LIBOR-swap market." However they find it does not do well at pricing volatility-dependent options. This is not surprising, in that this model does not include stochastic volatility.

3.2 Stochastic volatility affine models

A popular stochastic volatility affine model is the BDFS model. It is similar to the CIR model but with the reverting mean θ and the volatility each themselves following stochastic processes. This is thus a three-factor model. The model is, with (t) omitted:

$$dr = k[\theta - r]\,dt + v^{0.5}dB_r$$

$$d\theta = \alpha[\beta - \theta]\,dt + \eta dB_\theta$$

$$dv = a[b - v]\,dt + \phi v^{0.5}dB_v$$

$$\mathrm{cov}(dB_r\, dB_v) = \rho dt$$

Now only the volatility follows a square-root process, and it is correlated with the interest rate. A similar model, the Chen model, has the reverting mean also following a square-root process, but without any correlation. The bond price is considered almost closed form. It is given by:

$$P(0,T) = A(T)e^{-rB(T) - \theta C(T) - vD(T)}$$

Here the *B* and *C* functions are closed form functions of the parameters of the model, but *A* and *D* are the solutions of a system of differential equations. This is easy to solve by common numeric methods, so is considered almost closed form. The Chen model has a closed form solution for the bond prices but it uses functions that require numeric calculations, like the hypergeometric. Notation for affine models would identify the BDFS model as an $A_1(3)$ model. In this notation, as in actuarial notation, the "A" doesn't really mean anything. It is just a place to hang subscripts. The "3" indicates this is a 3-factor model, and the "1" means that 1 of the 3 factors (here v) affects the volatility of the processes.

Another popular stochastic volatility model is the $A_2(3)$ model. The Chen model is an example because θ is a function of its own volatility and the volatilities of both r and v are functions of v. The general version of $A_2(3)$ starts with two unobserved correlated square-root processes Y_1 and Y_2. Then both the reverting mean θ and the volatility v are taken as linear functions of Y_1 and Y_2. With constant terms, that takes 6 parameters for the 2 linear functions. The evolutions of the factors are:

$$dY_1 = k_{11}[\theta_1 - Y_1]dt + k_{12}[\theta_2 - Y_2]dt + \sqrt{Y_1}dB_1$$

$$dY_1 = k_{21}[\theta_1 - Y_1]dt + k_{22}[\theta_2 - Y_2]dt + \sqrt{\beta Y_2}dB_2$$

$$dr = k_{33}[\theta - r]dt + \sqrt{v}dB_r + p\sqrt{Y_1}dB_1 + q\sqrt{Y_2}dB_2$$

Here θ_1, θ_2, the ks, ϕ and β are constants and everything else is a function of t. This more general formulation allows for more interactions in the evolutions of the factors, which appears to improve model performance substantially. The bond-price formula is similar to that for the BDFS model.

4. OTHER FINANCIAL FACTORS

4.1 Risky bonds

Coordinating corporate bond yield curves with risk-free curves can be a tricky modeling exercise. There are all the same issues about shape of the curve, volatility term structure, etc. as well as correlations among risky and risk-free rates, and the distribution of credit spreads. Some of the stylized facts about corporate rates are:

- Longer yields are higher and have lower variance
- There is a wide variety of yield-curve shapes
- Spreads to Treasuries increase for longer tenors
- Higher short rate linked to a compression of spreads among longer maturities
- High correlation of Treasury and risky rates, but decreasing by maturity
- Correlation of lags of Treasury and risky rates still high but decreases by lag
- Spreads can move with or against Treasury rate movements, in part associated with inflation and economic activity

Duffie and Singleton (1999) propose adding the short-term credit spread to the $A_2(3)$ model as another linear combination of the unobserved factors Y_1 and Y_2. That would make the yield spreads as well as the volatility and reverting mean of the short rate linear combinations of the two drivers, with non-negative coefficients. With θ and v the reverting mean and volatility of r, take s as the risk-free to AA spread and let u be the AA to BBB spread. Writing these as linear functions of Y_1 and Y_2:

$$\theta(t) = a_\theta + b_\theta Y_1(t) + c_\theta Y_2(t)$$

$$v(t) = a_v + b_v Y_1(t) + c_v Y_2(t)$$

$$s(t) = a_s + b_s Y_1(t) + c_s Y_2(t)$$

$$u(t) = a_u + b_u Y_1(t) + c_u Y_2(t)$$

Adding the spreads to the risk-free short-rate gives the short risky rates, and these can be used in the bond-price formulas to get the risky yield curves. One constraint is that this theory links the market prices of risk to the factors, and since the spreads are deterministic functions of the Y_1 and Y_2 factors, the same market prices of risk have to be used for all the yield curves. This works better with the more flexible semi-affine and extended affine market prices of risk.

4.2 Cost levels

P&C insurance losses are affected by cost increases in medical and construction prices, among others. There are somewhat relevant price indexes that track some of these. Time series models are able to capture a good deal of the statistical properties of the indexes, such as the autocorrelation spectrum. For instance, the CPI medical index is reasonably well modeled as the sum of two unobserved AR1 processes. This and other cost indexes are highly autocorrelated. If they get high, they tend to stay high for quite a while. That means that an increase in cost levels is indicative of future similar increases. This has implications for both pricing and reserving.

One way to integrate multi-factor price-level and interest-rate modeling is to correlate interest and price-level individual factors separately. This can model the auto-correlation of both series and the observed correlations between them over short time periods. One problem with this approach, however, arises from the fact that the correlation between price-level changes and interest rates seems to be higher over longer time periods. This behavior cannot be produced by models of correlated short-term changes in the series if the changes are assumed to be independent across time periods. Some kind of structural link between prices and interest rates would need to be postulated.

A further complication arises if prices are needed for inflation-adjusted bonds, or other pricing where a risk-neutral inflation series is needed, such as the market value of insurance liabilities. Perhaps this could be accomplished by adding inflation as a 4^{th} factor to the $A_2(3)$ model, with reverting mean and volatility also functions of Y_1 and Y_2. Then this factor would also get a market price of risk for pricing purposes. This could be calibrated to inflation-adjusted bond prices, for instance.

4.3 Equity prices

Early equity-price models were based on geometric Brownian motion, which means that the log of the price follows a Brownian motion. The Black-Scholes options-pricing formula is derived from this model. However actual options prices suggest that some more complicated process is at work. Options that are for very short time periods and options with strike prices further away from current levels are both relatively more expensive than others compared to what this theory would imply. For instance, Figure 3, taken from the sctcm blog, shows the Black-Scholes volatilities implied by actual options prices on 8/6/2010. These price differences are evident in the figure.

An early model by Merton added lognormally distributed jumps from a Poisson process to the geometric-Brownian-motion model. This improves the model performance, but subsequent work has found that up jumps and down jumps have both different frequency and different severity. They are also more heavy-tailed than lognormal. In logs, exponentially distributed jumps seem to work well. See for example Kou and Wang (2004). Ramezani and Zeng (2007) estimate parameters for the S&P 500. Adding stochastic volatility still allows solving for the options prices and improves fits.

Poisson jumps have also been included in affine models of interest rates. Andersen et al. (2004) provide an example of such a model.

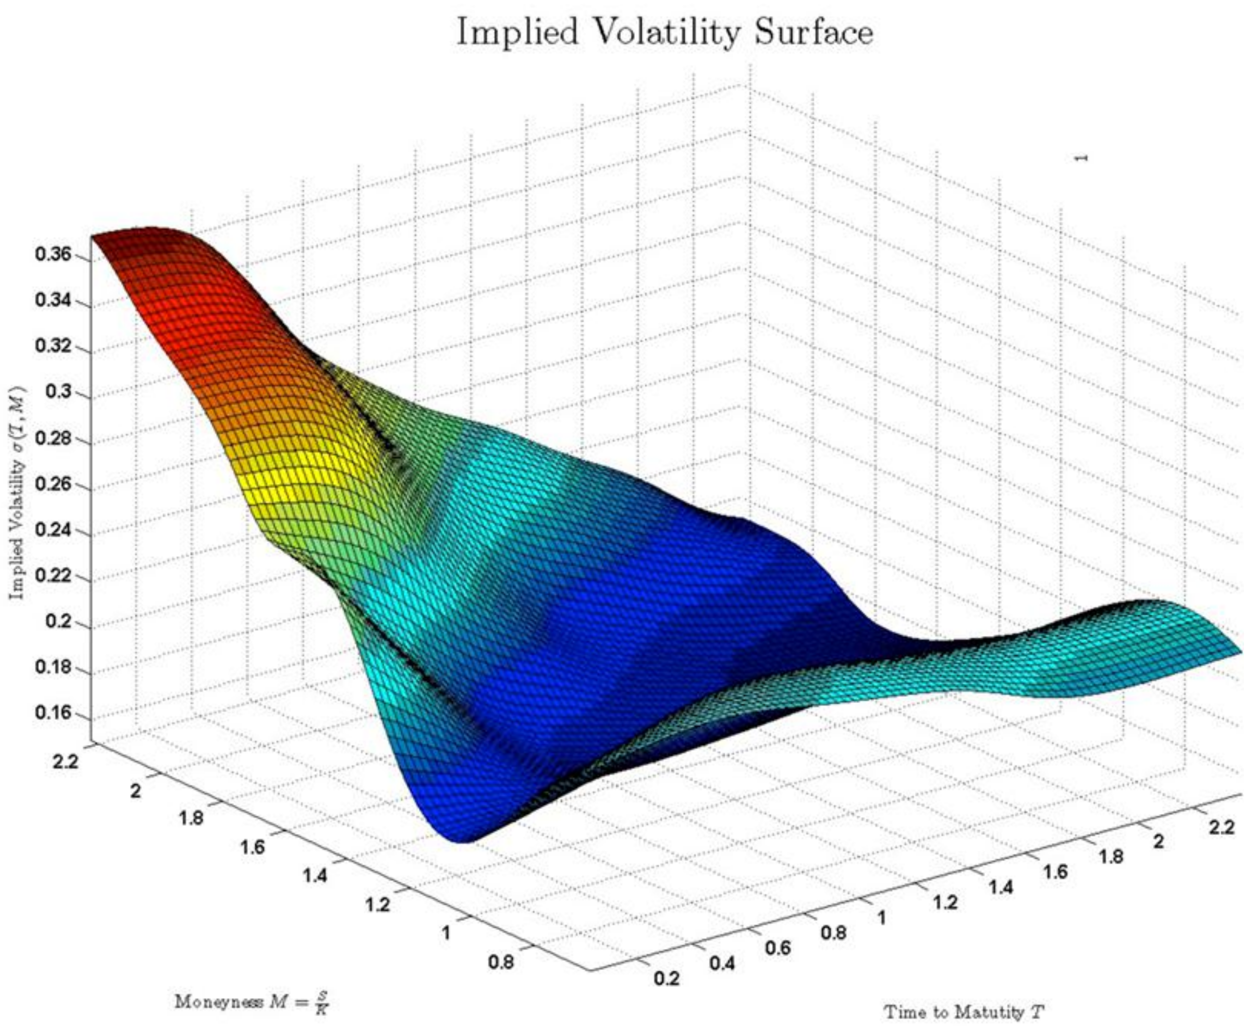

Figure 3: Implied Volatilities 8/6/10

5. FX MODELING

Popular models of exchange rates include interest-rate parity and purchasing-power parity. The former assumes that the expected change in an exchange rate is that which would equalize the returns on bond investments when the bond rates differ in the economies. The latter forecasts equal prices for the same products across economies. Neither works very well and in fact movements in exchange rates are often directly opposite to what these models would predict. Froot and Thaler (1990) discuss possible reasons for this. To some degree this is irrelevant for risk modeling, as volatility tends to dominate mean changes in the FX markets. Correlated AR1 processes can capture much of the movement in exchange rates.

It is convenient to compare exchange rates for the US $ across economies. These tend to be correlated, especially among countries close to each other, as illustrated in Figure 4.

Correlations for Monthly Change in Exchange Rate Against US $
1-1-71 to 1-1-007

	UK pound	Aus $	NZ $	Swiss Fr	Canada $	Japan Yen	S Kroner	D Kroner	N Kroner
UK pound	100%	25%	42%	66%	21%	44%	64%	68%	69%
Aus $	25%	100%	63%	25%	38%	29%	30%	27%	29%
NZ $	42%	63%	100%	39%	28%	38%	41%	44%	44%
Swiss Fr	66%	25%	39%	100%	16%	59%	72%	89%	80%
Canada $	21%	38%	28%	16%	100%	11%	25%	22%	24%
Japan Yen	44%	29%	38%	59%	11%	100%	45%	56%	51%
S Kroner	64%	30%	41%	72%	25%	45%	100%	81%	82%
D Kroner	68%	27%	44%	89%	22%	56%	81%	100%	88%
N Kroner	69%	29%	44%	80%	24%	51%	82%	88%	100%

10%	19%
20%	29%
30%	39%
40%	49%
50%	59%
60%	69%
70%	79%
80%	89%

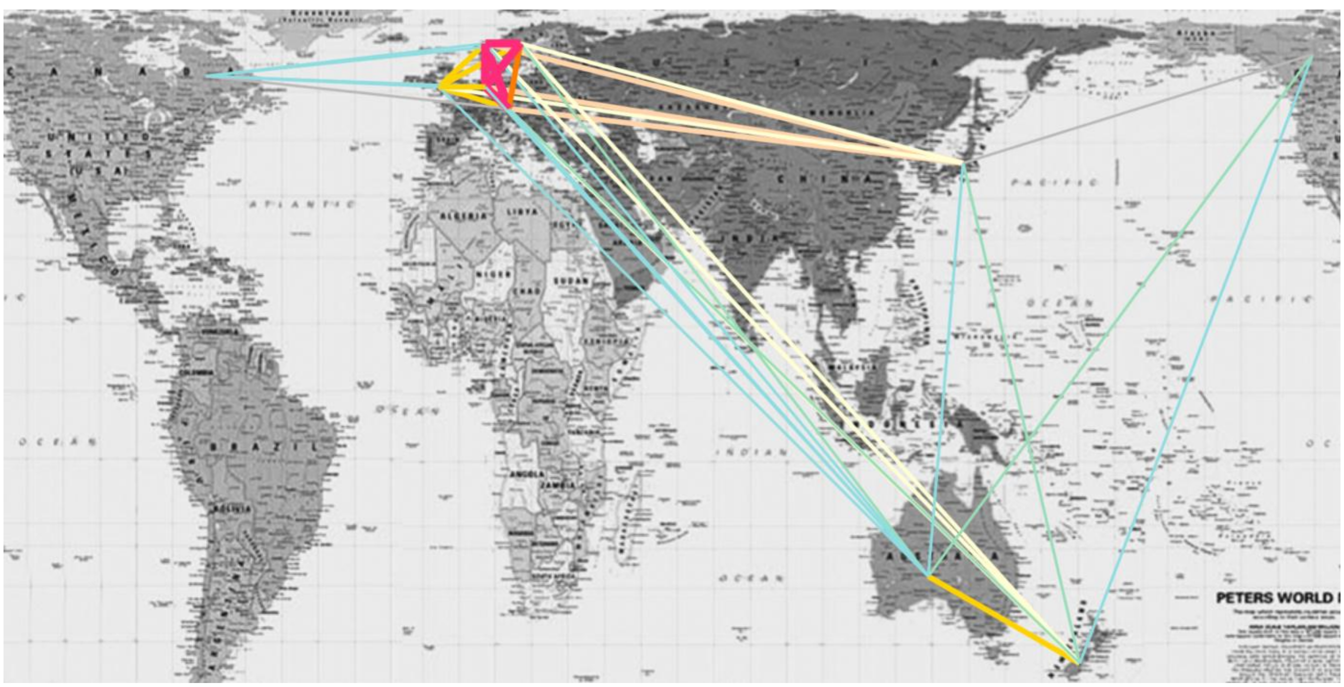

Figure 4: US$ Rate FX Correlations 1971-2006

These correlations do not remain stable over long periods, however. Experience indicates that about a decade is a reasonable period to use to estimate current correlations. However correlation alone does not measure all of the co-movements of FX rates. Some rates are correlated only in the cases of extreme movements and not much usually. This so-called tail correlation is measured by the tail dependence, which for two variates Y and Z for the right tail is the limiting value, as Z increases, of R(x) = Pr(Y>x|Z>x), and similarly for the left tail.

The normal copula has tail dependence of zero, and the t-copula has tail dependence that increases monotonically with the correlation. Neither behavior is observed in the market, as some currencies with fairly low correlation can have higher tail dependence. The grouped or individuated t-copula is able to model this situation. This IT copula, however, does not have a closed-form density. It is given by:

$h_n(y) =$ inverse chi-squared distribution with ν_n dof at y
$w_n =$ inverse of t-distribution with ν_n dof at u_n, over $\sqrt{\nu_n}$
$J_{nm} = n, m$ element of inverse of correlation matrix ρ
Then IT copula density at $(u_1, \ldots u_n)$ is:

$$\int_0^1 \frac{\prod_{n=1}^{N} \sqrt{(1+w_n^2)^{1+\nu_n} h_n(y)}\Gamma(\nu_n/2)\Gamma([1+\nu_n]/2)}{exp\left(\frac{1}{2}\sum_{n,m} J_{nm} w_n w_m \sqrt{h_n(y)h_m(y)}\right)\sqrt{det(\rho)2^N}} dy$$

Nonetheless it is not difficult to compute this numerically for MLE purposes. Dividing the currencies into two groups gave a reasonable fit – with the larger currencies having more tail dependence. Large joint movements in the $ rate probably has something to do with the $ more than other economies, so there tends to be tail dependence among larger currencies, whereas smaller currencies are more likely to have large idiosyncratic movements.

6. PARAMETERIZATION BY SIMULATED METHOD OF MOMENTS (SMM)

All the parameters of the short-rate affine models govern the short-rate process except for the market-price-of risk parameters, which affect the whole yield curve. Typically all of the parameters are fit simultaneously to the history of yield curves. However a drawback to this approach is that market prices of risk tend to change over time, and the models do not provide for this. The extended-affine approach is one response to this issue. However for the affine models an alternative approach is to fit the short-rate parameters to the short-rate observed process, and then fit the market prices of risk to the current yield curves when the projections are run. This would include calibrating the current values of the unobserved processes. We tried this approach for the $A_2(3)$ model.

SMM is a fitting method that first identifies key statistical features of the data set being modeled, then postulates trial parameters which are used to simulate a long future path for the series. The statistical properties of the simulated series are then measured and compared to the historical properties, and some measure of the distance between simulated and historical is selected, such as weighted sum of squared differences. Then the postulated parameters are changed using non-linear optimizers to minimize the distance. The statistical properties are called generalized moments, hence the name. The generalized moments and weights are selected based on properties that are important for the intended application, and so are not necessarily aimed at finding the true generating process.

For the US Treasury, AA, and BBB short rates we selected the autocorrelation spectra of the series, moments of changes of rates, moments and autocorrelation of absolute values of changes, and the cross-correlation spectra of the series as the generalized moments of interest. The autocorrelation spectrum of the absolute value of rate changes is a measure of stochastic volatility. If those values tend to be high for a while and low for other periods, that indicates that volatility is changing.

SMM has advantages and disadvantages. It is not efficient in the statistical sense. Other estimators like MLE and MCMC usually have lower parameter estimation variance. On the other hand, the robust estimation framework starts with the viewpoint that all the models are wrong anyway – there is actually zero probability that the data was generated by any of them. The actual processes are in fact much more complicated. The usefulness of the model comes from how well it represents interesting features of the data. Some quotes:

Cochrane (1996): “Efficient (estimation) may pay close attention to economically uninteresting but statistically well-measured moments.”

Brandt and Chapman (2002): “... the successes and failures of alternative models are much more transparent using economic moments. For example, it is easy to see that a particular model can match the observed cross- and auto-correlations but not the conditional volatility structure. In contrast, when models are estimated (by efficient methods), it is much more difficult to trace a model rejection to a particular feature of the data. In fact, the feature of the data responsible for the rejection may be in some obscure higher-order dimension that is of little interest to an economic researcher.”

SMM is not just a method for fitting interest-rate models. We tried it on the US medical care CPI, fitting a model that is a sum of two unobserved AR1 processes by SMM, MLE, and MCMC. This model is a special case of an ARMA(2,1) model, so that was what was fit by MLE. The selected moments for SMM were the autocorrelation spectrum of the series and moments of monthly changes and their absolute values. MCMC and MLE gave virtually the same fit, but SMM was a bit different. The empirical and fitted monthly autocorrelation spectra are graphed in Figure 5.

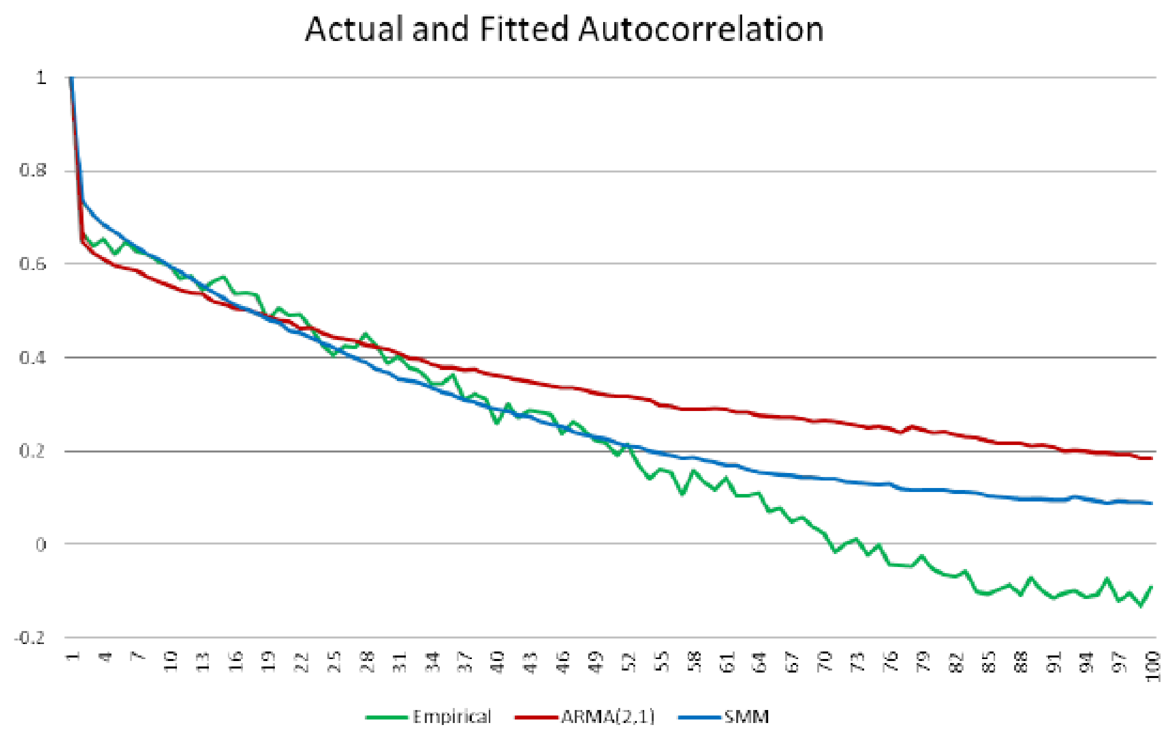

Figure 5: US Medical CPI Autocorrelation Empirical and Fitted

As can be seen, neither fit represents everything about the data. In particular after about 72 months the autocorrelations become consistently negative, suggesting some kind of cyclical behavior that this model does not capture. But the SMM fit better portrays the autocorrelations after 24 months or so. This could be critical in reserving, as the stickiness of inflation means that if it goes up, it tends to stay up for quite a while. Simulations of reserve risk need to reflect this behavior to get a good representation of what different paths could hold, but it should not be overstated either.

The interest-rate fitting showed similar problems with autocorrelation for Treasury, AA, and BBB rates. All of them showed some long-term cyclical behavior not well captured by the Brownian motion models. Cross-correlation measures correlations of lags of the series. For these three series that was one of the moments that we tried to match, both for the short-rates and changes in the short rates. Figure 6 shows actual vs. fitted, which seems ok.

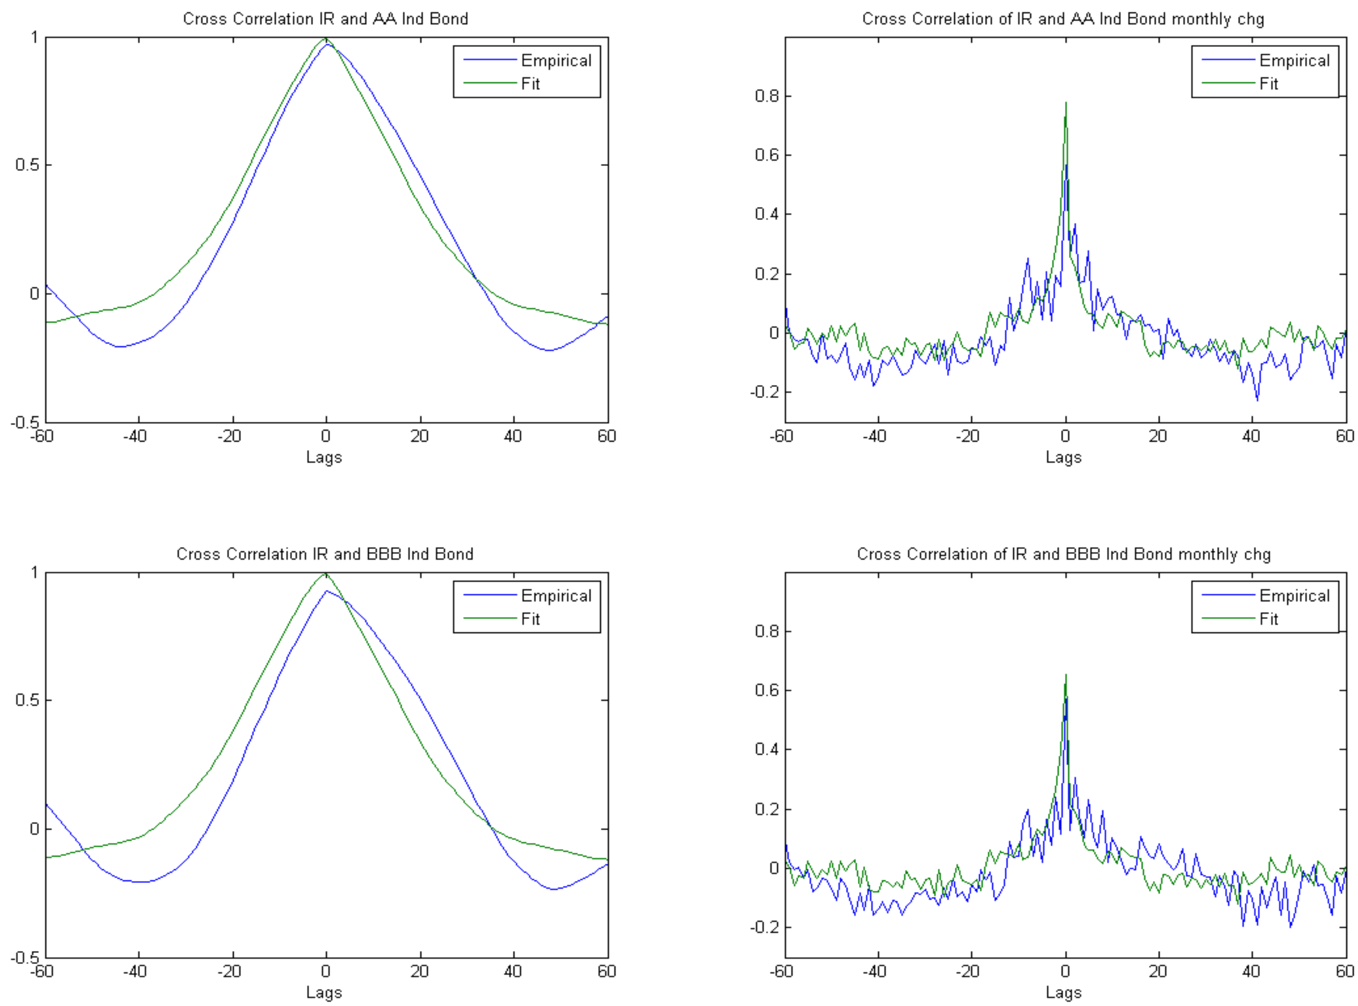

Figure 6: Cross-correlations of US Treasury (IR) and Corporate Short Rates Actual and Fitted

7. ILLUSTRATION OF TESTING ESG OUTPUT

As an illustration of how this methodology can be applied in model validation, we look at historical vs. modeled statistics for two sets of ESG output – one from the $A_2(3)$ model of risk-free and risky rates discussed above, called model A, and one we have output from for testing purposes, called model B. The output was only for the endpoint of a one-year horizon, so autocorrelation was not tested.

Figure 7 shows the mean spreads of the risky bonds to Treasuries by maturity for empirical and both models for AA and BBB rates. Empirically the spreads widen for the longer maturities. Model

A has this same behavior, but shows higher spreads than historical (1995-2011). That may not be a concern, depending on the current state of the market. Model B shows a quite different pattern that does not seem particularly reasonable.

Figure 8 graphs the standard deviations of the rates historically vs. across the simulated scenarios for Treasury and AA rates. Historically the two volatility series have very similar term structures,

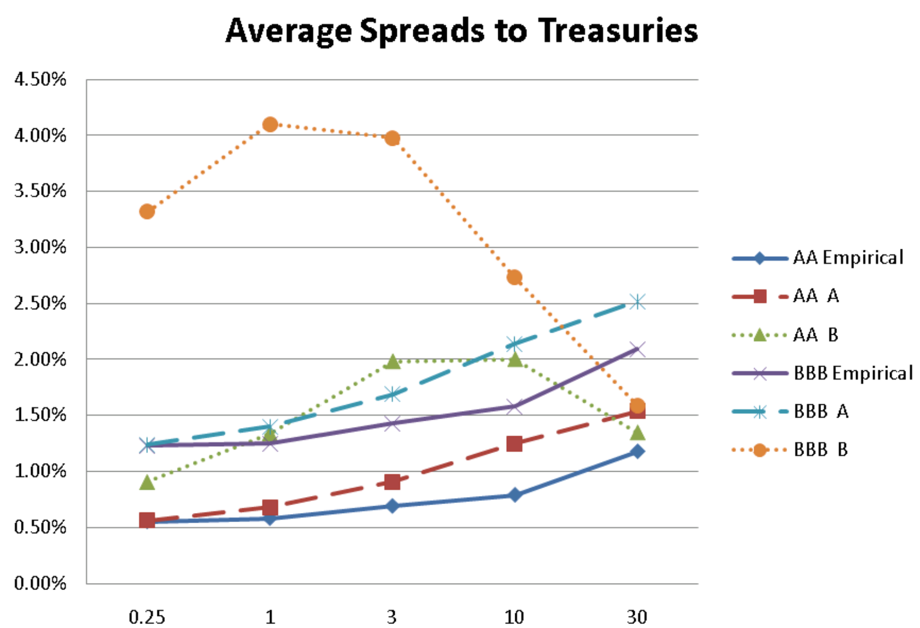

Figure 7: Mean spreads of risky rates to Treasury rates actual and projected

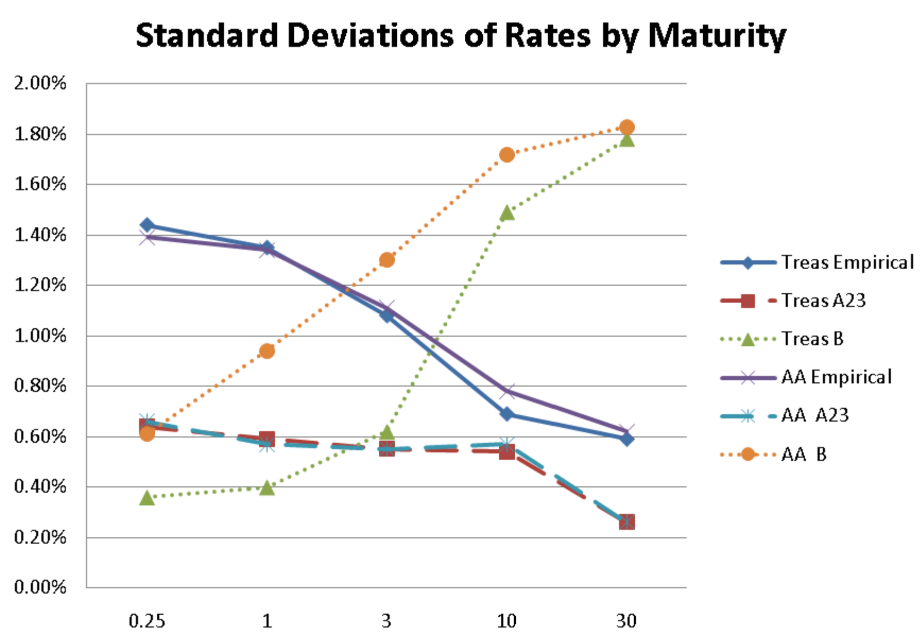

Figure 8: Standard deviations of Treasury and AA rates by maturity actual and projected

which are downward sloping. The $A_2(3)$ model also shows similar downward-sloping term structures, but has lower volatilities. That again may be ok for a one-year horizon. Model B, on the other hand, has upward-sloping term structures of volatility, with AA rates much more volatile. This seems difficult to reconcile with historical patterns.

Figure 9 compares correlations of AA and BBB rates with Treasuries. Empirically the term-structure is slightly hill-shaped, with the highest correlations in the middle of the curve. Correlations are higher for AA than for BBB. That last feature is shared by both models. However Model A does not have much separation between the curves. It also does not show correlations falling off at the 10-year and 30-year maturities the way the historical correlations do. However its correlations are much closer to historical than those from Model B, which are valley-shaped not hill-shaped, and has a very wide spread between AA and BBB at the short end, and little at the long end. Neither model really captures the correlations of the data, but Model A is closer.

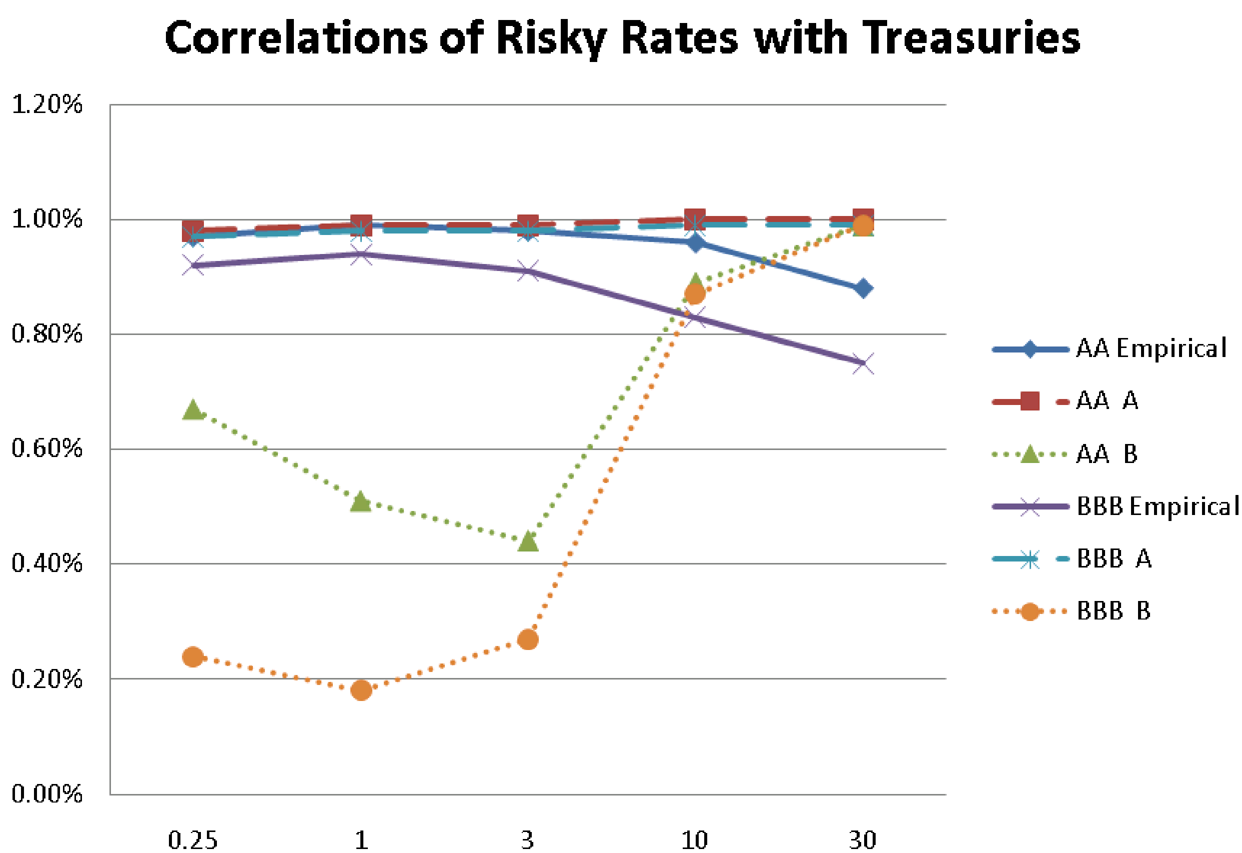

Figure 9: Correlations of risky and risk-free rates by maturity.

8. NEXT STEPS AND SUMMARY

We discussed several issues in the modeling of financial factors, and how to test the output of economic scenario generators against historical data. These tests can identify areas where the models perform reasonably and where they do not, but in themselves cannot fully evaluate model output. That testing needs to include addressing current and expected economic conditions, for example.

Next steps would include updating and expanding the historical statistical properties of the financial series. A number of the papers referenced are fairly old, and a refresh would be helpful.

REFERENCES

[1] Andersen, Torben G. and Jesper Lund. 1997. Stochastic Volatility and Mean Drift in the Short Rate Diffusion: Sources of Steepness, Level and Curvature in the Yield Curve. Working Paper No. 214, Department of Finance, J.L. Kellogg Graduate School of Management, Northwestern University, http://ssrn.com/abstract=7244.

[2] Andersen, Torben G., Luca Benzoni, and Jesper Lund. 2004. Stochastic volatility, mean drift, and jumps in the short-term interest rate.

[3] Backus, David, Chris Telmer, and Liuren Wu. 1999. Design and Estimation of Affine Yield Models. GSIA Working Papers No. 5, Tepper School of Business, Carnegie Mellon University.

[4] Black, Fischer and Piotr Karisinski.. 1991. Bond and Option Pricing when Short Rates are Lognormal. Financial Analysts Journal, 47(4):52{59, July-Aug.

[5] Brandt, Michael W. and David A. Chapman. 2002. Comparing Multifactor Models of the Term Structure, working paper https://faculty.fuqua.duke.edu/~mbrandt/papers/working/mfcompare.pdf

[6] Cairns, Andrew J.G. 2004. A Family of Term-Structure Models for Long-Term Risk Management and Derivative Pricing. Mathematical Finance, 14:415-444.

[7] Cheridito, Patrick, Damir Filipovic, and Robert L. Kimmel. 2007. Market price of risk specifications for affine models: Theory and evidence. Journal of Financial Economics, 83:123170.

[8] Cochrane, John H. 1996. A Cross-Sectional Test of an Investment-Based Asset Pricing Model, Journal of Political Economy.

[9] Dai, Qiang and Kenneth J. Singleton. 2000. Specification analysis of affine term structure models. Journal of Finance, LV (5):1943{1978.

[10] Duarte, Jefferson. 2004. Evaluating an Alternative Risk Preference in Affine Term Structure Models. Review of Financial Studies, 17(2):379-404.

[11] Duffie, Gregory R. 2002. Term Premia and Interest Rate Forecasts in Affine Models. Journal of Finance, 57:405-443.

[12] Duffie, Darrell and Kenneth J. Singleton. 1993. Simulated Moments Estimation of Markov Models of Asset Prices. Econometrica, 61:929-952.

[13] Duffie, Darrell and Kenneth J. Singleton. 1999. Modeling Term Structures of Defaultable Bonds. Review of Financial Studies, 12 (4):687-720.

[14] Feldhütter, Peter. 2008. Can Affine Models Match the Moments in Bond Yields? April. http://ssrn.com/abstract=915323.

[15] Froot, Kenneth A., and Richard H. Thaler. 1990. "Anomalies: Foreign Exchange." Journal of Economic Perspectives, 4(3): 179-192.

[16] Jagannathan, Ravi, Andrew Kaplin and Steve Sun. 2003. An evaluation of multi-factor CIR models using LIBOR, swap rates, and cap and swaption prices. Journal of Econometrics 116 (2003) 113 – 146.

[17] Kou, S. G. and Hui Wang. 2004. Option Pricing under a Double Exponential Jump Diffusion Model. *Management Science* Vol. 50, No. 9 (Sep.), pp. 1178-1192.

[18] Ramezani and Zeng, 2007. Maximum likelihood estimation of the double exponential jump-diffusion process, Annals of Finance 3:487.

[19] Venter, Gary G. 2004. Testing Distributions of Stochastically Generated Yield Curves. ASTIN Bulletin, 34(1):229-247. www.elsevier.com/locate/econbase