Yield Curve Risk Factors: Domestic And Global Contexts

Wesley Phoa

Quantitative Research
Capital Strategy Research
The Capital Group of Companies
11100 Santa Monica Boulevard
Los Angeles, CA 90025

phone (310) 996 6308
fax (310) 996 6022
email wkp@capgroup.com

To appear in

The Practitioner's Handbook of Financial Risk Management

CONTENTS

INTRODUCTION: HANDLING MULTIPLE RISK FACTORS

1. Methodological introduction

Traditional interest rate risk management focuses on duration and duration management. In other words, it assumes that only parallel yield curve shifts are important. In practice, of course, non-parallel shifts in the yield curve often occur, and represent a significant source of risk. What is the most efficient way to manage non-parallel interest rate risk?

This chapter is mainly devoted to an exposition of *principal component analysis*, a statistical technique that attempts to provide a foundation for measuring non-parallel yield curve risk, by identifying the "most important" kinds of yield curve shift that empirically occur. The analysis turns out to be remarkably successful. It gives a clear justification for the use of duration as the primary measure of interest rate risk, and it also suggests how one may design "optimal" measures of non-parallel risk.

Principal component analysis is a popular tool, not only in theoretical studies but in practical risk management applications. We discuss such applications at the end of the chapter. However, it is first important to understand that principal component analysis has limitations, and should not be applied blindly. In particular, it is important to distinguish between results that are economically meaningful, and results that are statistical artifacts without economic significance.

There are two ways to determine whether the results of a statistical analysis are meaningful. The first is to see whether they are consistent with theoretical results; the Appendix gives a sketch of this approach. The second is simply to carry out as much exploratory data analysis as possible, with different data sets and different historical time periods, to screen out those findings which are really robust. This chapter contains many examples

In presenting the results, our exposition will rely mainly on graphs rather than tables and statistics. This is not because rigorous statistical criteria are unnecessary – in fact, they are very important. However, in the exploratory phase of any empirical study it is critical to get a good feel for the results first, since statistics can easily mislead. The initial goal is to gain insight; and visual presentation of the results can convey the important findings most clearly, in a non-technical form.

It is strongly suggested that, after finishing this chapter, readers should experiment with the data themselves. Extensive hands-on experience is the only way to avoid the pitfalls inherent in any empirical analysis.

2. Non-parallel risk, duration bucketing and partial durations

Before discussing principal component analysis, we briefly review some more primitive approaches to measuring non-parallel risk. These have by no means been superseded: later on we will discuss precisely what role they continue to play in risk management.

The easiest approach is to *group securities into maturity buckets*. This is a very simple way of estimating exposure to movements at the short, medium and long ends of the yield curve. But it is not very accurate: for example, it ignores the fact that a bond with a higher coupon intuitively has more exposure to movements in the short end of the curve than a lower coupon bond with the same maturity.

Next, one could *group securities into duration buckets.* This approach is somewhat more accurate because, for example, it distinguishes properly between bonds with different coupons. But it is still not entirely accurate because it does not recognize that the different individual cashflows of a single security are affected in different ways by a non-parallel yield curve shift.

Next, one could *group security cashflows into duration buckets*. That is, one uses a finer-grained unit of analysis: the cashflow, rather than the security. This makes the results much more precise. However, bucketed duration exposures have no direct interpretation in terms of changes in some reference set of yields (i.e. a shift in some reference yield curve), and can thus be tricky to interpret. More seriously, as individual cashflows shorten they will move across bucket duration boundaries, causing discontinuous changes in bucket exposures which can make risk management awkward.

Alternatively, one could *measure partial durations*. That is, one directly measures how the value of a portfolio changes when a single reference yield is shifted, leaving the other reference yields unchanged; note that doing this at the security level and at the cashflow level gives the same results. There are many different ways to define partial durations: one can use different varieties of reference yield (e.g. par, zero coupon, forward rate), one can choose different sets of reference maturities, one can specify the size of the perturbation, and one can adopt different methods of interpolating the perturbed yield curve between the reference maturities.

The most popular partial durations are the *key rate durations* defined in [Ho]. Fixing a set of reference maturities, these are defined as follows: for a given reference maturity T, the T-year key rate duration of a portfolio is the percentage change in its value when one shifts the T-year zero coupon yield by 100 bp, leaving the other reference zero coupon yields fixed, and linearly interpolating the perturbed zero coupon curve between adjacent reference maturities (often referred to as a "tent" shift). Exhibit 1 shows some examples of key rate durations.

All of the above approaches must be used with caution when dealing with option-embedded securities such as callable bonds or mortgage pools, whose cashflow timing will vary with the level of interest rates. Option-embedded bonds are discussed in detail elsewhere in this volume.

3. Limitations of key rate duration analysis

Key rate durations are a popular and powerful tool for managing non-parallel risk, so it is important to understand their shortcomings.

First, key rate durations can be unintuitive. This is partly because "tent" shifts do not occur in isolation, and in fact have no economic meaning in themselves. Thus, using key rate durations requires some experience and familiarization.

Second, correlations between shifts at different reference maturities are ignored. That is, the analysis treats shifts at different points in the yield curve as independent, whereas different yield curve points tend to move in correlated ways. It is clearly important to take these correlations into account when measuring risk, but the key rate duration methodology does not suggest a way to do so.

Third, the key rate duration computation is based on perturbing a theoretical zero coupon curve rather than observed yields on coupon bonds, and is therefore sensitive to the precise method used to strip (e.g.) a par yield curve. This introduces some arbitrariness into the results, and more significantly makes them hard to interpret in terms of observed yield curve shifts. Thus swap dealers (for example) often look at partial durations computed by directly perturbing the swap curve (a par curve) rather than perturbing a zero coupon curve.

Fourth, key rate durations for mortgage-backed securities must be interpreted with special care. Key rate durations closely associated with specific reference maturities which drive the prepayment model can appear anomalous; for example, if the mortgage refinancing rate is estimated using a projected 10-year Treasury yield, 10-year key rate durations on MBS will frequently be negative. This is correct according to the definition, but in this situation one must be careful constructing MBS hedging strategies using key rate durations.

Fifth, key rate durations are unwieldy. There are too many separate interest rate risk measures. This leads to practical difficulties in monitoring risk, and inefficiencies in hedging risk. One would rather focus mainly on what is "most important".

To summarize: while key rate durations are a powerful risk management tool, it is worth looking for a more sophisticated approach to analyzing non-parallel risk that will yield deeper insights, and that will provide a basis for more efficient risk management methodologies.

Acknowledgements: The research reported here was carried out while the author was employed at Capital Management Sciences. The author has attempted to incorporate several useful suggestions provided by an anonymous reviewer.

I. PRINCIPAL COMPONENT ANALYSIS

1. Definition and examples from US Treasury market

As often occurs in finance, an analogy with physical systems suggests an approach. Observed shifts in the yield curve may seem complex and somewhat chaotic. In principle, it might seem that any point on the yield curve can move independently in a random fashion. However, it turns out that most of the observed fluctuation in yields can be explained by more systematic yield shifts: that is, bond yields moving 'together', in a correlated fashion, but perhaps in several different ways. Thus, one should not focus on fluctuations at individual points on the yield curve, but on shifts that apply to the yield curve as a whole. It is possible to identify these systematic shifts by an appropriate statistical analysis; as often occurs in finance, one can apply techniques inspired by the study of physical systems.

The following concrete example, taken from [Jennings & McKeown], may be helpful. Consider a plank with one end fixed to a wall. Whenever the plank is knocked, it will vibrate. Furthermore, when it vibrates it does not deform in a completely random way, but has only a few "vibration modes" corresponding to its natural frequencies. These vibration modes have different degrees of importance, with one mode – a simple back-and-forth motion – dominating the others: see Exhibit 2a.

One can derive these vibration modes mathematically, if one knows the precise physical characteristics of the plank. But one should also be able to determine them empirically by observing the plank. To do this, one attaches motion sensors at different points on the plank, to track the motion of these points through time. One will find that the observed disturbances at each point are correlated. It is possible to extract the vibration modes, and their relative importance, from the correlation matrix. In fact, the vibration modes correspond to the eigenvalues of the matrix: in other words, the eigenvectors, plotted in graphical form, will turn out to look exactly as in Exhibit 2a. The relative importance of each vibration mode is measured by the size of the corresponding eigenvectors.

Let us recall the definitions. Let **A** be a matrix. We say that **v** is an *eigenvector* of **A**, with corresponding *eigenvalue* l, if $\mathbf{A} \cdot \mathbf{v} = l\mathbf{v}$. The eigenvalues of a matrix must be mutually orthogonal, i.e. "independent". Note that eigenvectors are only defined up to a scalar multiple, but that eigenvalues are uniquely defined.

Suppose **A** is a correlation matrix, e.g. derived from some time series of data; then it must be *symmetric* and also *positive definite* (i.e. $\mathbf{v} \cdot \mathbf{A} \cdot \mathbf{v} > 0$ for all vectors **v**). One can show that all the eigenvalues of such a matrix must be real and positive. In this case it makes sense to compare their relative sizes, and to regard them as "weights" which measure the importance of the corresponding eigenvectors.

For a physical system such as the cantilever, the interpretation is as follows. The eigenvectors describe the independent vibration modes: each eigenvector has one component for each sensor, and the component is a (positive or negative) real number which describes the relative displacement of that sensor under the given vibration mode. The corresponding eigenvalue measures how much of the observed motion of the plank can be attributed to that specific vibration mode.

This suggests that we can analyze yield curve shifts analogously, as follows. Fix a set of reference maturities for which reasonably long time series of, say, daily yields are available: each reference maturity on the yield curve is the analog of a motion sensor on the plank. Construct the time series of daily changes in yield at each reference maturity, and compute the correlation matrix. Next, compute the eigenvectors and eigenvalues of this matrix. The eigenvectors can then be interpreted as independent "fundamental yield curve shifts", analogous to vibration modes; in other words, the actual change in the yield curve on any particular day may be regarded as a combination of different, independent, fundamental yield curve shifts. The relative sizes of the eigenvalues tells us which fundamental yield curve shifts tend to dominate.

For a toy example, see Exhibit 2b. The imaginary data set consists of five days of observed daily yield changes at four unnamed reference maturities; for example, on days 1 and 3 a perfectly parallel shift occurred. The correlation matrix shows that yield shifts at different maturity points are quite correlated. Inspecting the eigenvalues and eigenvectors shows that, at least according to principal component analysis, there is a dominant yield curve shift, eigenvector [D], which represents an almost parallel shift: each maturity point moves by about 0.5. The second most important eigenvector [C] seems to represent a slope shift or "yield curve tilt". The third eigenvector [B] seems to appear because of the inclusion of day 5 in the data set.

Note that the results might not perfectly reflect one's intuition. First, the dominant shift [D] is not perfectly parallel, even though two perfectly parallel shifts were included in the data set. Second, the shift that occurred on day 2 is regarded as a combination of a parallel shift [D] and a slope shift [C], not a slope shift alone; shift [C] has almost the same shape as the observed shift on day 2, but it has been "translated" so that shifts of type [C] are uncorrelated with shifts of type [D]. Third, eigenvector [A] seems to have no interpretation. Finally, the weight attached to [D] seems very high – this is because the actual shifts on all five days are regarded as having a parallel component, as we just noted.

A technical point: In theory, one could use the covariance matrix rather than the correlation matrix in the analysis. However using the correlation matrix is preferable when observed correlations are more stable than observed covariances – which is usually the case in financial data where volatilities are quite unstable. (For further discussion, see [Buhler and Zimmermann].) In the example of Exhibit 2b, very similar results are obtained using the covariance matrix.

Exhibit 3 shows the result of a principal component analysis carried out on actual US Treasury bond yield data from 1993-98. In this case the dominant shift is a virtually parallel shift, which explains over 90% of observed fluctuations in bond yields. The second most important shift is a slope shift or tilt in which short yields fall and long yields rise (or vice versa). The third shift is a kind of curvature shift, in which short and long yields rise while mid-range yields fall (or vice versa); the remaining eigenvectors have no meaningful interpretation and are statistically insignificant.

Note that meaningful results will only be obtained if a consistent set of yields is used: in this case, constant maturity Treasury yields regarded as a proxy for a Treasury par yield curve. Yields on physical bonds should not be used, since the population of bonds both ages and changes composition over time. The analysis here has been carried out using CMT yields reported by the US Federal Reserve Bank.

An alternative is to use a dataset consisting of historical swap rates, which are par yields by definition. The results of the analysis turn out to be very similar.

2. Meaningfulness of factors: dependence on dataset

It is extremely tempting to conclude that (a) the analysis has determined that there are exactly three important kinds of yield curve shift, that (b) that it has identified them precisely, and that (c) it has precisely quantified their relative importance.

But we should not draw these conclusions without looking more carefully at the data. This means exploring data sets drawn from different historical time periods, from different sets of maturities, and from different countries. *Risk management should only rely on those results which turn out to be robust.*

Exhibit 4 shows a positive finding. Analyzing other 5-year historical periods, going back to 1963, we see that the overall results are quite consistent. In each case the major yield curve shifts turn out to be parallel, slope and curvature shifts; and estimates of the relative importance of each kind of shift are reasonably stable over time, although parallel shifts appear to have become more dominant since the late 1970s.

Exhibits 5a and 5b show that some of the results remain consistent when examined in more detail: the estimated form of both the parallel shift and the slope shift are very similar in different historical periods. Note that in illustrating each kind of yield curve shift, we have carried out some normalization to make comparisons easier: for example, estimated slope shifts are normalized so that the 10-year yield moves 100 bp relative to the 1-year yield, which remains fixed. See below for further discussion of this point.

However, Exhibit 5c does tentatively indicate that the form of the curvature shift has varied over time – a first piece of evidence that results on the curvature shift may be less robust than those on the parallel and slope shifts.

Exhibit 6 shows the effect of including 3-month and 6-month Treasury bill yields in the 1993-98 data set. The major yield curve shifts are still identified as parallel, slope and curvature shifts. However, an analysis based on the data set including T-bills attaches somewhat less importance to parallel shifts, and somewhat more importance to slope and curvature shifts. Thus, while the estimates of relative importance remain qualitatively significant, they should not be regarded as quantitatively precise.

Exhibits 7a and 7b show that the inclusion of T-bill yields in the data set makes almost no difference to the estimated form of both the parallel and slope shifts. However,

Exhibit 7c shows that the form of the curvature shift is totally different. Omitting T-bills, the change in curvature occurs at the 3-5 year part of the curve; including T-bills, it occurs at the 1-year part of the curve. There seem to be some additional dynamics associated with yields on short term instruments, which become clear once parallel and slope shifts are factored out; this matter is discussed further in [Phoa].

The overall conclusions are that parallel and slope shifts are unambiguously the most important kinds of yield curve shift that occur, with parallel shifts being dominant; that the forms of these parallel and slope shifts can be estimated fairly precisely and quite robustly; but that the existence and form of a third, "curvature" shift are more problematic, with the results being very dependent on the data set used in the analysis. Since the very form of a curvature shift is uncertain, and specifying it precisely requires making a subjective judgment about which dataset is "most relevant", the curvature shift is of more limited use in risk management. However, we do give an example of its use in Example 15.

The low weight attached to the curvature factor also suggests that it may be less important than other (conjectural) phenomena which might somehow have been missed by the analysis. The possibility that the analysis has failed to detect some important yield curve risk factors, which potentially outweigh curvature risk, is discussed further below.

International bond yield data are analyzed in the next section. The results are broadly consistent, but also provide further grounds for caution. The appendix provides some theoretical corroboration for the positive findings.

We have glossed over one slightly awkward point. The fundamental yield curve shifts estimated by a principal component analysis – in particular, the first two principal components representing parallel and slope shifts – are, by definition, uncorrelated. But normalizing a "slope shift" so that the 1-year yield remains fixed introduces a possible correlation. This kind of normalization is convenient both for data analysis, as above, and for practical applications; but it does mean that one then has to estimate the correlation between parallel shifts and normalized slope shifts. This is not difficult in principle, but, as show in [Phoa], this correlation is time-varying and indeed exhibits secular drift. This corresponds to the fact that, while the estimated (non-normalized) slope shifts for different historical periods have almost identical shapes, they have different "pivot points". The issue of correlation risk is discussed further below.

3. Correlation structure and other limitations of the approach

It is now tempting to concentrate entirely on parallel and slope shifts. This approach forms the basis of most useful two factor interest rate models: see [Brown and Schaefer]. However, it is important to understand what is being lost when one focuses only on two kinds of yield curve shift.

First, there is the question of *whether empirical correlations are respected*. Exhibit 8a shows, graphically, the empirical correlations between daily Treasury yield shifts at different maturity points. It shows that, as one moves to adjacent maturities, the correlations fall away rather sharply. In other words, even adjacent yields quite often shift in uncorrelated ways.

Exhibit 8b shows the correlations which would have been observed if only parallel and slope shifts had taken place. These slope away much more gently as one moves to adjacent maturities: uncorrelated shifts in adjacent yields do not occur. This observation is due to [Rebonato and Cooper], who prove that the correlation structure implied by a two factor model must always take this form.

What this shows is that, even though the weights attached to the "other" eigenvectors seemed very small, discarding these other eigenvectors radically changes the correlation structure. Whether or not this matters in practice will depend on the specific application.

Second, there is the related question of *the time horizon of risk*. Unexplained yield shifts at specific maturities may be unimportant if they quickly "correct"; but this will clearly depend on the investor's time horizon. If some idiosyncratic yield shift occurs, which has not been anticipated by one's risk methodology, this may be disastrous for a hedge fund running a highly leveraged trading book with a time horizon of hours or days; but an investment manager with a time horizon of months or quarters, who is confident that the phenomenon is transitory and who can afford to wait for it to reverse itself, might not care as much.

This is illustrated in Exhibit 9. It compared the observed 10-year Treasury yield from 1953-96 to the yield which would have been predicted by a model in which parallel and slope risk fully determine (via arbitrage pricing theory) the yields of all Treasury bonds. The actual yield often deviates significantly from the theoretical yield, as yield changes unrelated to parallel and slope shifts frequently occurred. But deviations appear to mean revert to zero over periods of around a few months to a year; this can be justified more rigorously by an analysis of autocorrelations. Thus, these deviations matter over short time frames, but perhaps not over long time frames. See [Phoa] for further details.

Third, there is the question of *effects due to market inhomogeneity*. In identifying patterns of yield shifts by maturity, principal component analysis implicitly assumes that the only relevant difference between different reference yields is maturity, and that the market is homogeneous in every other way. If it is not – for example, if there are differences in liquidity between different instruments which, in some circumstances, lead to fluctuations in relative yields – then this assumption may not be sound.

The US Treasury market in 1998 provided a very vivid example. Yields of on-the-run Treasuries exhibited sharp fluctuations relative to off-the-run yields, with "liquidity spreads" varying from 5 bp to 25 bp. Furthermore, different on-the-run issues were affected in different ways in different times. A principal component analysis based on constant maturity Treasury yields would have missed this source of risk entirely; and in fact, even given yield data on the entire population of Treasury bonds, it would have been extremely difficult to design a similar analysis which would have been capable of identifying and measuring some systematic "liquidity spread shift". In this case, risk management for a Treasury book based on principal component analysis needs to be supplemented with other methods.

Fourth, there is the possibility that *an important risk factor has been ignored*. For example, suppose there is an additional kind of fundamental yield curve shift, in which 30- to 100-year bond yields move relative to shorter bond yields. This would not be identified by a principal component analysis, for the simple reason that this maturity range is represented by only one point in the set of reference maturities. Even if the 30-year yield displayed idiosyncratic movements – which it arguably does – the analysis would not identify these as statistically significant. The conjectured "long end" risk factor would only emerge if data on other longer maturities were included; but no such data exists for Treasury bonds.

An additional kind of "yield curve risk", which could not be detected at all by an analysis of CMT yields, is the varying yield spread between liquid and illiquid issues. This was a major factor in the US Treasury market in 1998; in fact, from an empirical point of view, fluctuations at the long end of the curve and fluctuations in the spread between on-the-run and off-the-run Treasuries were, in that market, more important sources of risk than curvature shifts – and different methods were required to measure and control the risk arising from these sources.

To summarize, a great deal more care is required when using principal component analysis in a financial, rather than physical, setting. One should always remember that the rigorous justifications provided by the differential equations of physics are missing in financial markets, and that seemingly analogous arguments such as those presented in the Appendix are much more heuristic. The proper comparison is with biology or social science rather than physics or engineering.

II. INTERNATIONAL BONDS

1. Principal component analysis for international markets

All our analysis so far has used US data. Are the results applicable to international markets? To answer this question, we analyze daily historical bond yield data for a range of developed countries, drawn from the historical period 1986–96.

In broad terms, the results carry over. In almost every case, the fundamental yield curve shifts identified by the analysis are a parallel shift, a slope shift and some kind of curvature shift. Moreover, as shown in Exhibit 10, the relative importance of these different yield curve shifts is very similar in different countries – although there is some evidence that parallel shifts are slightly less dominant, and slope shifts are slightly more important, in Europe and Japan than in USD bloc countries.

It is slightly worrying that Switzerland appears to be an exception: the previous results simply do not hold, at least for the data set used. This proves that one cannot simply take the results for granted; they must be verified for each individual country. For example, one should not assume that yield curve risk measures developed for use in the US bond market are equally applicable to some emerging market.

Exhibit 11a shows the estimated form of a parallel shift in different countries. Apart from Switzerland, the results are extremely similar. In other words, duration is an equally valid risk measure in different countries.

Exhibit 11b shows the estimated form of a slope shift in different countries; in this case, estimated slope shifts have been normalized so that the 3-year yield remains fixed and the 10-year yield moves by 100 bp. Unlike the parallel shift, there is some evidence that the slope shift takes different forms in different countries; this is consistent with the findings reported in [Brown and Schaefer]. For risk management applications it is thus prudent to estimate the form of the slope shift separately for each country rather than, e.g., simply using the US slope shift. Note that parallel/slope correlation also varies between countries, as well as over time.

Estimated curvature shifts are not shown, but they are quite different for different countries. Also, breaking the data into sub-periods, the form of the curvature shift typically varies over time as it did with the US data. This is further evidence that there is no stable "curvature shift" which can reliably be used to define an additional measure of non-parallel risk.

2. Co-movements in international bond yields

So far we have only used principal component analysis to look at *data within a single country*, to identify patterns of co-movement between yields at different maturities. We derived the very useful result that two major kinds of co-movement explain most variation in bond yields.

It is also possible to analyze data *across countries*, to identify patterns of co-movements between bond yields in different countries. For example, one could carry out a principal component analysis of daily changes in 10-year bond yields for various countries. Can any useful conclusions be drawn?

The answer is yes, but the results are significantly weaker. Exhibit 10 shows the dominant principal component identified from three separate data sets: 1970-79, 1980-89 and 1990-98. As one might hope, this dominant shift is a kind of "parallel shift", i.e. a simultaneous shift in bond yields, with the same direction and magnitude, in each country. In other words, the notion of "global duration" seems to make sense: the aggregate duration of a global bond portfolio is a meaningful risk measure, which measures the portfolio's sensitivity to an empirically identifiable global risk factor.

However, there are three important caveats. First, the "global parallel shift" is not as dominant as the term structure parallel shift identified earlier. In the 1990s, it explained only 54% of variation in global bond yields; in the 1970s, it explained only 29%. In other words, while duration captures most of the interest rate risk of a domestic bond portfolio, "global duration" captures only half, or less, of the interest rate risk of a global bond portfolio.

Second, the shift in bond yields is not perfectly equal in different countries. It seems to be lower for countries like Japan and Switzerland, perhaps because bond yields have tended to be lower in those countries.

Third, the "global parallel shift" is not universal: not every country need be included. For example, it seems as if Australian and French bond yields did not move in step with other countries' bond yields in the 1970s, and only partially did so in the 1980s. Thus, the relevance of a global parallel shift to each specific country has to be assessed separately.

Apart from the global parallel shift, the other eigenvectors are not consistently meaningful. For example, there is some evidence of a "USD bloc shift" in which US, Canadian, Australian and NZ bond yields move while other bond yields remain fixed, but this result is far from robust.

To summarize, principal component analysis provides some guidelines for global interest rate risk management, but it does not simplify matters as much as it did for yield curve risk. The presence of currency risk is a further complication; we return to this topic below.

3. Correlations: between markets, between yield and volatility

Recall that principal component analysis uses a single correlation matrix to identify dominant patterns of yield shifts. The results imply something about the correlations themselves: for instance, the existence of a global parallel shift that explains around 50% of variance in global bond yields suggests that correlations should, on average, be positive.

However, in global markets, correlations are notoriously time-varying: see Exhibit 13. In fact, short term correlations between 10-year bond yields in different countries are significantly less stable than correlations between yields at different maturities within a single country. This means that, at least for short time horizons, one must be especially cautious in using the results of principal component analysis to manage a global bond position.

We now discuss a somewhat unrelated issue: the relationship between yield and volatility, which has been missing from our analysis so far. Principal component analysis estimates the form of the dominant yield curve shifts, namely parallel and slope shifts. It says nothing useful about the size of these shifts, i.e. about parallel and slope volatility. These can be estimated instantaneously, using historical or implied volatilities. But for stress testing and scenario analysis, one needs an additional piece of information: whether there is a relationship between volatility and (say) the outright level of the yield curve. For example, when stress testing a trading book under a +100 bp scenario, should one also change one's volatility assumption?

It is difficult to answer this question either theoretically or empirically. For example, most common term structure models assume that basis point (parallel) volatility is either independent of the yield level, or proportional to the yield level; but these assumptions are made for technical convenience, rather that being driven by the data.

Here are some empirical results. Exhibits 14a–c plot 12-month historical volatilities, expressed as a percentage of the absolute yield level, against the average yield level itself. If basis point volatility were always proportional to the yield level, these graphs would be horizontal lines; if basis point volatility were constant, these graphs would be hyperbolic.

Neither seems to be the case. The Japanese data set suggests that when yields are under around 6%–7%, the graph is hyperbolic. All three data sets suggest that when yields are in the 7%–10% range, the graph is horizontal. And the US data set suggests that when yields are over 10%, the graph actually slopes upward: when yields rise, volatility rises more than proportionately. But in every case, the results are confused by the presence of volatility spikes.

The conclusion is that, when stress testing a portfolio, it is safest to assume that when yields fall, basis point volatility need not fall; but when yields rise, basis point volatility will also rise. Better yet, one should run different volatility scenarios as well as interest rate scenarios.

III. PRACTICAL IMPLICATIONS

1. Risk management for a leveraged trading desk

This section draws some practical conclusions from the above analysis, and briefly sketches some suggestions about risk measurement and risk management policy; more detailed proposals may be found elsewhere in this volume.

Since parallel and slope shifts are the dominant yield curve risk factors, it makes sense to focus on measures of parallel and slope risk; to structure limits in terms of maximum parallel and slope risk rather than more rigid limits for each point of the yield curve; and to design flexible hedging strategies based on matching parallel and slope risk. If the desk as a whole takes proprietary interest rate risk positions, it is most efficient to specify these in terms of target exposures to parallel and slope risk, and leave it to individual traders to structure their exposures using specific instruments.

Rapid stress testing and value-at-risk estimates may be computed under the simplifying assumption that only parallel and slope risk exist. This approach is not meant to replace a standard VaR calculation using a covariance matrix for a whole set of reference maturities, but to supplement it.

A simplified example of such a VaR calculation appears in Exhibit 15, which summarizes both the procedure and the results. It compares the value-at-risk of three positions, each with a net market value of $100m: a *long portfolio* consisting of a single position in a 10-year par bond; a *steepener portfolio* consisting of a long position in a 2-year bond and a short position in a 10-year bond with offsetting durations, i.e. offsetting exposures to parallel risk; and a *butterfly portfolio* consisting of long/short positions in cash and 2-, 5- and 10-year bonds with zero net exposure to both parallel and slope risk. For simplicity, the analysis assumes a 'total volatility' of bond yields of about 100 bp p.a., which is broadly realistic for the US market.

The long portfolio is extremely risky compared to the other two portfolios; this reflects the fact that most of the observed variance in bond yields comes from parallel shifts, to which the other two portfolios are immunized. Also, the butterfly portfolio appears to have almost negligible risk: by this calculation, hedging both parallel and slope risk removes over 99% of the risk. However, it must be remembered that the procedure assumes that the first three principal components are the only sources of risk.

This calculation was oversimplified in several ways: for example, in practice the volatilities would be estimated more carefully, and risk computations would probably be carried out on a cashflow-by-cashflow basis. But the basic idea remains straightforward. Because the calculation can be carried out rapidly, it is easy to vary assumptions about volatility/yield relationships and about correlations, giving additional insight into the risk profile of the portfolio. Of course, the calculation is approximate, and in practice large exposures at specific maturities should not be ignored. That would tend to understate the risk of butterfly trades, for example.

However, it is important to recognize that a naïve approach to measuring risk, which ignores the information about co-movements revealed by a principal component analysis, will tend to overstate the risk of a butterfly position; in fact, in some circumstances a butterfly position is no riskier than, say, an exposure to the spread between on-the-run and off-the-run Treasuries. In other words, the analysis helps risk managers gain some sense of perspective when comparing the relative importance of different sources of risk.

Risk management for a global bond book is harder. The results of the analysis are mainly negative: they suggest that the most prudent course is to manage each country exposure separately. For value-at-risk calculations, the existence of a "global parallel shift" suggests an alternative way to estimate risk, by breaking it into two components: (a) risk arising from a global shift in bond yields, and (b) country-specific risk relative to the global component.

This approach has some important advantages over the standard calculation, which uses a covariance matrix indexed by country. First, the results are less sensitive to the covariances, which are far from stable. Second, it is easier to add new countries to the analysis. Third, it is easier to incorporate an assumption that changes in yields have a heavy-tailed (non-Gaussian) distribution, which is particularly useful when dealing with emerging markets. Again, the method is not proposed as a replacement for standard VaR calculations, but as a supplement.

2. Total return management and benchmark choice

For an unleveraged total return manager, many of the proposals are similar. It is again efficient to focus mainly on parallel and slope risk when setting interest rate risk limits, implementing an interest rate view, or designing hedging strategies. This greatly simplifies interest rate risk management, freeing up the portfolio manager's time to focus on monitoring other forms of risk, on assessing relative value, and on carrying out more detailed scenario analysis.

Many analytics software vendors, such as CMS, provide measures of slope risk. Investment managers should ensure that such a measure satisfies two basic criteria. First, it should be consistent with the results of a principal component analysis: a measure of slope risk based on an unrealistic slope shift is meaningless. Second, it should be easy to run, and the results should be easy to interpret: otherwise, it will rarely be used, and slope risk will not be monitored effectively.

The above comments on risk management of global bond positions apply equally well in the present context. However, there is an additional complication. Global bond investors tend to have some performance benchmark, but it is most unclear how an "optimal" benchmark should be constructed, and how risk should be measured against it. For example, some US investors simply use a US domestic index as a benchmark; many use a currency unhedged global benchmark.

(Incidentally, the weights of a global benchmark are typically determined by issuance volumes. This is somewhat arbitrary: it means that a country's weight in the index depends on its fiscal policy and on the precise way public sector borrowing is funded. Mason has suggested using GDP weights; this tends to lower the risk of the benchmark.)

Exhibits 16a–c may be helpful. They show the risk/return profile, in USD terms, of a US domestic bond index; currency unhedged and hedged global indexes; and the full range of *post hoc* efficient currency unhedged and hedged portfolios. Results are displayed separately for the 1970s, 1980s and 1990s data sets. The first observation is that the US domestic index has a completely different (and inferior) risk/return profile to any of the global portfolios. It is not an appropriate benchmark.

The second observation is that hedged and unhedged portfolios behave in completely different ways. In the 1970s, hedged portfolios were unambiguously superior; in the 1980s, hedged and unhedged portfolios behaved almost like two different asset types; and in the 1990s, hedged and unhedged portfolios seemed to lie on a continuous risk/return scale, with hedged portfolios at the less risky end. If a benchmark is intended to be conservative, a currency hedged benchmark is clearly appropriate.

What, then, is a suitable global benchmark? None of the *post hoc* efficient portfolios will do, since the composition of efficient portfolios is extremely unstable over time – essentially because both returns and covariances are unstable. The most plausible candidate is the currency hedged global index. It has a stable composition, has relatively low risk, and is consistently close to the efficient frontier.

Once a benchmark is selected, principal component analysis may be applied as follows. First, it identifies countries which may be regarded as particularly risk relative to the benchmark; in the 1970s and 1980s this would have included Australia and France (see Exhibit 12). Note that this kind of result is more easily read off from the analysis than by direct inspection of the correlations.

Second, it helps managers translate country-specific views into strategies. That it, by estimating the proportion of yield shifts attributable to a global parallel shift (around 50% in the 1990s) it allows managers will a bullish or bearish view on a specific country to determine an appropriate degree of overweighting.

Third, it assists managers who choose to maintain open currency risk. A more extensive analysis can be used to identify "currency blocs" (whose membership may vary over time) and to estimate co-movements between exchange rates and bond yields. However, all such results must be used with great caution.

3. Asset/liability management and the use of risk buckets

For asset/liability managers, the recommendations are again quite similar. One should focus on immunizing parallel risk (duration) and slope risk. If these two risk factors are well matched, then from an economic point of view the assets are an effective hedge for the liabilities. Key rate durations are a useful way to measure exposure to individual points on the yield curve; but it is probably unnecessary to match all the key rate durations of assets and liabilities precisely. However, one does need to treat both the short and the very long end of the yield curve separately.

Regarding the long end of the yield curve, it is necessary to ensure that really long-dated liabilities are matched by similarly long-dated assets. For example, one does not want to be hedging 30-year liabilities with 10-year assets, which would be permitted if one focused only on parallel and slope risk. Thus, it is desirable to ensure that 10-year to 30-year key rate durations are reasonably well matched.

Regarding the short end of the yield curve, two problems arise. First, for maturities less than about 18–24 months – roughly coinciding with the liquid part of the Eurodollar futures strip – idiosyncratic fluctuations at the short end of the curve introduce risks additional to parallel and slope risk. It is safest to identify and hedge these separately, either using duration bucketing or partial durations.

Second, for maturities less than about 12 months, it is desirably to match actual cashflows and not just risks. That is, one needs to generate detailed cashflow forecasts rather than simply matching interest rate risk measures.

To summarize, an efficient asset/liability management policy might be described as follows: *from 0–12 months, match cashflows in detail; from 12-24 months, match partial durations or duration buckets in detail; from 2–15, match parallel and slope risk only; beyond 15 years, ensure that partial durations are roughly matched too.*

Finally, one must not forget optionality. If the assets have very different option characteristics from the liabilities – which may easily occur when callable bonds or mortgage-backed securities are held – then it is not sufficient to match interest rate exposure in the current yield curve environment. One must also ensure that risks are matched under different interest rate and volatility scenarios. Optionality is treated in detail elsewhere in this volume.

In conclusion: Principal component analysis suggests a simple and attractive solution to the problem of efficiently managing non-parallel yield curve risk. It is easy to understand, fairly easy to implement, and various off-the-shelf implementations are available. However, there are quite a few subtleties and pitfalls involved. Therefore, risk managers should not rush to implement policies, or to adopt vendor systems, without first deepening their own insight through experimentation and reflection.

APPENDIX: ECONOMIC FACTORS DRIVING THE CURVE

1. Macroeconomic explanation of parallel and slope risk

This appendix presents some theoretical explanations for why (a) parallel and slope shifts are the dominant kinds of yield curve shift that occur, (b) curvature shifts are observed but tend to be both transitory and inconsistent in form, and (c) the behavior of the short end of the yield curve is quite idiosyncratic. The theoretical analysis helps to ascertain which empirical findings are really robust and can be relied upon: that is, an empirical result is regarded as reliable if it has a reasonable theoretical explanation. For reasons of space, the arguments are merely sketched.

We first explain why parallel and slope shifts emerge naturally from a macroeconomic analysis of interest rate expectations. For simplicity, we use an entirely standard linear macroeconomic model, shown in Exhibit 17; see [Frankel] for details.

The model is used in the following way. Bond yields are determined by market participants' expectations about future short-term interest rates. These in turn are determined by their expectations about the future path of the economy: output, prices and the money supply. It is assumed that market participants form these expectations in a manner consistent with the macroeconomic model. Now, the model implies that the short-term interest rate must evolve in a certain fixed way; thus, market expectations must, "in equilibrium", take a very simple form.

To be precise, it follows from the theorem stated in Exhibit 17 that if i_0 is the current short-term nominal interest rate, i_t is the currently expected future interest rate at some future time *t*, and i_∞ is the long-term expected future interest rate, then rational interest rate expectations must take the following form in equilibrium:

$$i_t = i_\infty + \left(i_0 - i_\infty\right)e^{-d}$$

In this context, a slope shift corresponds to a change in either i_∞ or i_0, while a parallel shift corresponds to a simultaneous change in both. Exhibit 18 shows, schematically, the structure of interest rate expectations as determined by the model. The expected future interest rate at some future time is equal to the expected future rate of inflation, plus the expected future real rate. (At the short end, some distortion is possible, of which more below.)

In this setting, yield curve shifts occur when market participants revise their expectations about future interest rates – that is, about future inflation and output growth. A *parallel shift* occurs when both short-term and long-term expectations change at once, by the same amount. A *slope shift* occurs when short-term expectations change but long-term expectations remain the same, or vice versa.

Why are parallel shifts so dominant? The model allows us to formalize the following simple explanation: in financial markets, changes in long-term expectations are primarily driven by short-term events, which of course also drive changes in short-term expectations. For a detailed discussion of this point, see [Keynes].

Why is the form of a slope shift relatively stable over time, but somewhat different in different countries? In this setting, the shape taken by a slope shift is determined by d, and thus by the elasticity parameters g, f, l, r of the model. These parameters depend in turn on the flexibility of the economy and its institutional framework – which may vary from country to country – but not on the economic cycle, or on the current values of economic variables. So d should be reasonably stable.

Finally, observe that there is nothing in the model which ensures that parallel and slope shifts should be uncorrelated. In fact, using the most natural definition of "slope shift", there will almost certainly be a correlation – but the value of the correlation coefficient is determined by how short-term events affect market estimates of the different model variables, not by anything in the underlying model itself. So the model does not give us much insight into correlation risk.

2. Volatility shocks and curvature risk

We have seen that, while principal component analysis seems to identify curvature shifts as a source of non-parallel risk, on closer inspection the results are somewhat inconsistent. That is, unlike parallel and slope shifts, curvature shifts do not seem to take a consistent form, making it difficult to design a corresponding risk measure.

The main reason for this is that "curvature shifts" can occur for a variety of quite different reasons. A change in mid-range yields can occur because (a) market volatility expectations have changed, (b) the "term premium" for interest rate risk has changed, (c) market segmentation has caused a temporary supply/demand imbalance at specific maturities, or (d) a change in the structure of the economy has caused a change in the value d above. We briefly discuss each of these reasons, but readers will need to consult the references for further details.

Regarding (a): The yield curve is determined by forward short-term interest rates, but these are not completely determined by expected future short-term interest rates; forward rates have two additional components. First, forward rates display a downward "convexity bias", which varies with the square of maturity. Second, forward rates display an upward "term premium", or risk premium for interest rate risk, which (empirically) rises at most linearly with maturity. The size of both components obviously depends on expected volatility as well as maturity.

A change in the market's expectations about future interest rate volatility causes a curvature shift for the following reason. A rise in expected volatility will not affect short maturity yields since both the convexity bias and the term premium are negligible. Yields at intermediate maturities will rise, since the term premium dominates the convexity bias at these maturities; but yields at sufficiently long maturities will fall, since the convexity bias eventually dominates. The situation is illustrated in Exhibit 19. The precise form taken by the curvature shift will depend on the empirical forms of the convexity bias and the term premium, neither of which are especially stable.

Regarding (b): The term premium itself, as a function of maturity, may change. In theory, if market participants expect interest rates to follow a random walk, the term premium should be a linear function of maturity; if they expect interest rates to range trade, or mean revert, the term premium should be sub-linear (this seems to be observed in practice). Thus, curvature shifts might occur when market participants revise their expectations about the nature of the dynamics of interest rates, perhaps because of a shift in the monetary policy regime. Unfortunately, effects like this are nearly impossible to measure precisely.

Regarding (c): Such manifestations of market ineffiency do occur, even in the US market. They do not assume a consistent form, but can occur anywhere on the yield curve. Note that, while a yield curve distortion caused by a short term supply/demand imbalance may have a big impact on a leveraged trading book, it might not matter so much to a typical mutual fund or asset/liability manager.

Regarding (d): It is highly unlikely that short-term changes in d occur, although it is plausible that this parameter may drift over a secular time scale. There is little justification for using "sensitivity to changes in d" as a measure of curvature risk.

Curvature risk is clearly a complex issue, and it may be dangerous to attempt to summarize it using a single stylized "curvature shift". It is more appropriate to use detailed risk measures such as key rate durations to manage exposure to specific sections of the yield curve.

3. The short end and monetary policy distortions

The dynamics of short maturity money market yields is more complex and idiosyncratic than that of longer maturity bond yields. We have already seen a hint of this in Exhibit 7c, which shows that including T-bill yields in the data set radically changes the results of a principal component analysis; the third eigenvector represents, not a "curvature shift" affecting 3-5 year maturities, but a "hump shift" affecting maturities around 1 year. This is confirmed by more careful studies.

As with curvature shifts, hump shifts might be caused by changes in the term premium. But there is also an economic explanation for this kind of yield curve shift: it is based on the observation that market expectations about the path of interest rates in the near future can be much more complex than longer term expectations.

For example, market participants may believe that monetary policy is "too tight" and can make detailed forecasts about when it may be eased. Near term expected future interest rates will not assume the simple form predicted by the macroeconomic model of Exhibit 16 if investors believe that monetary policy is "out of equilibrium". This kind of bias in expectations can create a hump or bowl at the short end of the yield curve, and is illustrated schematically in Exhibit 18.

One would not expect a "hump factor" to take a stable form, since the precise form of expectations, and hence of changes in expectations, will depend both on how monetary policy is currently being run and on specific circumstances. Thus, one should not feed money market yields to a principal component analysis and expect it to derive a reliable "hump shift" for use in risk management.

For further discussion and analysis, see [Phoa]. The overall conclusion is that when managing interest rate risk at the short end of the yield curve, measures of parallel and slope risk must be supplemented by more detailed exposure measures. Similarly, reliable hedging strategies cannot be based simply on matching parallel and slope risk, but must make use of a wider range of instruments such as a whole strip of Eurodollar futures contracts.

LITERATURE GUIDE AND REFERENCES

The following brief list of references is provided merely as a starting point for further reading, which might be structured as follows.

For general background on matrix algebra and matrix computations, both [Jennings & McKeown] and the classic [Press &al.] are useful, though there are a multitude of alternatives. On principal components analysis, [Litterman & Scheinkman] and [Garbade] are still well worth reading, perhaps supplemented by [Phoa] which contains further practical discussion. This should be followed with [Buhler & Zimmermann] and [Hiraki &al.] which make use of additional statistical techniques not discussed in the present chapter.

However, at this point it is probably more important to gain hands-on experience with the techniques and, especially, the data. Published results should not be accepted unquestioningly, even those reported here! For numerical experimentation, a package such as Numerical Python or MATLAB™ is recommended; attempting to write one's own routines for computing eigenvectors is emphatically *not* recommended. Finally, historical bond data for various countries may be obtained from central banking authorities, often via the World Wide Web.[*]

Brown, R. and Schaefer, S., "Interest rate volatility and the shape of the term structure", in Howison, S., Kelly, F. and Wilmott, P. (eds.), *Mathematical Models in Finance*, Chapman and Hall, 1995.

Buhler, A. and Zimmermann, H., "A statistical analysis of the term structure of interest rates in Switzerland and Germany", *J. Fixed Income*, December 1996.

Frankel, J., *Financial Markets and Monetary Policy*, MIT Press, 1995.

Garbade, K., *Fixed Income Analytics*, MIT Press, 1996.

Hiraki, T., Shiraishi, N. and Takezawa, N., "Cointegration, common factors and the term structure of Yen offshore interest rates", *J. Fixed Income*, December 1996.

Ho, T., "Key rate durations: measures of interest rate risk", *Journal of Fixed Income*, September 1992.

Jennings, A. and McKeown, J., *Matrix Computation (2nd Edition)*, Wiley, 1992.

Keynes, J. M., *The General Theory of Employment, Interest and Money*, Macmillan, 1936.

Litterman, R. and Scheinkman, J., "Common factors affecting bond returns", *J. Fixed Income*, June 1991.

Phoa, W., *Advanced Fixed Income Analytics*, Frank J. Fabozzi Associates, 1998.

Phoa, W., *Foundations of Bond Market Mathematics*, CMS Research Report, 1998.

Press, W., Teukolsky, S., Vetterling, W. and Flannery, B., *Numerical Recipes in C: The Art of Scientific Computing (2nd Edition)*, Cambridge University Press, 1992.

Rebonato, R., and Cooper, I., "The limitations of simple two-factor interest rate models", *J. Financial Engineering*, March 1996.

---

[*] The international data sets used here were provided by Sean Carmody and Richard Mason of Deutsche Bank Securities. The author would also like to thank them for many useful discussions.

LIST OF EXHIBITS

EXHIBITS

EXHIBIT 1 ■ Key rate durations of non-callable Treasury bonds

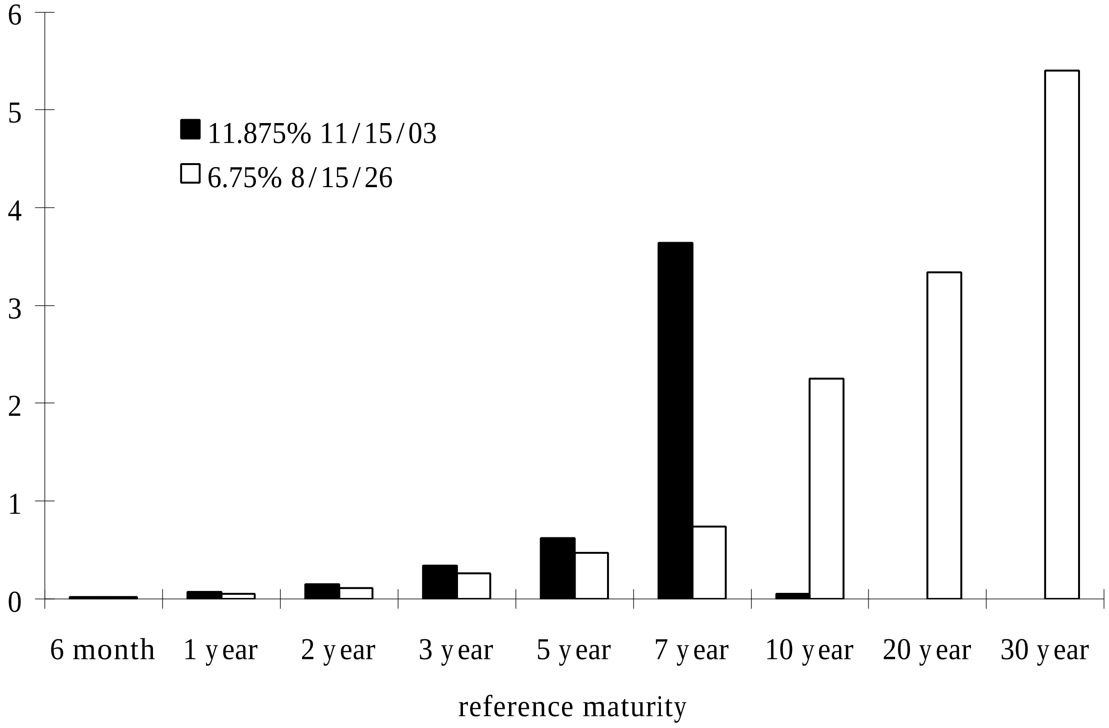

EXHIBIT 2a ■ Vibration modes of the cantilever

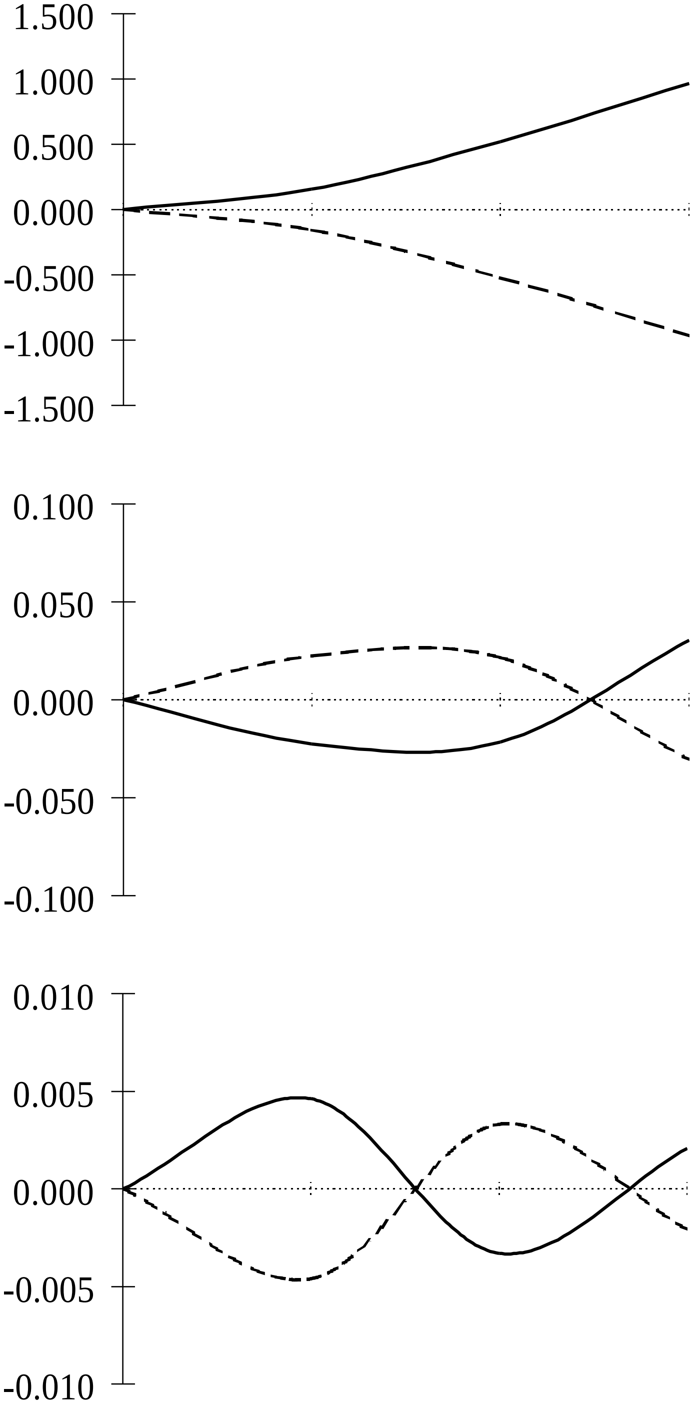

EXHIBIT 2b ■ Principal component analysis of a sample financial data set

Daily changes

1	1	1	1
0	1	2	3
-2	-2	-2	-2
5	4	3	2
0	0	1	0

Correlation matrix

1.00	0.97	0.83	0.59
0.97	1.00	0.92	0.77
0.83	0.92	1.00	0.90
0.59	0.77	0.90	1.00

Eigenvalues and eigenvectors

[A]	0.000	(0%)	0.607	–0.762	0.000	0.225
[B]	0.037	(1%)	–0.155	–0.263	0.827	–0.471
[C]	0.462	(11%)	–0.610	–0.274	0.207	0.715
[D]	3.501	(88%)	0.486	0.524	0.522	0.465

EXHIBIT 3 ■ Principal component analysis of US Treasury yields, 1993-98

	1 year	2 year	3 year	5 year	7 year	10 year	20 year	30 year
0.3%	0.00	0.05	-0.20	0.31	-0.63	0.50	0.32	-0.35
0.3%	0.00	-0.08	0.49	-0.69	0.06	0.27	0.30	-0.34
0.2%	0.01	-0.05	-0.10	0.25	0.30	-0.52	0.59	-0.48
0.4%	-0.05	-0.37	0.65	0.27	-0.45	-0.34	0.08	0.22
0.6%	0.21	-0.71	0.03	0.28	0.35	0.34	-0.27	-0.26
1.1%	0.70	-0.30	-0.32	-0.30	-0.19	-0.12	0.28	0.32
5.5%	-0.59	-0.37	-0.23	-0.06	0.14	0.20	0.44	0.45
91.7%	0.33	0.35	0.36	0.36	0.36	0.36	0.35	0.35

[Historical bond yield data provided by the Federal Reserve Board.]

EXHIBIT 4 ■ Relative importance of principal components, 1963-98

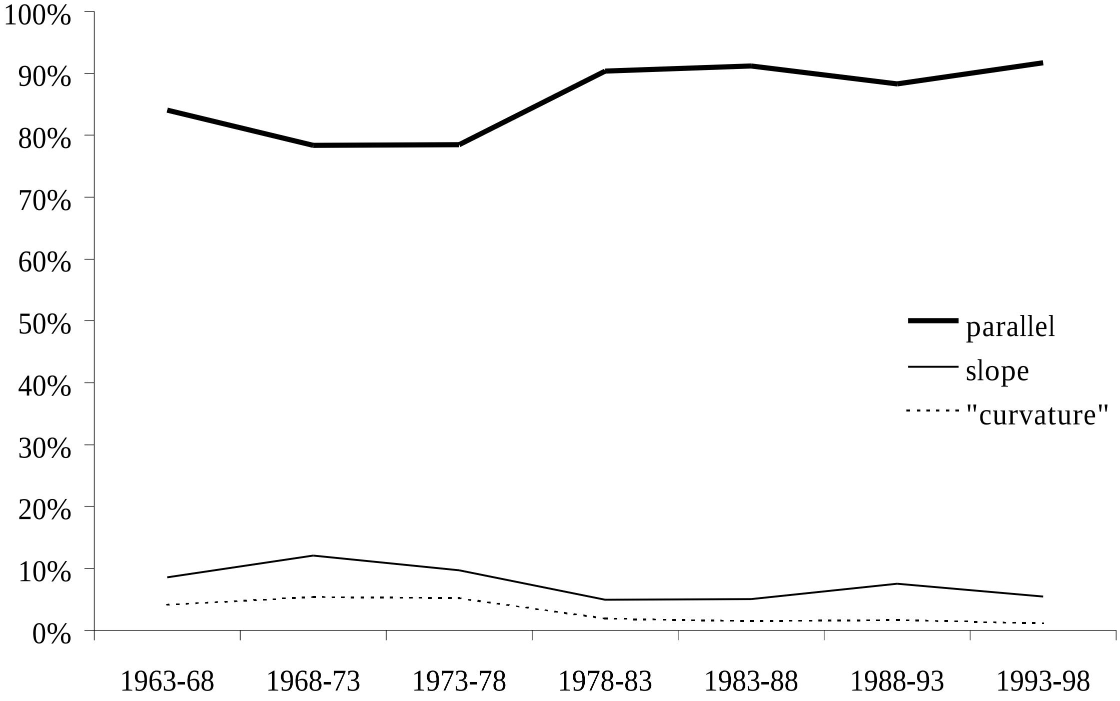

EXHIBIT 5a ■ Shape of "parallel" shift, 1963-98

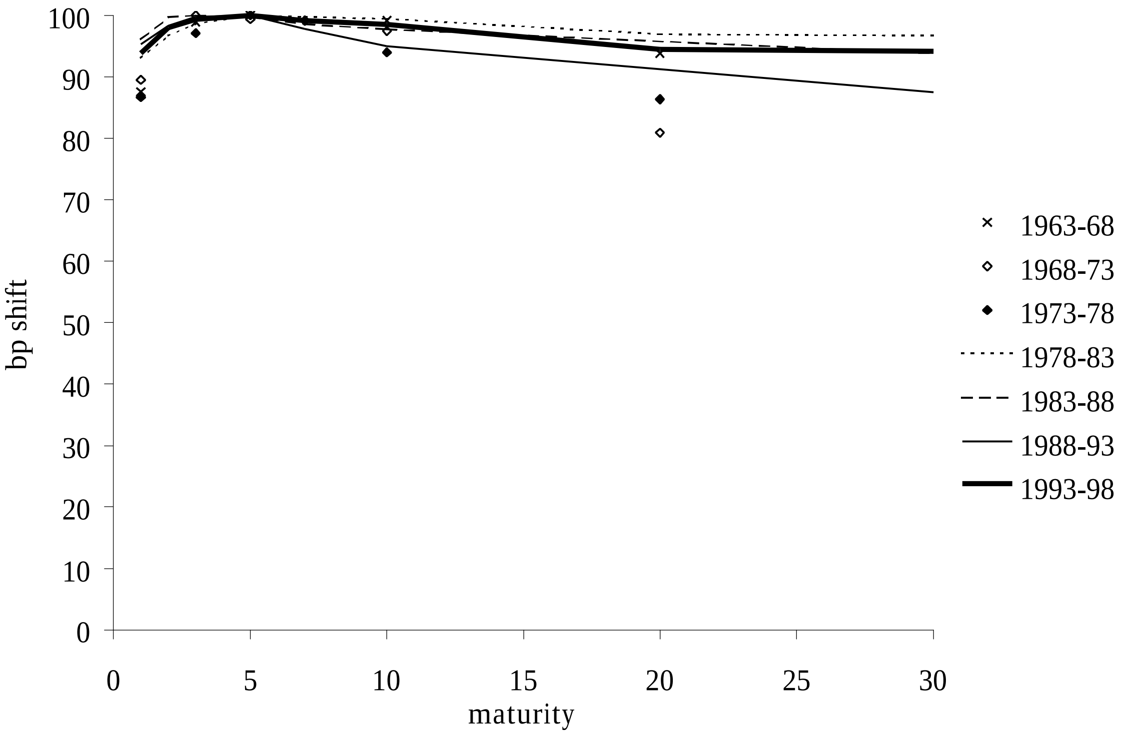

EXHIBIT 5b ■ Shape of "slope" shift, 1963-98

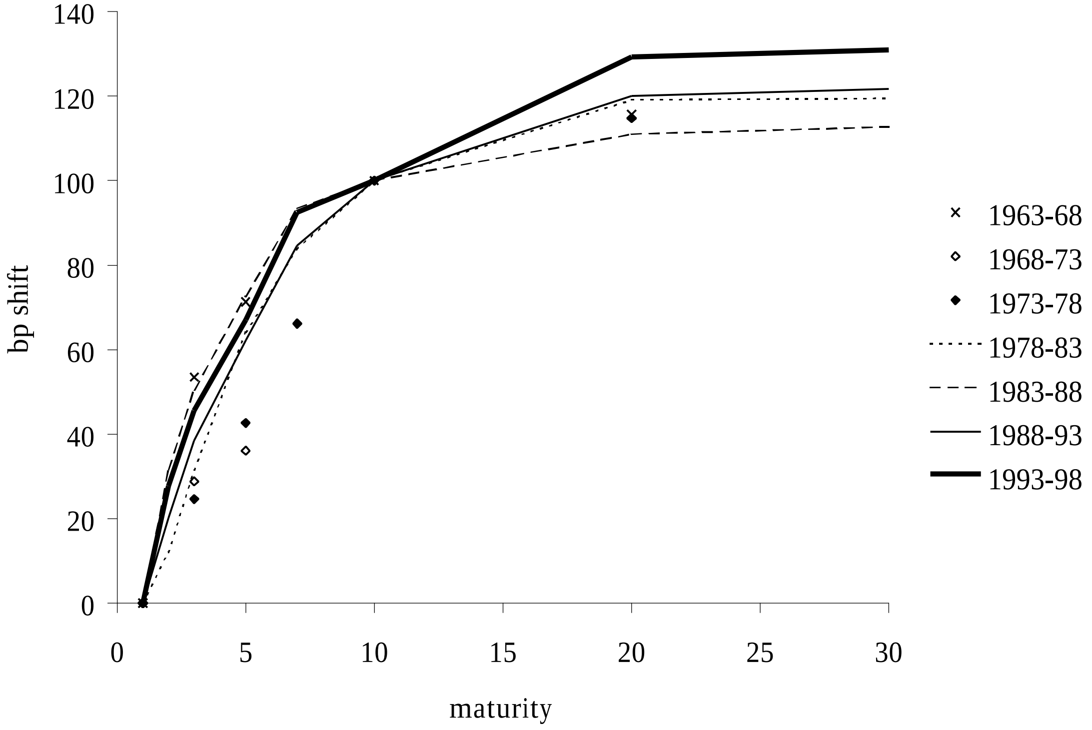

EXHIBIT 5c ■ Shape of "curvature" shift, 1963-98

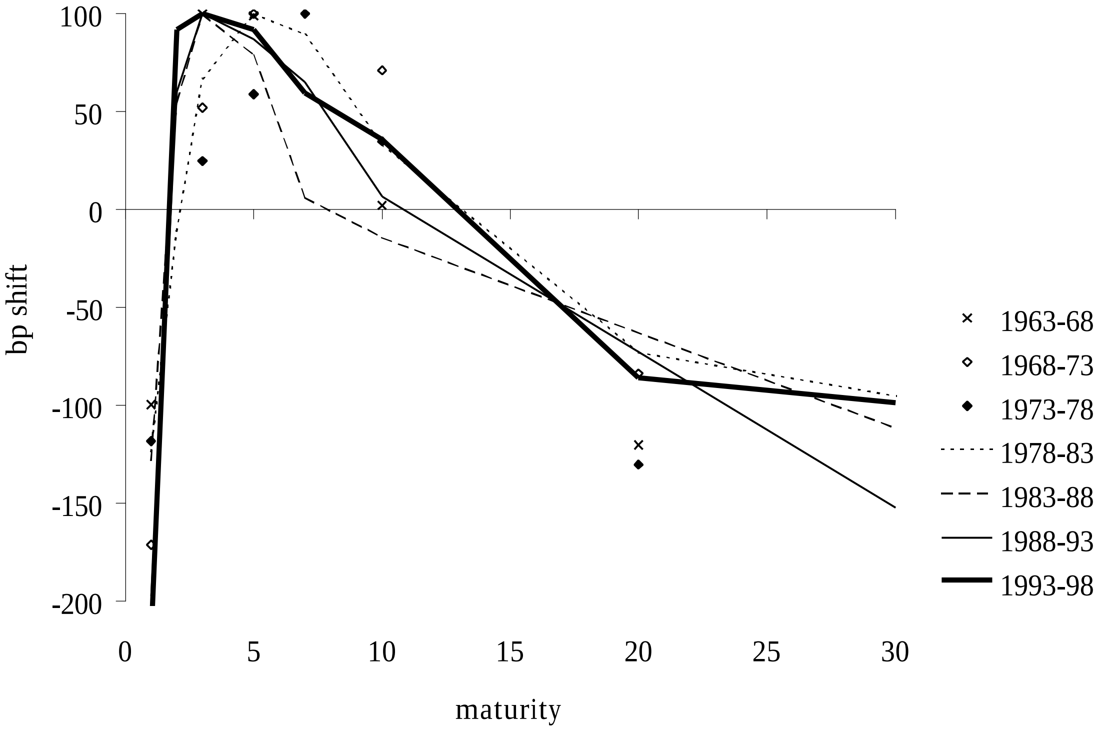

EXHIBIT 6 ■ Relative importance of principal components, with / without T-bills

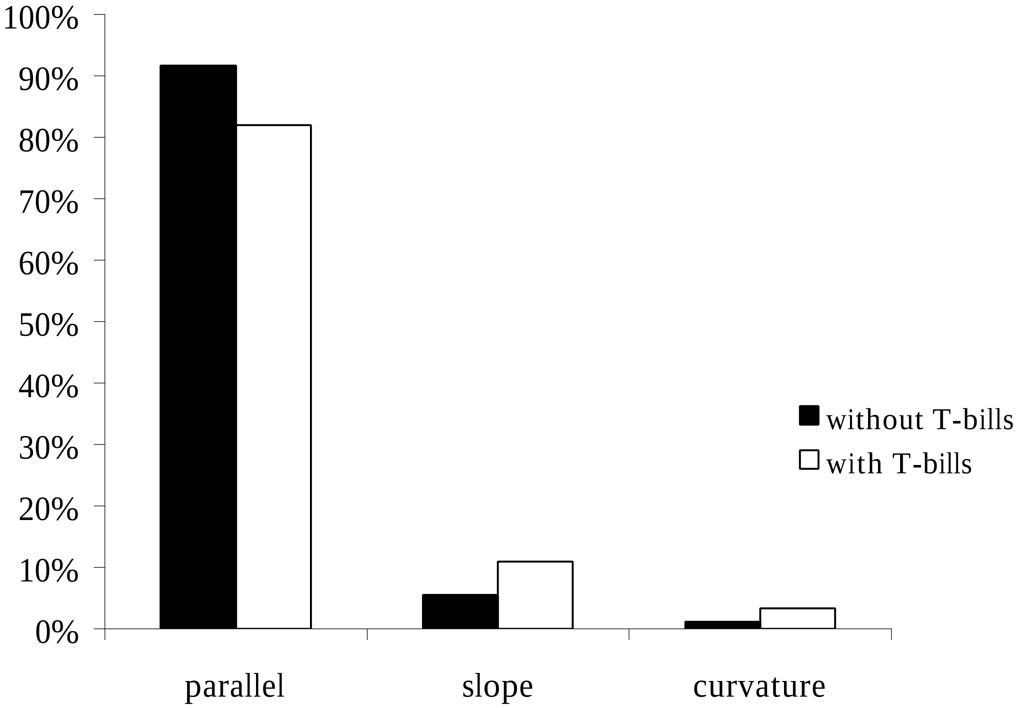

EXHIBIT 7a ■ Estimated "parallel" shift, with/without T-bills

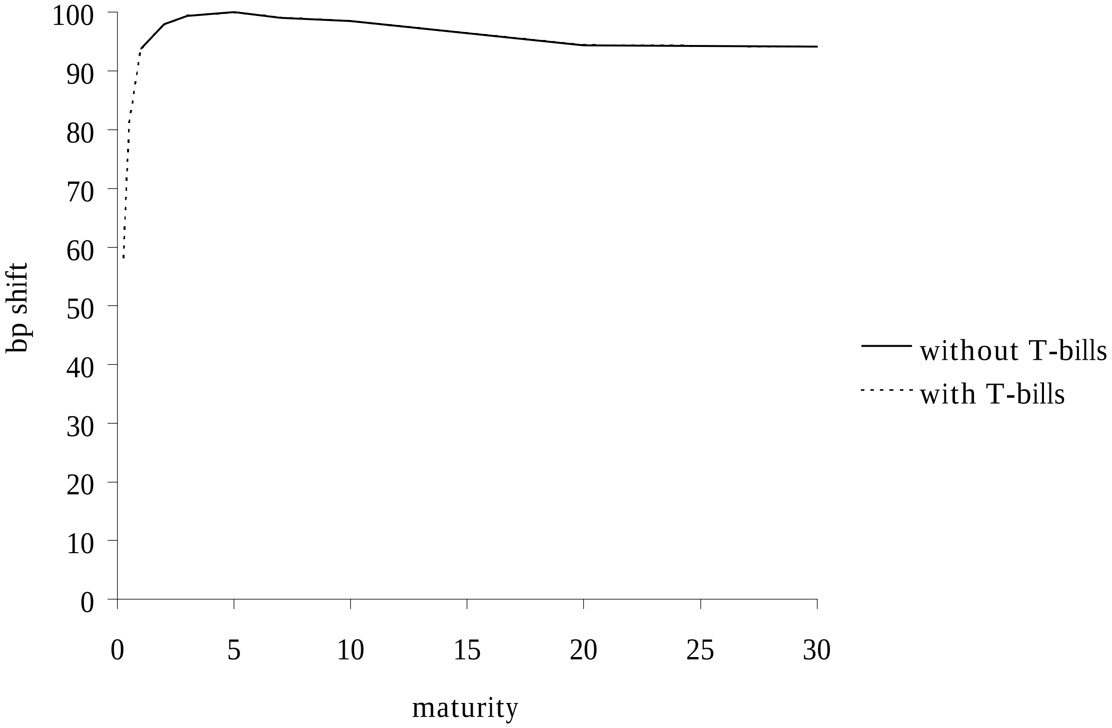

EXHIBIT 7b ■ Estimated "slope" shift, with/without T-bills

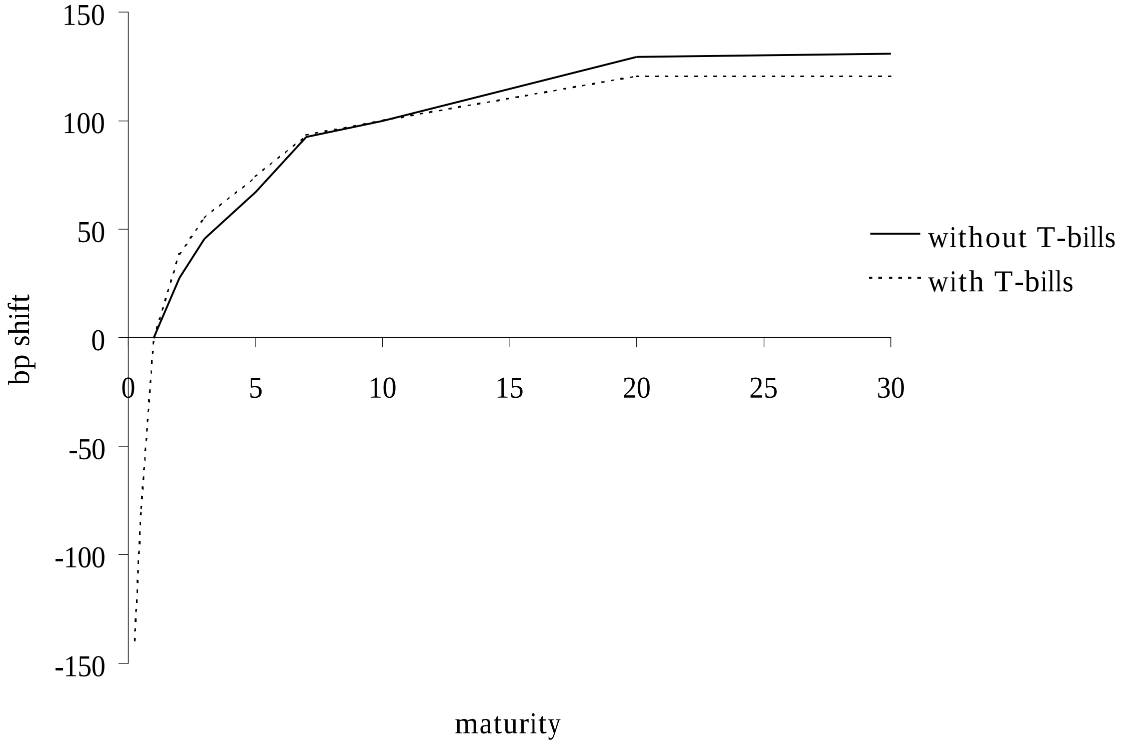

EXHIBIT 7c ■ Estimated “curvature” shift, with / without T-bills

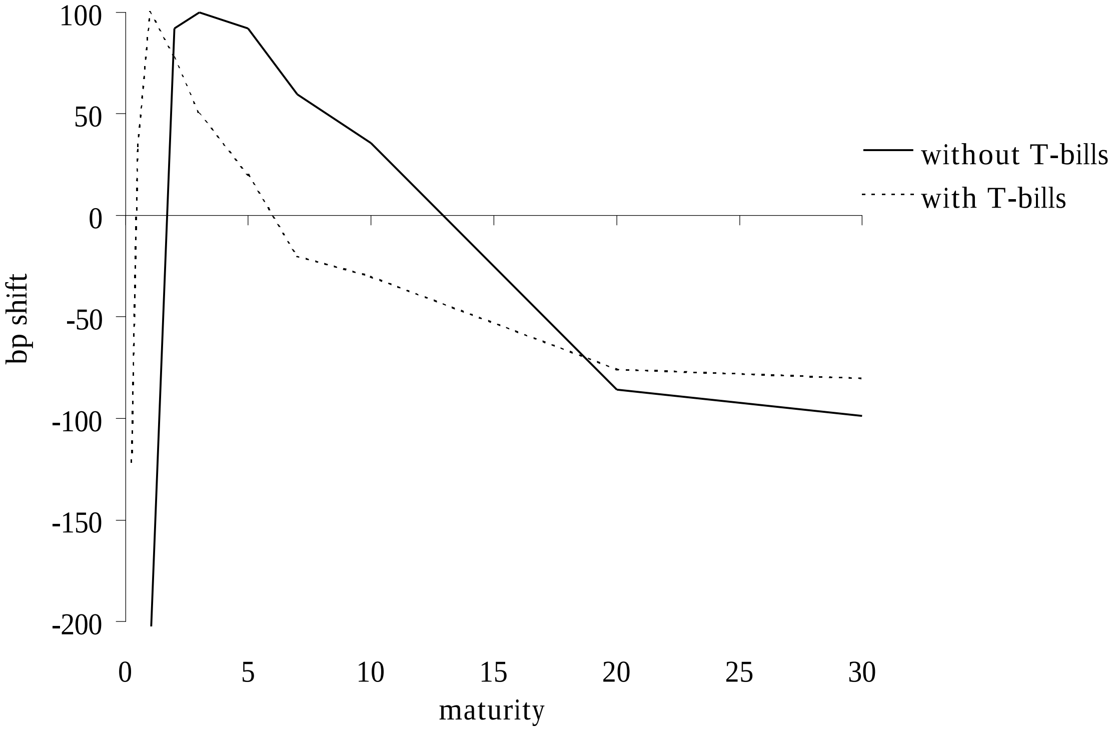

EXHIBIT 8a ■ Empirical Treasury yield correlations

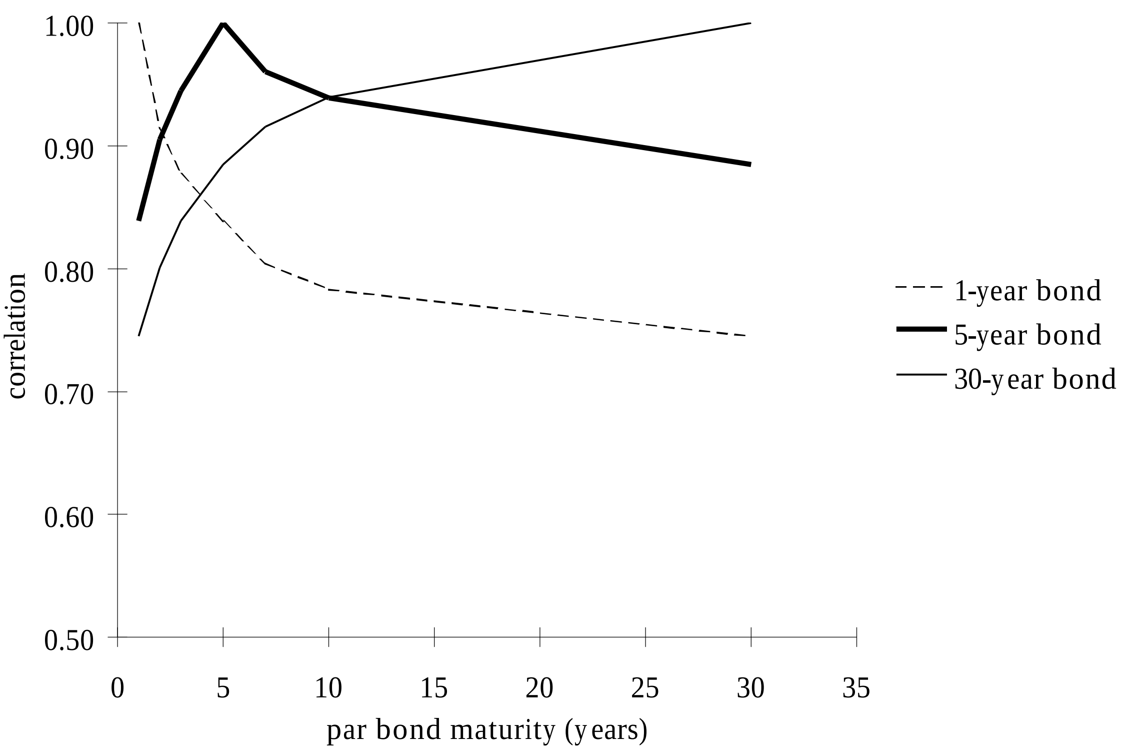

EXHIBIT 8b ■ Theoretical Treasury yield correlations, two factor model

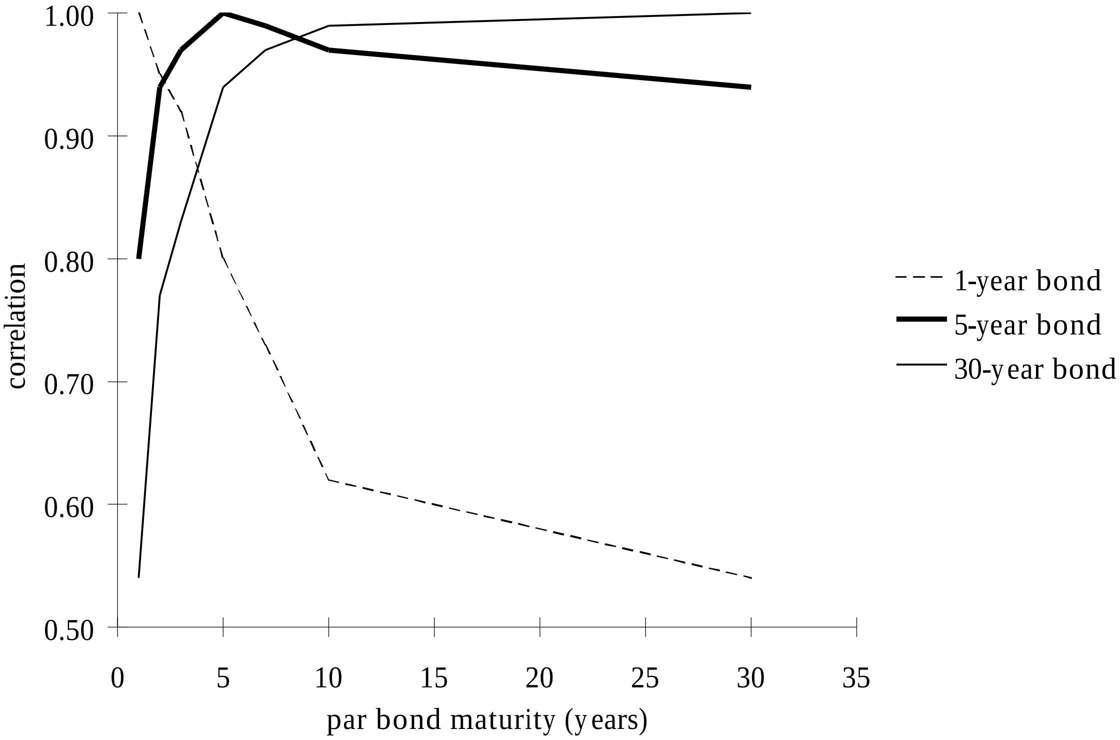

EXHIBIT 9 ■ Actual Treasury yield vs. yield predicted by two factor model

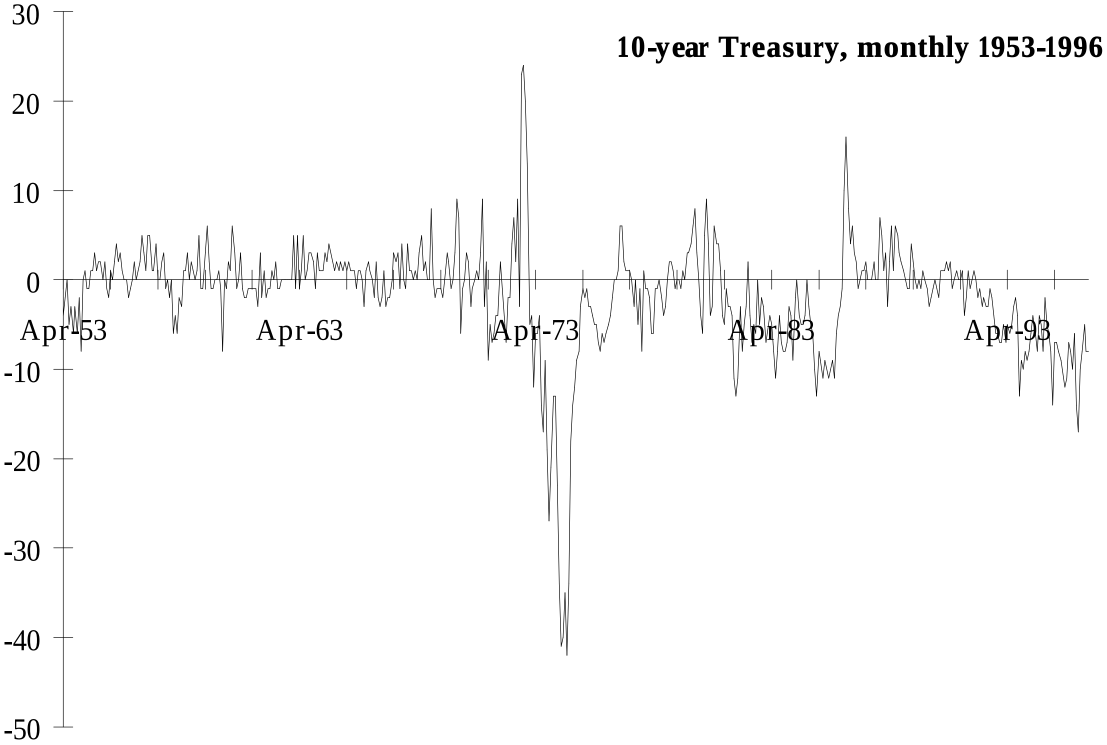

EXHIBIT 10 ■ Relative importance of principal components in various countries

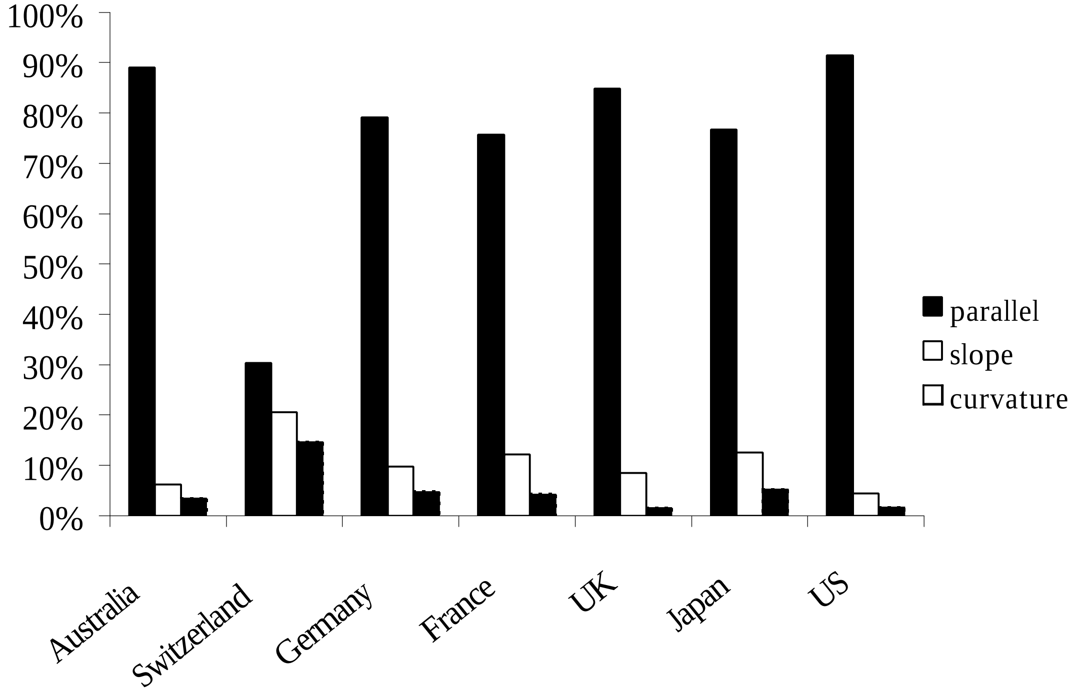

[Historical bond yield data provided by Deutsche Bank Securities.]

EXHIBIT 11a ■ Shape of "parallel" shift in different countries

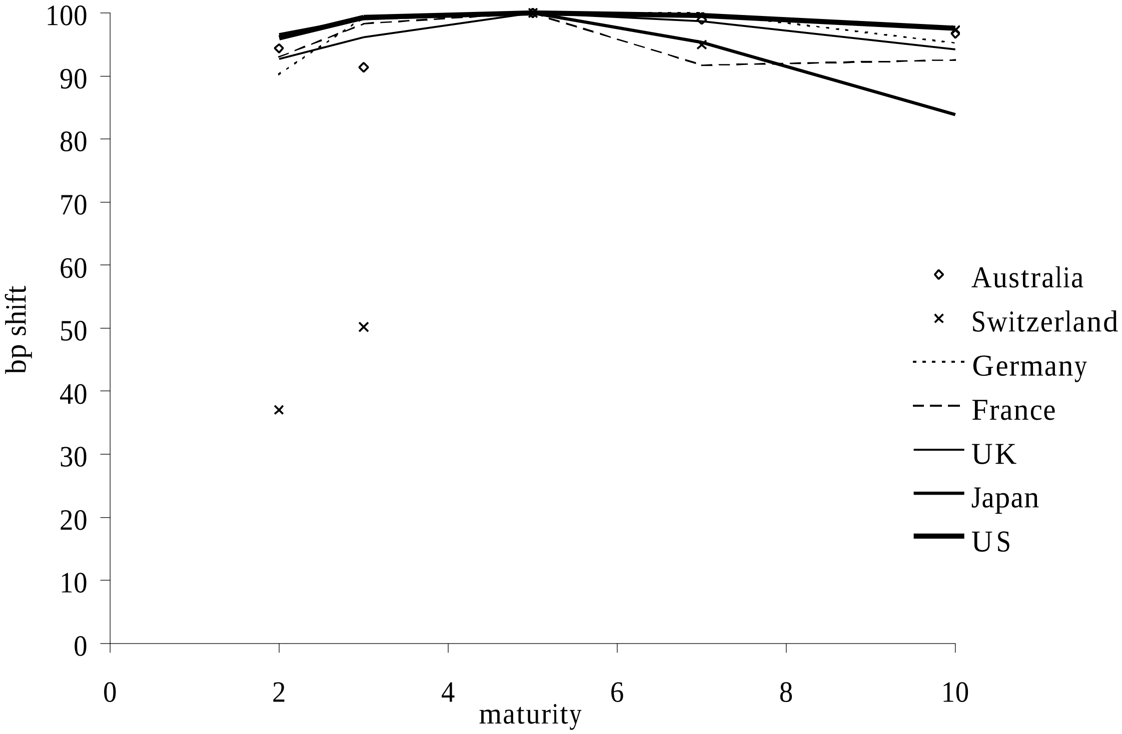

EXHIBIT 11b ■ Shape of "slope" shift in different countries

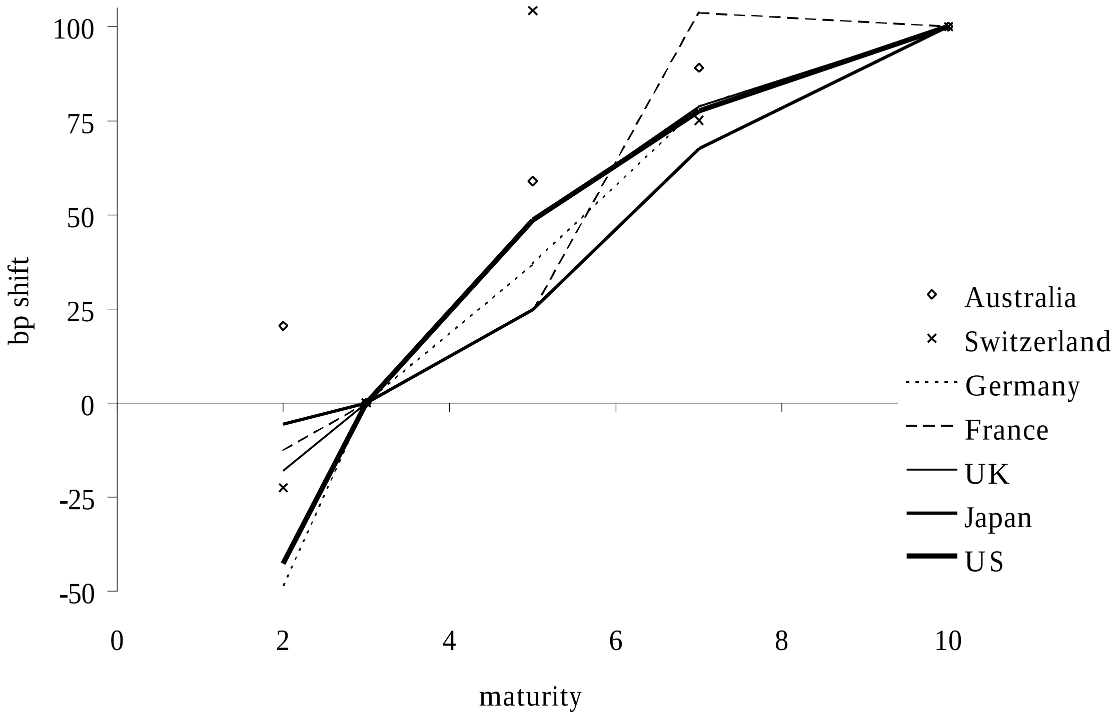

EXHIBIT 12 ■ Dominant principal component, global 10-year bond yields

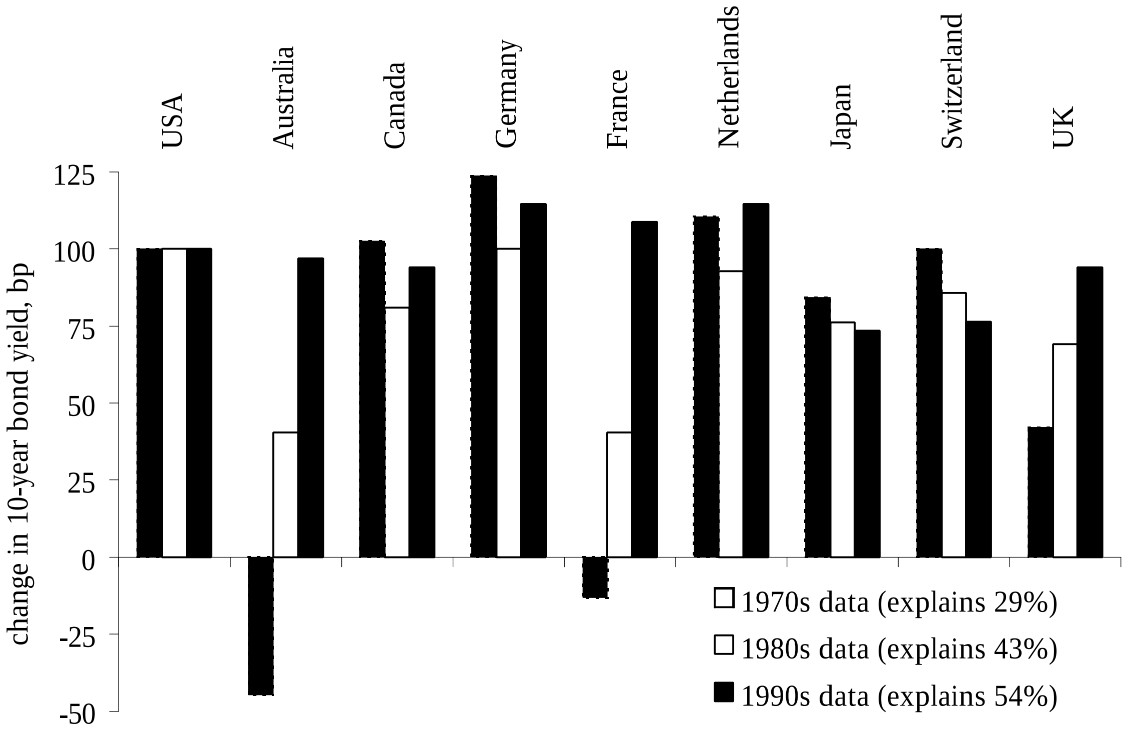

EXHIBIT 13 ■ Historical 12-month correlations between 10-year bond yields

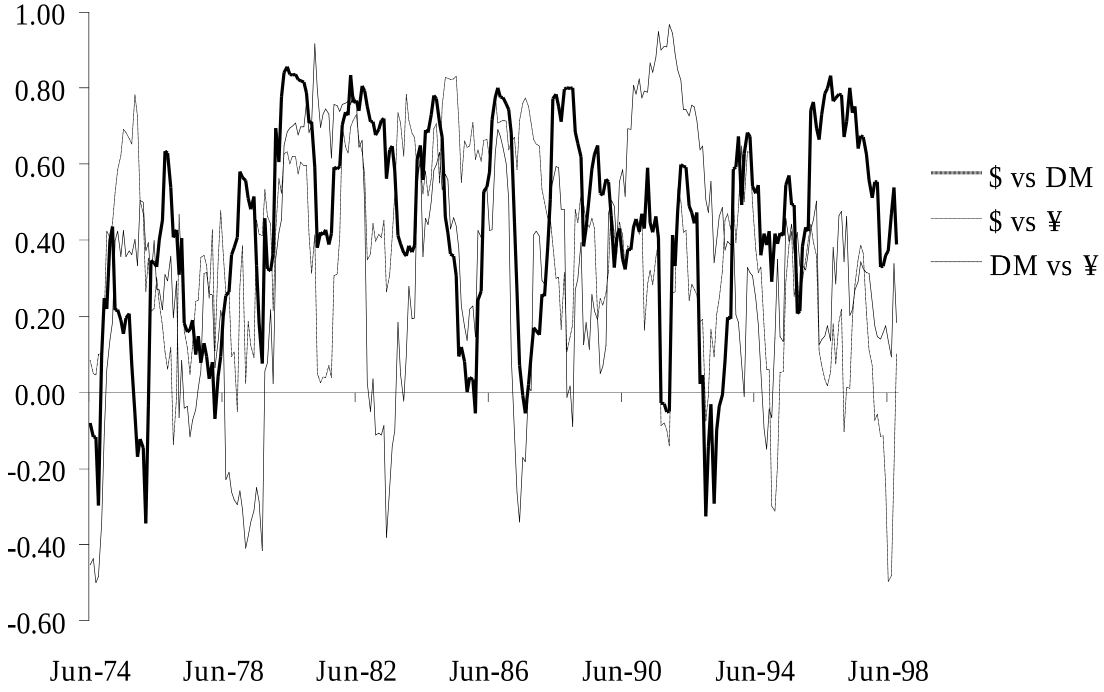

EXHIBIT 14a ■ US yield / volatility relationship

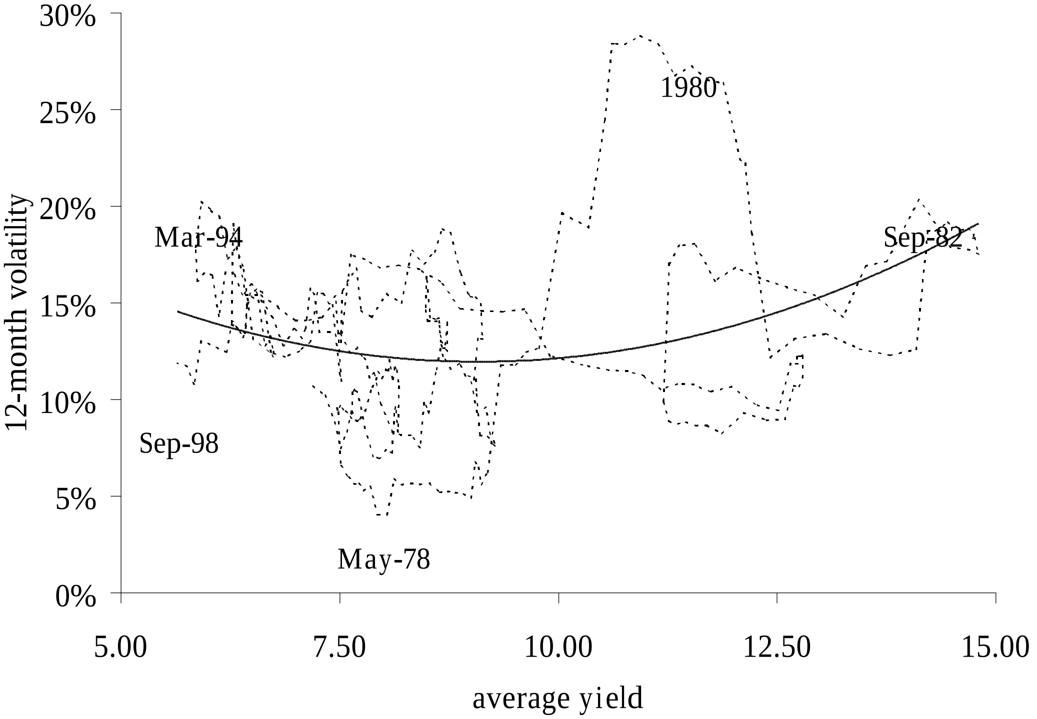

EXHIBIT 14b ■ Japan yield / volatility relationship

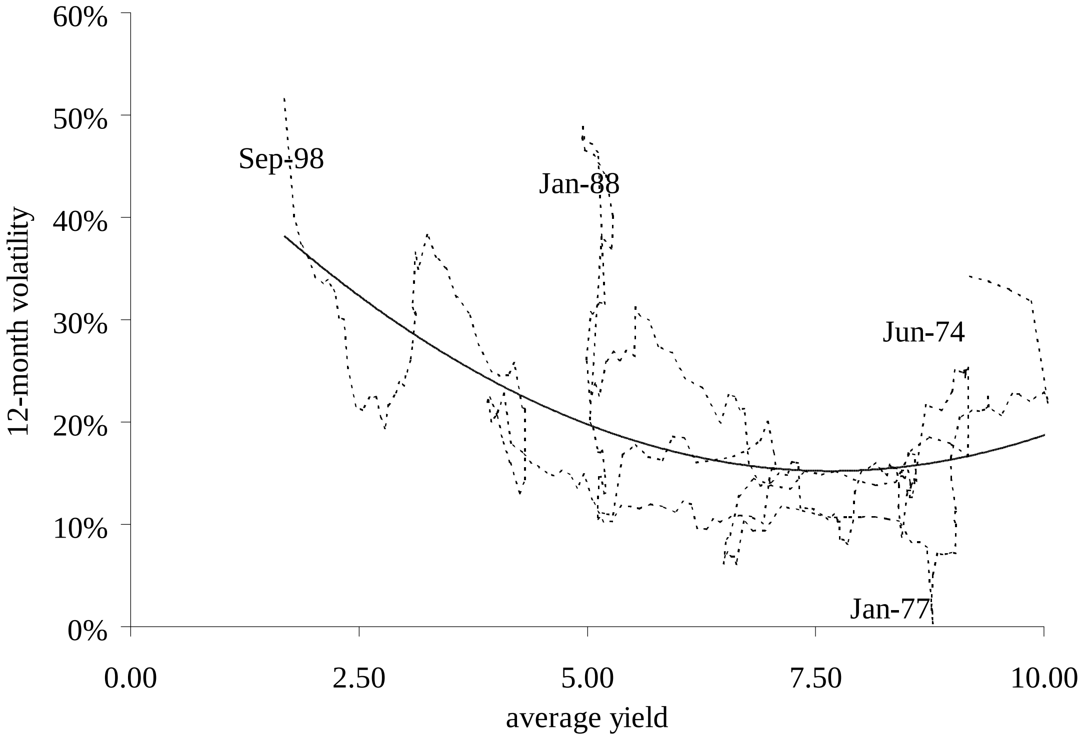

EXHIBIT 14c ■ Germany yield/volatility relationship

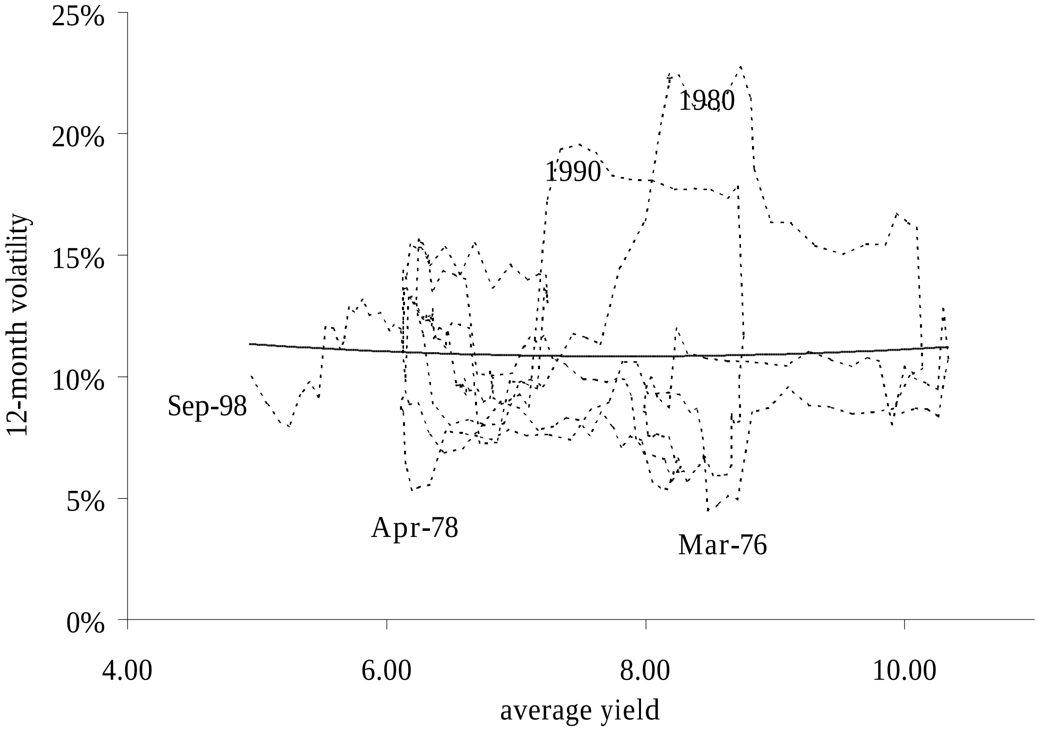

EXHIBIT 15 ■ Simplified value-at-risk calculation using principal components

Definitions

d_i	"duration" relative to factor i	$\propto$ duration $\cdot$ (factor i shift)
v_i	variance of factor i	$\propto$ factor weight
s_i	bp volatility of factor i	$= v_i^{1/2}$
VaR_i	value-at-risk due to factor i	$\propto s_i \cdot d_i \cdot \sqrt{T}$
VaR	aggregate value-at-risk	$= \left(\sum_i \text{VaR}_i^2 \right)^{1/2}$

Long portfolio	$100m	10-year par bond
Steepener portfolio	$131m	2-year par bond
	–$31m	10-year par bond
Butterfly portfolio	$64m	cash
	$100m	2-year par bond
	–$93m	5-year par bond
	$29m	10-year par bond

Calculations

Assume 100 bp p.a. 'total volatility', factors and factor weights as in Exhibit 3. Ignore all but the first three factors (those shown in Exhibits 5a–c).

		Parallel	*Slope*	*Curvature*	*1 s.d. risk*	*Daily VaR*
Long	10yr durn	7.79	1.50	–0.92		
	Total durn	7.79	1.50	–0.92		
	Risk (VaR)	5.75%	0.27%	–0.07%	5.95%	$376,030
Steepener	2yr durn	2.39	–0.87	–0.72		
	10yr durn	–2.39	–0.46	0.28		
	Total durn	0.00	–1.33	–0.44		
	Risk (VaR)	0.00	–0.24	–0.04	0.28%	$17,485
Butterfly	Cash durn	0.00	0.00	0.00		
	2yr durn	1.83	–0.67	–0.55		
	5yr durn	–4.08	0.24	1.20		
	10yr durn	2.25	0.43	–0.26		
	Total durn	0.00	0.00	0.38		
	Risk (VaR)	0.00%	0.00%	0.03%	0.03%	$1,954

EXHIBIT 16a ■ Global bond efficient frontier and hedged index, 1970s

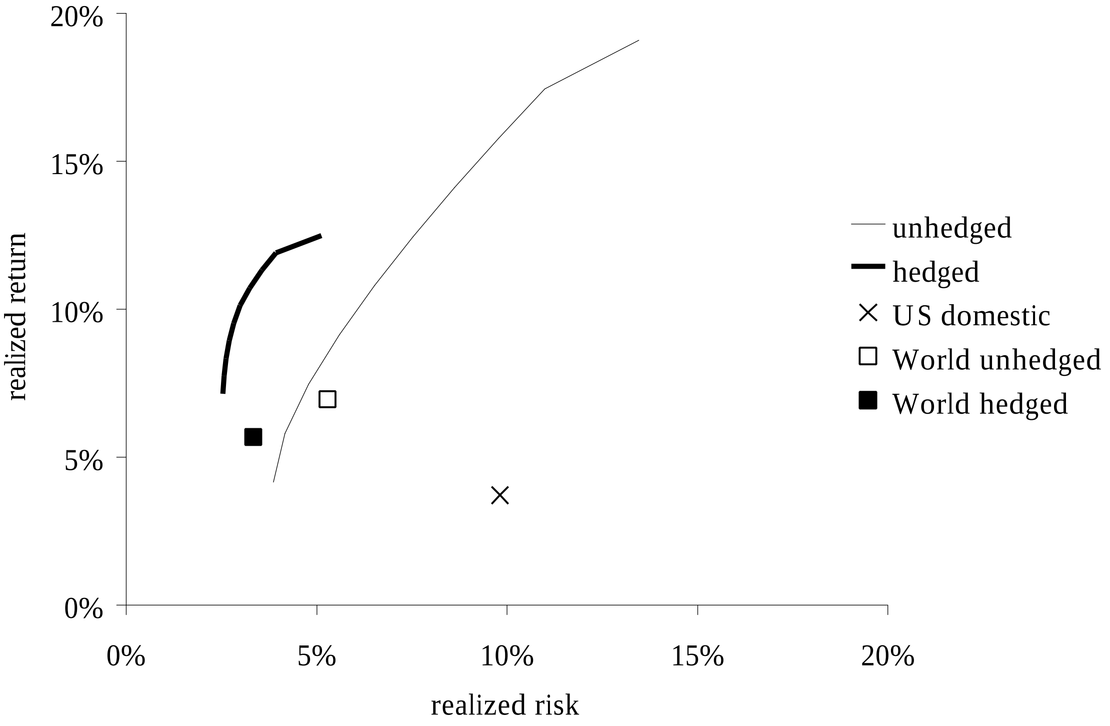

[Historical bond and FX data provided by Deutsche Bank Securities.]

EXHIBIT 16b ■ Global bond efficient frontier and hedged index, 1980s

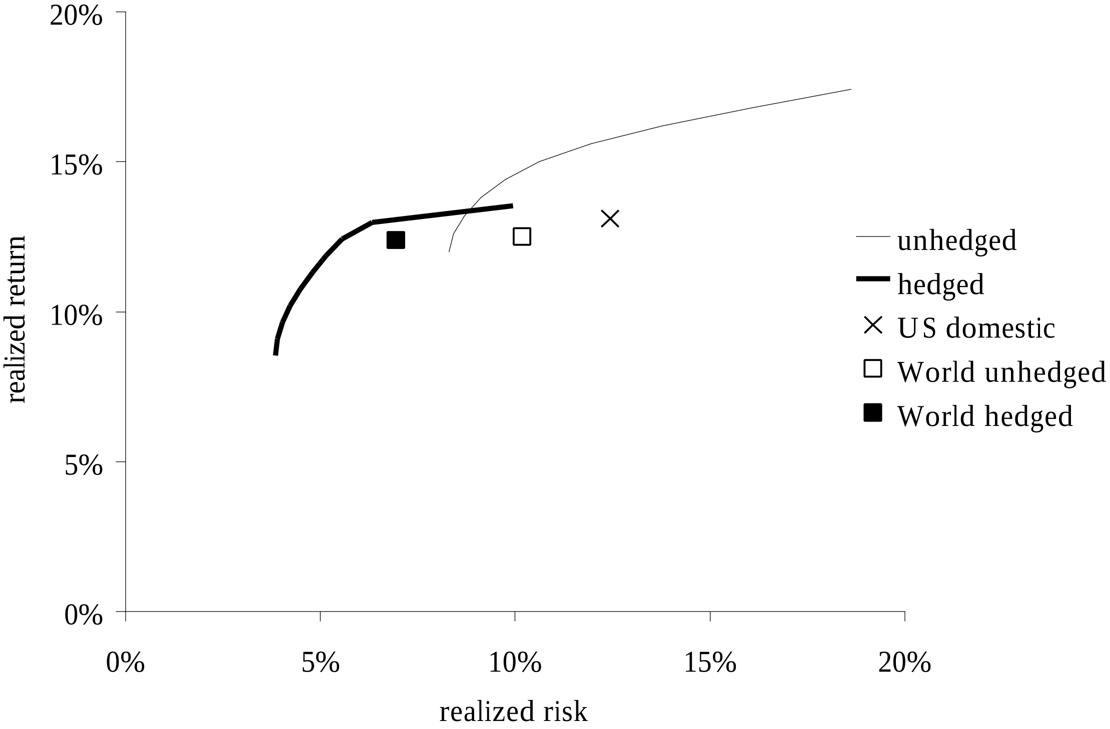

EXHIBIT 16c ■ Global bond efficient frontier and hedged index, 1990s

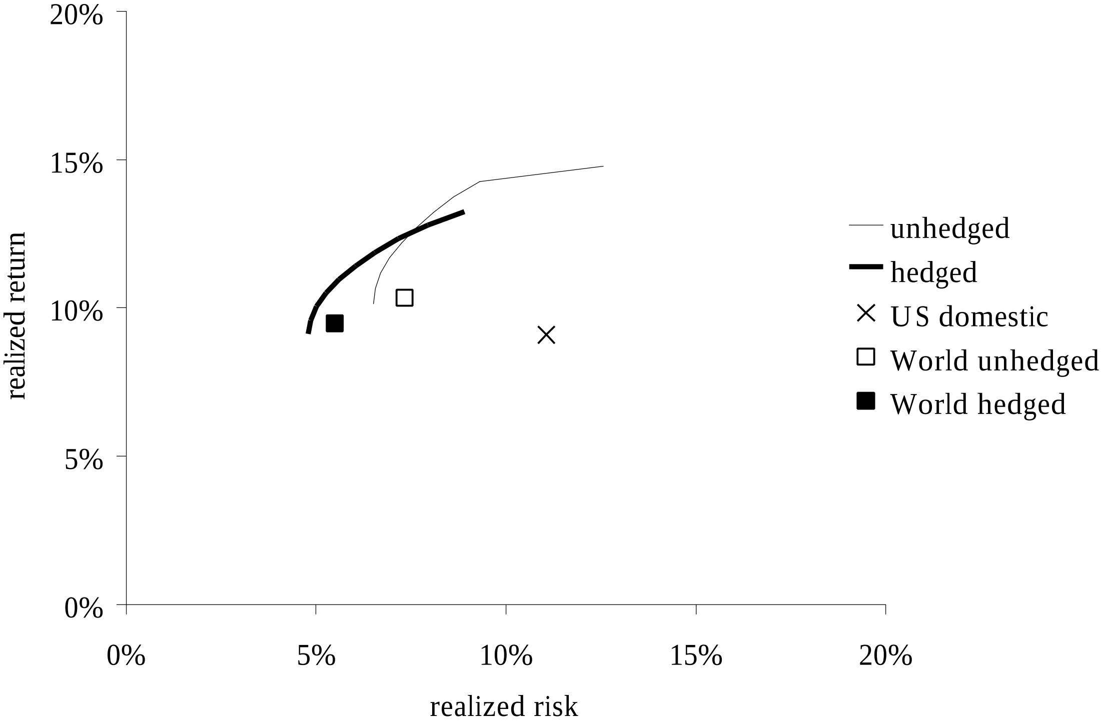

EXHIBIT 17 ■ A macroeconomic model of interest rate expectations

Model definitions:

i	short-term nominal interest rate
p^e	expected long-term inflation rate
r^e	expected long-term real interest rate
y	log of output
$\bar{y}$	log of normal or potential output
m	log of the money supply
p	log of the price level
g, f, l, r	constant model parameters (elasticities)

Model assumptions:

The output gap is related to the current real interest rate through investment demand:

$$y - \bar{y} = -g\left(i - p^e - r^e\right)$$

Real money demand depends positively on income and negatively on the interest rate:

$$m - p = f\,y - l\,i$$

Price changes are determined by excess demand and expected long-term inflation:

$$\frac{dp}{dt} = r\left(y - \bar{y}\right) + p^e$$

Theorem: (Frankel) The expected rate of change of the interest rate is given by:

$$\frac{di}{dt} = -d\left(i - p^e - r^e\right), \text{ where } d = \frac{r\,g}{f\,g + l}\,.$$

EXHIBIT 18 ■ Schematic breakdown of interest rate expectations

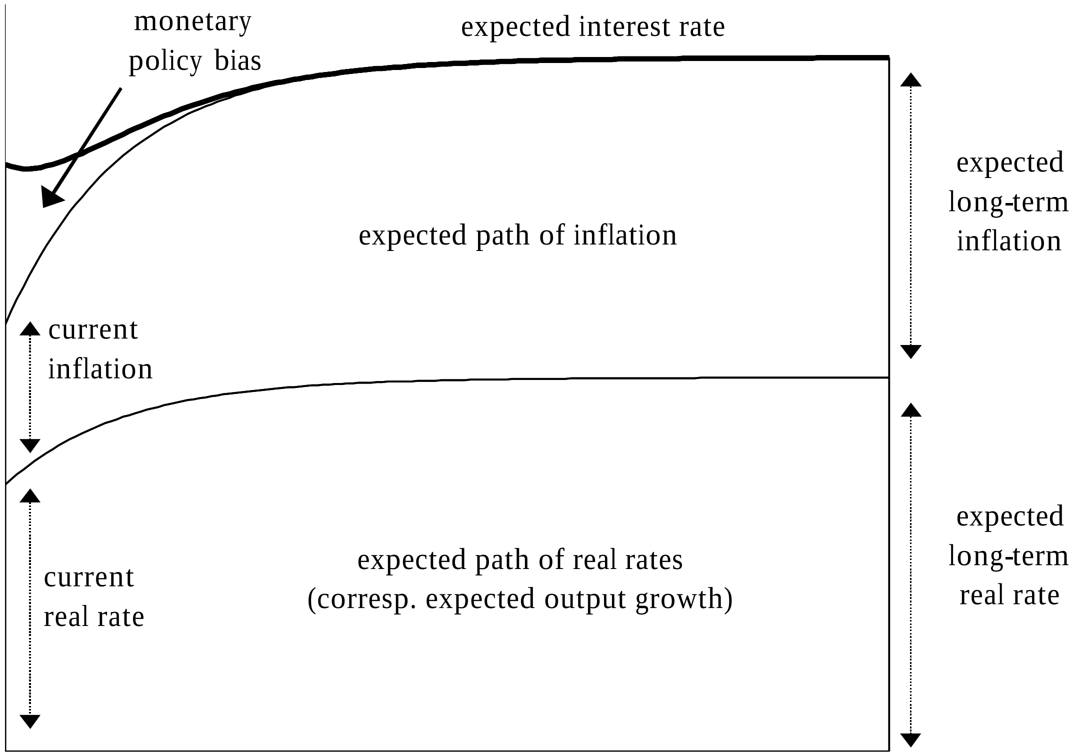

EXHIBIT 19 ■ Curvature shift arising from changing volatility expectations