The Market Price of Interest-rate Risk: Measuring and Modelling Fear and Greed in the Fixed-income Markets

Riaz Ahmad, 7city Learning, 4 Chiswell St, London
Paul Wilmott, Empirica Laboratory Ltd, 4 Chiswell St, London, paul@wilmott.com

Abstract:
In this paper we examine the statistical properties of the spot interest rate and the yield curve, using US data, to identify the behaviour of the market price of interest rate risk. This is then also examined statistically so that a two-factor interest rate model is developed.

Keywords:
Spot rate model; Market price of risk; Two-factor model; Market price of market price of risk risk; Fear; Greed

1 Introduction

What should govern the pricing of fixed-income instruments? Forget everything you know about single-factor, two- or three-factor models, Vasicek, CIR, Hull & White, Ho & Lee, Heath, Jarrow & Morton, Brace, Gatarek & Musiela, etc., and suppose you were to address interest rate modelling from first principles. How would you go about the modelling?

The following shows how you might approach this task. Note that we aren't going forget the tools of the trade in terms of stochastic calculus nor the financial principles relating risk and return. The model we end up with is still quite classical in nature, inhabiting the modelling world of stochastic differential equations.

We also aren't going to worry about calibration. We know that is quite a shocking thing to say these days, but if it helps you read the rest of this paper, think of the ideas as inhabiting an 'equilibrium world' rather than a 'no-arbitrage world.' Some of the resulting ideas might find use in trading rather than pricing exotics.

1.1 Warning

What follows combines modelling as it should be done with some statistical data analysis. The model is intended to 'explain,' to some extent, how pricing should work rather than 'describe' what happens.

This approach is doomed.

Explanatory models in finance have historically been far less successful than descriptive or phenomenological models. Famous examples of such failures are equilibrium models generally, Modern Portfolio Theory (less successful in practice than in theory), fundamental analysis ("The market can stay irrational longer than you can stay solvent," said Keynes), and dynamical systems or chaos theory.

The most successful models have answered 'what' rather than 'why.' The simple random walk has been incredibly successful as a model, yet it only describes the behaviour of markets. Few really care why models follow these random walks.

Contrast this situation with models of physical systems such as fluid mechanics. The successful models both explain and describe. The Navier–Stokes and Euler equations explain why fluids behave as they do (being based on conservation of momentum and mass) and also describe what happens (and are therefore used for designing planes).

Well, we are going to attempt to model what *ought to* happen in the fixed-income world, with the expectation that the model will not describe what *does* happen.

2 Implied modelling: One factor

We are going to build up a model for interest rates based on analyzing historical yield-curve data.

- We build up our model in stages and try not to be too sidetracked by 'tractability.'
- We want our model to be simple, accurate in terms of comparison with historical data, and sensible.

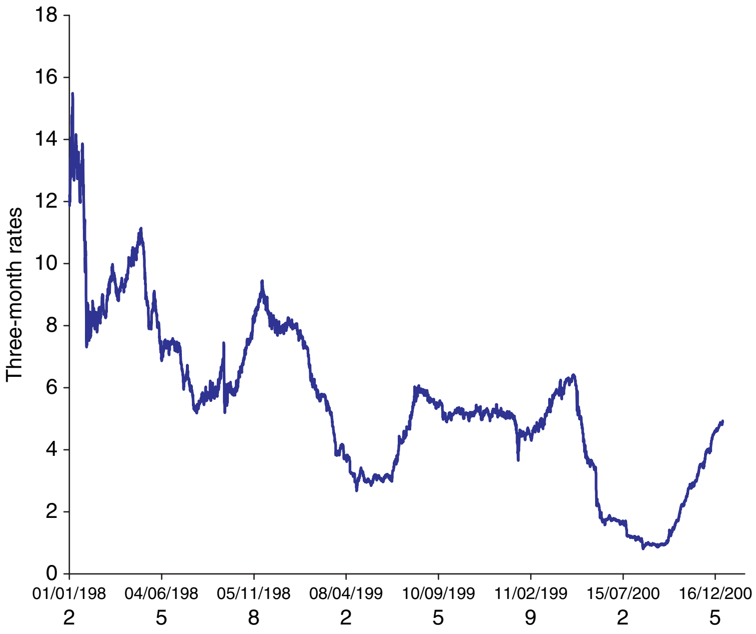

Figure 1: Short-term interest rates 1982–2006.

- We will prefer to choose a model which makes sense over one which calibrates to one day's yield curve.
- Our spot rate model will be

$$dr = u(r)dt + v(r)dX. \tag{1}$$

Note that we are assuming time-independent parameters. There are two reasons for this: This is the most parsimonious assumption; If it were time dependent then there would probably be no sensible way of determining the *future* time dependence. (Time dependence is usually determined by calibration to the yield curve on a single day. As we shall see, this is highly unstable.) Note that Henrotte (2004) said "Departure from time homogeneity may be the sign of serious modelling deficiency.'' Indeed, time *in*homogeneity is often used to make bad models look good via calibration.

In Figure 1 are shown the US three-month interest rates, daily, for the period 1982–2006. This is the data that we use in our analysis. The ideas that we introduce can be applied to any currency, but here we use US data for illustration.

There are three key stages in determining the spot rate model:

1. By differencing spot rate time series data we determine the volatility dependence on the spot rate $v(r)$.
2. By examining the steady-state probability density function for the spot rate we determine the functional form of the drift rate $u(r)$.
3. We examine the slope of the yield curve to determine the market price of risk λ.

This first part of the paper recreates the analysis of Apabhai, Choe, Khennach & Wilmott (1995).

3 The Volatility Structure

Many popular one-factor models take the form

$$dr = u(r)dt + \nu r^{\beta}\, dX$$

where the 'volatility,' the coefficient of the Wiener process, is a power of the spot rate. Examples of such models are the Ho & Lee ($\beta = 0$), Vasicek ($\beta = 0$), Cox, Ingersoll & Ross ($\beta = 1/2$), and Black & Karasinski ($\beta = 1$) models.

We can think of models with

$$dr = \cdots + c\, dX$$

as having a 'normal' volatility structure and those with

$$dr = \cdots + cr\, dX$$

as having a 'lognormal' volatility structure.

We know that getting the volatility structure right is key to pricing derivatives, so which is the correct functional form for $w(r)$? Are any of the popular models accurate?

3.1 The Method

Here we are going to use a very simple, not at all sophisticated, method for determining the volatility structure. From the time-series data divide the changes in the interest rate, δr, into buckets covering a range of r values. Then calculate the average value of $(\delta r)^2$, for each bucket.

If the model (1) is correct we would expect

$$E[(\delta r)^2] = w(r)^2 \delta t$$

to leading order in the time step δt, which for our data is one day.

In anticipation of a power law, $w(r) = \nu r^{\beta}$, being a reasonable model we have plotted $\ln(E[(\delta r)^2])$ against $\ln r$ using the US data in Figure 2. The slope of this 'line' gives an estimate for 2β.

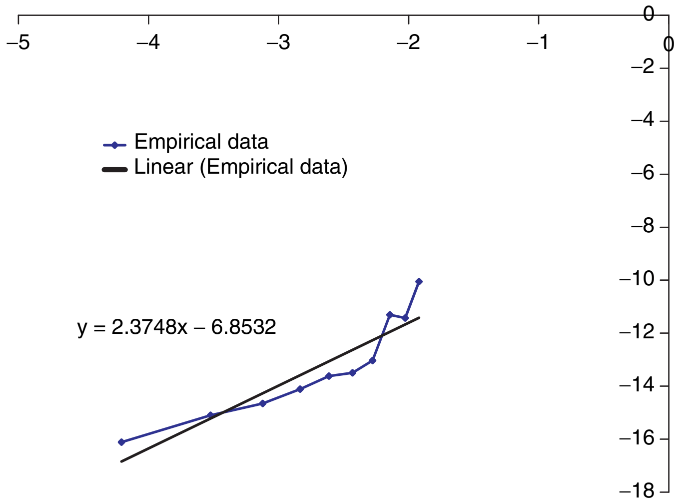

Figure 2: Estimation of β.

From this calculation it is estimated that

$$\beta = 1.2 \text{ and} \nu = 0.52.$$

This confirms that the spot rate volatility is close to lognormal in nature. (Note the curvature in the empirical data. This suggests that for smaller values of r the power should be less than 1.2, but greater than 1.2 for higher r. We will come back to this point later.)

4 The Drift Structure

It is statistically harder to estimate the drift term from the data; this term is smaller than the volatility term and thus subject to larger relative errors. If we were to use a naive method to determine the drift, we may find ourselves with a model that behaves well for short times but behaves poorly in the long term. We will therefore take an alternative, more stable, approach involving the empirical and analytical determination of the steady-state probability density function for r.

4.1 The method

If r satisfies the stochastic differential equation (1) with $w(r) = \nu r^{\beta}$ then the probability density function $p(r, t)$ for r satisfies the forward Fokker–Planck equation

$$\frac{\partial p}{\partial t} = \frac{1}{2}\nu^2 \frac{\partial^2}{\partial r^2}(r^{2\beta} p) - \frac{\partial}{\partial r}(u(r)p). \tag{2}$$

In this r and t are the forward variables, usually denoted by r' and t'.

It is possible for equation (2) to have a steady-state distribution as a solution. This distribution is that to which the probability density function will evolve from any initial condition. We can estimate this steady state from the empirical data and thus find a solution of (2).

This steady state $p_\infty(r)$ will satisfy

$$\frac{1}{2}\nu^2 \frac{d^2}{dr^2}(r^{2\beta} p_\infty) - \frac{d}{dr}(u(r)p_\infty) = 0. \tag{3}$$

If this steady-state probability density function is empirically determined, then by integrating (3) we find that

$$u(r) = \phi^2 \beta r^{2\beta-1} + \frac{1}{2} r^{2\beta} \frac{d}{dr}(\ln p_\infty).$$

Not only is this method for finding the drift more stable in the long run, but also the steady-state probability density function is something simple to focus attention on as opposed to the less meaningful drift function. This probability density function is something that it may be possible to estimate, or at least take an educated guess at. At the same time, it is harder to have an intuitive feel for the drift coefficient $u(r)$.

We will also know that our model cannot behave too outrageously, that the distribution of the model and of the real data will match. By choosing a model with a sensible steady-state distribution, we can guarantee that the model will not allow the spot rate to do anything unrealistic such as grow unboundedly.

Again, looking at US data, we can determine a plausible functional form for $p_\infty(r)$ from three-month US rates, see Figure 3. The steady-state distribution is determined by dividing r into buckets and observing the frequency with which each bucket is reached.

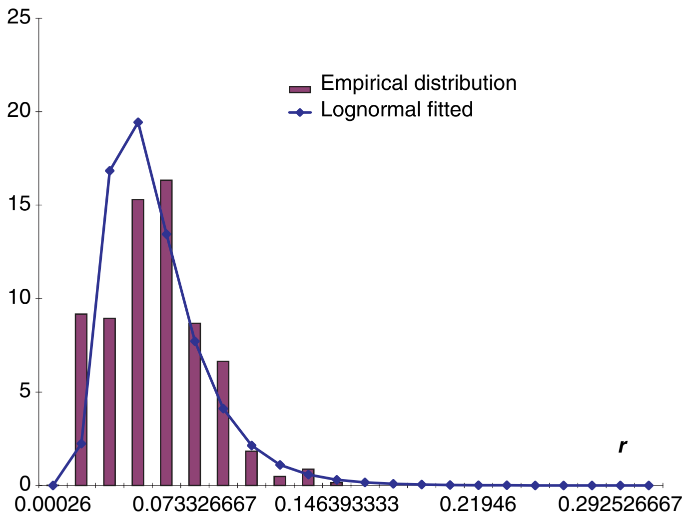

Figure 3: Steady-state empirical and fitted distributions.

The shape of this graph is reminiscent of a lognormal curve. For this reason, and because it has a simple formula with just two parameters, we choose $p_\infty(r)$ to be a lognormal curve that best fits the empirical data; this curve is also shown in the figure.

Our choice for $p_\infty(r)$ is

$$\frac{1}{ar\sqrt{2\pi}} \exp\left(-\frac{1}{2a^2}(\ln(r/\overline{r}))^2\right)$$

where $a = 0.46$ and $\overline{r} = 0.048$. From this we find that for the US market

$$u(r) = \nu^2 r^{2\beta-1}\left(\beta - \frac{1}{2} - \frac{1}{2a^2}\ln(r/\overline{r})\right).$$

The real spot rate is therefore mean-reverting to 5.3%.

5 But There's More

Modelling the spot rate is not enough for pricing interest-rate derivatives. Unlike equities where a model for the 'underlying' leads to the Black–Scholes equation, fixed income has a twist. Because the spot rate is not traded it is not possible to eliminate interest-rate risk by dynamic hedging. Contrast this with equity derivatives for which it is theoretically possible to eliminate market risk by delta hedging. If we cannot eliminate risk then we must know how to price it. This amounts to modelling how much extra expected return is required for a 'unit' amount of interest-rate risk. Once this is specified then we use this same 'market price of risk' to price all fixed-income contracts in a consistent way. See Wilmott (2006) for details.

We will denote this market price of interest-rate risk by λ.

With our single-factor spot-rate model, and with a *deterministic* λ, the pricing equation for non-path dependent contracts is

$$\frac{\partial V}{\partial t} + \frac{1}{2}w^2 \frac{\partial^2 V}{\partial r^2} + (u - \lambda w)\frac{\partial V}{\partial r} - rV = 0. \tag{4}$$

6 The Slope of the Yield Curve and the Market Price of Risk

Now we have found $w(r)$ and $u(r)$, it only remains for us to find λ. The model will then be complete.

Because there is no information about the market price of risk in the spot-rate process, we must look to interest-rate derivatives to see how the market is valuing risk. The simplest way to do this is to look at the yield curve, because in the present context even bonds can be thought of as derivatives of the spot rate.

In particular, the short end of the yield curve contains easily accessed information about the market price of risk.

Let's look at Equation (4) as applied to a zero-coupon bond with maturity $t = T$. This will satisfy the equation together with the final condition $V(r, T) = 1$, the 1 representing the payment of the principal.

We can expand $V(r, t)$ in a Taylor series about $t = T$, this is the short end of the yield curve, to find that

$$V(r, t) \sim 1 - r(T - t) + \frac{1}{2}(T - t)^2(r^2 - u + \lambda w) + \ldots \text{ as } t \to T$$

for any model $w(r)$, $u(r)$ and λ. From this we have

$$-\frac{\ln V}{T - t} \sim r + \frac{1}{2}(u - \lambda w)(T - t) + \ldots \text{ as } t \to T. \quad (5)$$

Equation (5) says that the slope of the yield curve at the short end in this one-factor model, with deterministic λ, is simply $(u - \lambda w)/2$. We can use this result together with time-series data to determine the λ empirically.

Knowing r each day, and the slope of the yield curve each day, we can back out a time series for λ. This is shown in Figure 4. This figure shows λ as a function of time.

Note that on balance λ is negative. This is correct, technically this is because the 'risk' in a bond associated with the spot rate is proportional to the sensitivity of the bond price to the spot rate, and $\frac{\partial V}{\partial r} < 0$ for zero-coupon bonds.

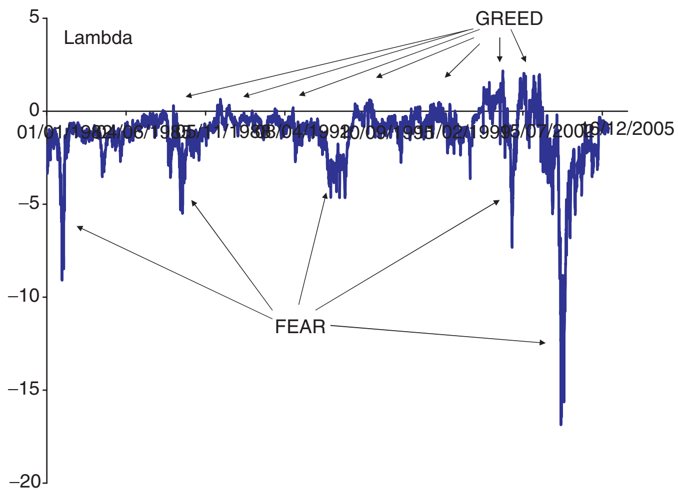

Figure 4: λ time series.

But this figure shows many interesting features.

- λ is not constant (or even a deterministic function of r, for a plot of λ versus r would show little structure)
- λ is not always negative
- λ varies wildly from day to day

7 Observations on the Market Price of Risk

We can make several observations concerning the results shown in the previous figure.

First, the average level of the market price of risk (MPR), λ, is about -1.2. This *average* is relatively stable through time. More recently the MPR has become larger in magnitude. This may not actually be the case, but instead may be because the function for the spot rate volatility, $v(r)$ is just too inaccurate for the recent low levels of r. As hinted at earlier a lower value for β for the smaller values of r would give a better fit for the volatility function *and* a better-behaved MPR.

The positive peaks in the figure mean that people are willing to pay to take risk. These are accordingly labeled 'Greed' in the plot. The negative troughs are when investors need extra return for taking risk. These are labeled 'Fear.'

The MPR does appear to be mean reverting over a short timescale (at least short compared with the period analyzed). Therefore it suggests possible statistical arbitrage strategies based on exploiting the fear/greed factor of the market. These would amount to plays on the slope of the yield curve at the short end.

8 Why should the MPR not be Constant?

- Investors are not always rational

If investors were always rational you might expect the λ plot to be flat. Or perhaps it would be slowly varying in time representing the varying attitude towards risk of different generations. Or perhaps it would be decreasing and piecewise constant representing investors sudden appreciation of previously unknown risks. All sorts of shapes for this plot could be hypothesized based upon rationality, but 'random' as seen here is unlikely to be one of them. You would not expect the MPR to rise and fall in the way it seems to, unless investors have very short memories.

Of course, there is always 'liquidity,' or rather lack of it, to explain short-term anomalies in the MPR.

Further, why would the MPR ever be of the 'wrong' sign? Why would investors pay to take risk? There are several possible explanations for this. First, they will only know that they are paying to take risk if they have performed the above or similar analysis. Or they could well be taking a view on the market, perhaps anticipating interest rate jumps, which may or may not turn out to be correct. They may be using fixed-income instruments for hedging purposes, but whether this makes sense depends on their whole portfolio and correlations. However, there are times when investors do consciously enter a negative-expectation game. This is seen whenever anyone buys a lottery ticket or plays roulette. But in the 'rational' world of investment banking would you expect to see this behaviour? It depends on your view of what is rational. If a hedge fund is sitting on a billion dollars of investors' money yet cannot find any trading opportunities they may find themselves with the decision to either invest and lose a little in a negative-expectation game, or sit on the money, the rational thing to do, and watch investors redeem their money that is not being utilized. Why would investors pay a 2% management fee to have a hedge fund sit on their money? Of course, to some extent this can be mitigated by hedge funds only charging management fee on the money that they utilize, so slightly reducing the incentive for investors to redeem.

Now, of course, it could be that the spot rate model itself is wrong, and the correct model would show constant MPR. This seems like extreme wishful thinking. The model could no doubt be easily improved, and rational arguments made for the behaviour of the MPR. However, having identified a mechanism for estimating the MPR and having any understanding of human behaviour the conclusion that investors are not always rational is clearly the correct and obvious one. Whatever rational explanations can be made, there will always be an irrational element left, and fine tuning of the model could identify this.

Finally, given the data, the MPR is clearly a prime candidate for modelling as a stochastic quantity. So we shall shortly seek a two-factor stochastic spot rate and stochastic market price of risk model.

9 Philosophy

Market price of risk always appears in pricing equations when you model something that isn't traded (such as spot interest rate or volatility). It measures how much extra return you need for taking unhedgeable risk.

You can see the market price of interest rate risk by looking at the slope of the yield curve. It changes a lot, appearing to be random. Sometimes very high, sometimes very low. Sometimes the market is fearful, sometimes greedy.

One of the consequences of this is that when you calibrate a model you are calibrating to a snapshot of the market and its price of risk at an instant. If that market is extremely fearful or greedy then you are making an assumption when you calibrate about how the market's fear/greed will evolve.

This is just a simple way of thinking of what calibration means, and why it shouldn't work. Just because everyone else is sticking their finger in the fire now doesn't mean that you should or that they will do so in the future.

If markets were rational and the market price of risk were always constant then calibration would be fine. But they are not, and it isn't.

10 Analyzing the market price of risk

In the following stochastic market price of risk model we have a spot interest rate which is random, and fear/greed, as measured by the market price of risk, is also random.

The analysis of the market price of risk follows the same first two steps as with the analysis of the spot interest rate.

1. Find the volatility of the MPR process
2. Find the drift of the MPR process

11 The MPR Process

Let's write

$$d\lambda = p(\lambda)\, dt + q(\lambda) dX_2.$$

Note that this is a very simple model. It does not, for example, include any r dependence in the coefficients. Just like our r model does not contain any λs in the u or w functions. This is probably not realistic but for the moment has the advantage of being simple to determine from the data. One could be more general in the modelling at the expense of more complicated historical analysis. For now, since we are as much presenting a philosophy and a warning as a model, the simple assumptions will suffice.

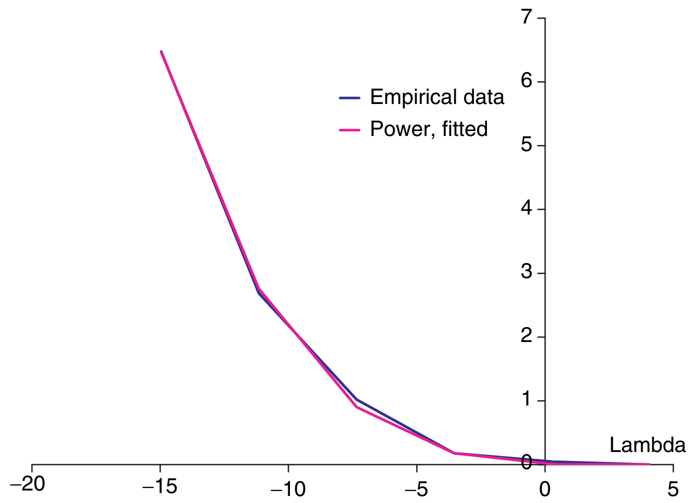

Figure 5: Estimation of volatility of λ.

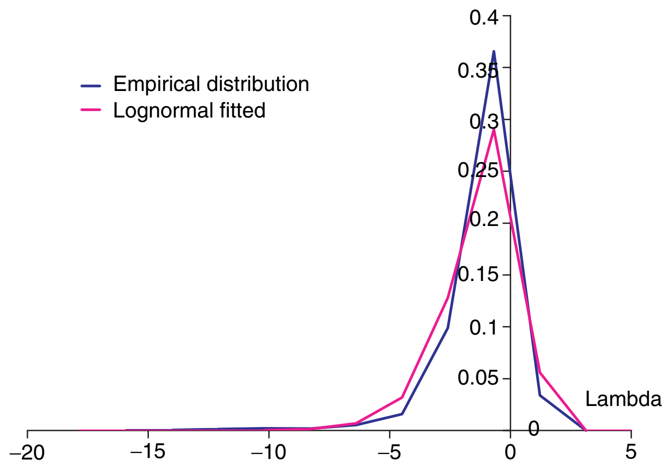

Figure 6: Distribution of λ.

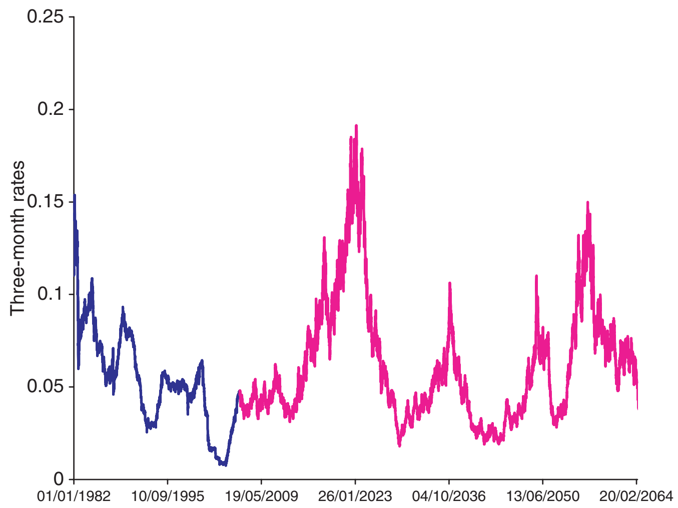

Figure 7: Simulation of *r*.

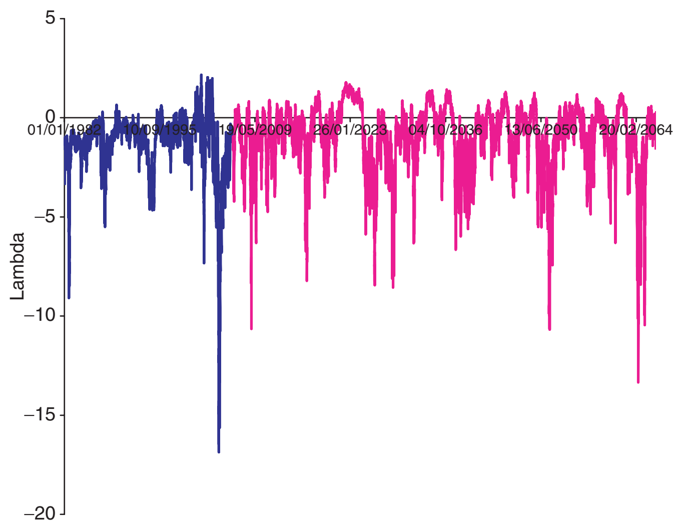

Figure 8: Simulation of λ.

The details can be skipped, and we now present the results. Refer to Figures 5 and 6 for the fitting of the q and p functions. (Note that the Taylor series leading to (5) is still valid to the necessary order even if λ is stochastic.)

From the data we conclude that

$$q(\lambda) = l\,(\lambda_{\max} - \lambda)^m,$$

with $l = 0.21$, $\lambda_{\max} = 3.2$ and $m = 1.8$.

From the data we also conclude that

$$\text{PDF} = \frac{1}{c(\lambda_{\max} - \lambda)\sqrt{2\pi}} \exp\left(-\frac{1}{2c^2}(\ln((\lambda_{\max} - \lambda)/\bar{\lambda}))^2\right)$$

where $c = 0.35$, $\bar{\lambda} = 4.1$ and $\lambda_{\max} = 3.2$. From this we find that

$$p(\lambda) = l^2(\lambda_{\max} - \lambda)^{2m-1}\left(-m + \frac{1}{2} + \frac{1}{2c^2}\ln((\lambda_{\max} - \lambda)/\bar{\lambda})\right).$$

The correlation between dX_1, for the spot interest rate, and dX_2, for the MPR, is, from the data and assuming it to be a simple constant, 0.54.

In Figures 7 and 8 are shown simulations of r and λ using this model.

12 The Market Price of Risk

Of course, we can't find the whole yield curve or price derivatives unless we know how the market values risk in this new non-traded parameter, the MPR.

Technically this means we need to know *the market price of market price of risk risk*!

Possible choices for this parameter are

1. Zero
2. A constant
3. A function of time
4. λ

Justifications and thoughts concerning these choices are as follows.

1. **Zero:** This risk is too obscure to be priced in at all. This is appealing because it is very simple. However, there is no theoretical reason for this.
2. **A constant:** The market is totally rational about the MPR, and prices it consistently from day to day. Again, if the world were a rational place then this would be fine. Experience suggests otherwise.
3. **A function of time:** A function to be chosen so that the yield curve is fitted. This is standard calibration. Cheating. This is supposed to model what prices *should be* given the two fundamental variables spot rate and price of risk. Calibration would be a violation of our basic philosophy.
4. **λ:** All market prices of risk ought to be the same. The most elegant choice. After all, why should different types of risk have different prices, risk is risk is risk (at least when measured in monetary terms like here). This is our prefered choice.

Our final model can now be written in risk-neutral terms as

$$dr = (u(r) - \lambda v(r))dt + v(r)dX_1$$

and

$$d\lambda = (p(r) + \lambda q(r))dt + q(r)dX_2,$$

with a correlation of ρ. (Note the + sign in front of the $\lambda q(r)$.This is because $\frac{\partial V}{\partial \lambda} > 0$.)

We have determined the parameter functions that give sensible models for r and λ, such that they match market data. The pricing equation is

$$\frac{\partial V}{\partial t} + \frac{1}{2}v^2\frac{\partial^2 V}{\partial r^2} + pvq\frac{\partial^2 V}{\partial r \partial \lambda} + \frac{1}{2}q^2\frac{\partial^2 V}{\partial \lambda^2}$$
$$+ (u - \lambda v)\frac{\partial V}{\partial r} + (p + \lambda q)\frac{\partial V}{\partial \lambda} - rV = 0.$$

13 Final thoughts

The above model focuses on the two financial quantities that *ought to* be modelled in the fixed-income world, the spot rate and the market price of interest-rate risk. Although the principle of this approach is sound, the details can be criticized on several grounds. So let us be the first to point some of these out.

Clearly, the last few years have been special with such a sustained low level of interest rates. In our analysis this coincides with larger (negative) λ. Either the market price of risk really has grown more negative, representing people being more afraid of risk, or more likely, our model for $w(r)$ needs fine tuning when r is small.

One of the consequences of the above model is that long-term rates will be constant, representing in some sense the average behaviour of r and λ, and since the parameters are independent of time. Changing long-term rates could either be interpreted as being simply wrong, another sign of the irrationality of the market, or of changing parameters or a regime change.

Because of the role of intervention in interest rates, the dynamics of r may not capture the 'jumpiness' seen in practice. This could be represented by a jump-diffusion model. Note that trending due to intervention is different from serial autocorrelation. The data we used for US three-month rates showed very little serial autocorrelation.

The model should be applied to different economies, to see if there are any similarities between magnitude of the MPR and its dynamics from one country to another. Is there perhaps any relationship between the MPR and rate of inflation? And what is significant about the dates on which λ took its more extreme values?

The market price of risk unifies all investments, it levels the playing field so that pure equity, options (with stochastic volatility), fixed income, credit, etc. can all be compared and analyzed. After all, sensible investing is about getting a good return after allowing for risk taken, and this exactly what the market price of risk measures.

Acknowledgments

PW would like to thank Pat Hagan, Fouad Khennach, the delegates at a BRIT training course in Sandton, August 2006, the delegates on a training course given with Nassim Nicholas Taleb and students on the Certificate in Quantitative Finance, for their helpful feedback.

REFERENCES

■ Apabhai, MZ, Choe, K, Khennach, F & Wilmott, P 1995 Spot-on modelling. *Risk magazine*, December **8** (11) 59–63
■ Henrotte, P 2004 The case for time homogeneity. *Wilmott* magazine, January pp 71–5
■ Wilmott, P 2006 *Paul Wilmott On Quantitative Finance, second edition*. John Wiley & Sons.