Government of Canada Yield-Curve Dynamics, 1986–2003

Grahame Johnson, Financial Markets Department

- *In this article, the author introduces a historical database of daily constant-maturity Government of Canada zero-coupon yield curves. The behaviour and evolution of these yield curves is examined both over the full period of 1986 to 2003 and two subperiods, 1986 to 1996, and 1997 to 2003.*
- *The differences between the actual market prices of government bonds and the price generated by the model used to build the yield curves decreased materially over the term of the database, indicating more consistent pricing of the individual cash flows over time.*
- *The behaviour of government bond yields was found to be significantly different in the second subperiod. By almost any measure, the government bond market appears to have become a "safer" market during this period, exhibiting sharply lower levels of volatility. While the level of excess returns earned was slightly lower in the second subperiod, the volatility of those returns fell much more sharply, resulting in superior risk-adjusted returns.*
- *Variations in the yield curve over time could be almost completely explained by three independent factors—level, slope, and curvature. The relative importance of these factors changed significantly between subperiods.*
- *Contrary to the assumptions implicit in a number of pricing and risk-management models, none of the yield-curve measures exhibited daily changes that fit a normal distribution. For all measures, the distribution had both a larger number of observations close to the mean and a larger number of extreme outliers than would occur under the assumption of normality.*

his article[1] introduces a database of historical Government of Canada zero-coupon yield curves that has been developed at the Bank of Canada. Zero-coupon interest rates (or spot rates) represent the fundamental building blocks of fixed-income markets. Defined as the interest rate on a single cash flow maturing at a given time (with no interest payments prior to maturity), these rates are used extensively to price, structure, and hedge a variety of financial products. For most of these applications, it is appropriate to use spot interest rates that have been derived from the interest rates on bank-deposit contracts and the fixed leg of interest rate swap contracts (referred to as the swap curve). The generation of spot rates from this yield curve is fairly straightforward, as each maturity point on the curve has a specific and unique yield associated with it.[2] The resulting spot rates, however, represent a chartered bank credit, and as such are not truly free of credit risk. There are a large number of applications that require a spot rate that is free of credit risk as an input, making the use of spot rates derived from the swap yield-to-maturity curve unsuitable. For these applications, a term structure of zero-coupon interest rates needs to be generated from an underlying yield curve that is based on government bonds.

Generating a spot-rate curve from an underlying government bond yield-to-maturity curve is more problematic than generating it from the interest rate swap curve. The Canadian government bond market contains a large number of issues of varying coupon rates, maturities, and yields. Only a small proportion

---

1. This article is a summary of Bolder, Johnson, and Metzler (2004), which provides a much more technical discussion of the results.

2. For a detailed description of the construction of swap yield curves see Ron (2000)

of these issues are actively traded in the secondary market, and it has not been unusual for cash flows that fall on the same date to have different yields.[3] As a result, there is no single government bond yield for a specific date. This necessitates the use of numerical curve-fitting techniques. Essentially, these techniques use mathematical models to estimate a yield curve that minimizes the difference between the bond prices predicted by the model and those observed in the market. The specific estimation algorithm used to generate the historical government zero-coupon curves used in this article is based on the Merrill Lynch exponential spline model introduced in Li et al. (2001). In this model, a numerical curve-fitting technique is used to estimate a discount-function curve. The discount-function curve is then used to price all of the government bonds in the sample, and the final curve is the one that generates the smallest differences between the model prices and the observed market prices. The final step is to transform the discount-function curve into a zero-coupon yield curve. The parameters of the curve are re-estimated each business day. Full details can be found in Bolder and Gusba (2002) and Bolder, Johnson, and Metzler (2004).

This curve-fitting algorithm was used to build a database of historical constant-maturity zero-coupon yield curves.

This curve-fitting algorithm was used to build a database of historical constant-maturity zero-coupon yield curves using daily closing prices for Government of Canada bonds over the period January 1986 to May 2003. While historical term-structure databases (such as McCulloch and Kwon)[4] exist for the U.S. Treasury market, to the best of our knowledge, this represents the first such database for Canadian zero-coupon rates in the public domain. The database

will be updated on a regular basis and is available to the public on the Bank of Canada website.

The purpose of this article is both to highlight the availability of these data and to perform some initial statistical analysis, examining the behaviour and evolution of the spot rates over the full period and two distinct subperiods. The article will specifically address the following areas:

- the evolution of the levels of key interest rates and yield-curve measures over the sample

- the daily changes in these key interest rates and yield-curve measures

- the identification of a relatively small number of factors that drove the evolution of the yield curve, and

- the total returns that would have been realized by holding bonds of different maturities for a given holding period.

The Sample Period

The full data sample used for the analysis in this article covers the period between January 1986 and May 2003, and comprises over 4,300 daily yield curves. Both the Canadian economy and the Canadian fixed-income markets underwent significant changes over the period covered by the database, however; and the full sample of over seventeen years can be thought of as being composed of at least two distinct subperiods. The first subperiod, covering the time from January 1986 to December 1996, can be characterized as follows:

- relatively high and volatile inflation, particularly for the first half of the subperiod; longer-term inflation expectations were also high and volatile, adjusting to the reduction in measured inflation with only a lag

- large government borrowing requirements

- a fragmented bond market, characterized by a large number of relatively small and illiquid issues, and

- restrictions on the stripping and reconstituting of individual cash flows from the underlying bonds.

The second subperiod, from January 1997 to the end of May 2003, had very different conditions. Specifically:

---

3. Cash flows often have different yields, depending on whether they represent an interest payment (the coupon) or the principal repayment (the residual).

4. The database on term structures for U.S. Treasury bonds is available on J.J. McCulloch's website at <http://www.econ.ohio-state.edu/jhm/jhm.html>.

- Inflation and inflation expectations were low and stable.

- Beginning in 1996, the Government of Canada began to run a sequence of budgetary surpluses, sharply reducing net government borrowing requirements.

- Numerous steps were taken by the Department of Finance and the Bank of Canada to help make the government bond market more efficient. These included the introduction of an official benchmark-bond program with explicit issuance targets and regular issuance patterns, regular and formal consultations with market participants to discuss potential changes to the debt program, and the implementation of a bond buyback program.

- The Canadian Depository for Securities (CDS) implemented several initiatives that helped to increase the efficiency of the bond market. In 1993, reconstituted, packaged, and generic CUSIPS[5] were introduced for book-entry strip bonds. This enabled individual coupon payments with the same maturity date to be fully fungible, allowing for increased arbitrage between rich and cheap bond issues. In 1999, any cash flow of similar type[6] that shared a maturity date became fully fungible, and in 2001 it became possible to reconstitute a bond beyond its original issue size. These developments ensured that cash flows having the same issuer and maturity were valued identically, regardless of the underlying issue from which they originated.

- Computerized trading strategies and quantitative valuation approaches gained increased popularity in financial markets. Hedge funds, many of which specialize in fixed-income relative-value arbitrage, also became much more significant factors in the fixed-income market.

---

5. CUSIP stands for Committee on Uniform Securities Identification Procedures. The CUSIP reference is a combination of nine characters, both letters and numbers, that represent a unique identifier for a given security. The first six characters identify the issuer, the next two identify the type of security, and the final digit is a check.

6. Fungible cash flows have to be either interest payments or principal payments. Interest and principal payments are not yet fungible with each other.

This article will effectively account for two different regime shifts. The first is a fiscal and macroeconomic shift, highlighted by the achievement of low inflation and a balanced fiscal position. The second is a shift in the operation of the actual fixed-income markets themselves, including changes in the issuance pattern, changes made by CDS, and the growing importance of quantitative trading strategies. No specific date marks a perfect break between these two regime shifts. The selection of January 1997 as the date for the break between the two regimes is somewhat arbitrary, and all of the changes highlighted above actually took place either before or after that date. The main point, however, is that the period between the late 1980s and the early 1990s had very different characteristics from the late 1990s and early 2000s, and by the beginning of 1997, most of those changes were evident.

This article will effectively account for two different regime shifts.

The Yield-Curve Model

The database of historical yield curves was built using a mathematical curve-fitting model that finds a specific zero-coupon yield curve on a given day that minimizes the difference between bond prices predicted by the model and those actually observed in the market. An examination of the size of the pricing errors generated by these "best-fit" curves offers some insight into how consistent the pricing of government bonds was at a given time. If the full universe of government securities were priced on a consistent basis, with cash flows of a similar term to maturity trading at similar yields (regardless of which specific bond issue those cash flows were associated with), then the model should produce a very accurate fit, with little pricing error. If, on the other hand, individual bond issues exhibited idiosyncratic pricing, with cash flows of similar maturities trading at significantly different yields, then the pricing errors produced by the model should be relatively large. Furthermore, given the evolution of market conditions outlined in the previous section, it would be reasonable to expect a reduction in the pricing error over time, since a

Chart 1

Average Pricing Error

RMSE (bps)
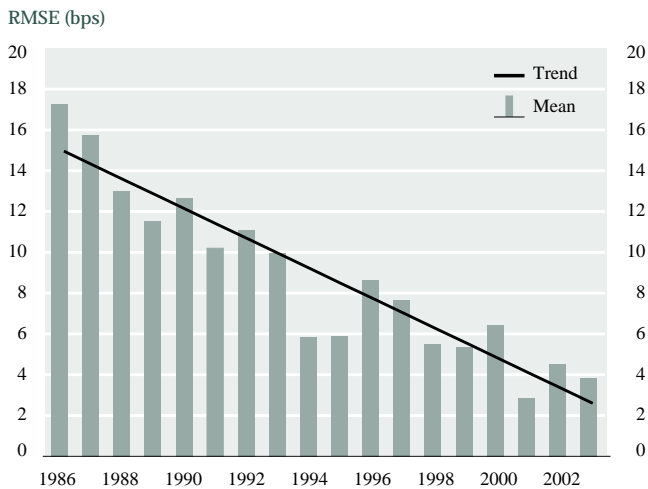
number of changes were introduced to make the government bond market more efficient.

Chart 1 shows the evolution of the goodness-of-fit of the model.[7] The root mean-square error for each year is shown, along with a time trend.

Individual cash flows of similar maturity are now valued in a more consistent fashion.

The goodness-of-fit of the model has clearly increased substantially over time. The average error over the first subperiod was 11.1 basis points, while it fell to 5.2 basis points for the second subperiod. Bond yields for the full universe of Government of Canada securities are now far more consistent with the prices predicted by a theoretical yield curve. This indicates that individual cash flows of similar maturity are now valued in a more consistent fashion, regardless of the specific bond from which the cash flow originated, or whether the cash flow represents an interest payment or a return of principal. While it is difficult to assign causal-

ity, changes of this nature are consistent with both the initiatives undertaken by the CDS to make similar cash flows fully interchangeable and the increasing prevalence of quantitative valuation methods and computerized trading strategies in fixed-income markets.

General Attributes of the Yield Curves

A first step in the analysis of the evolution of the yield curve is to examine what the "average" yield curve looked like, both over the full time period covered by the data set, and over each of the two subperiods. Charts 2 and 3 depict these yield curves (surrounded by confidence bands of plus and minus one standard deviation).

As Chart 2 shows, the average spot-yield curve over the entire period was upward sloping, with a 3-month interest rate of approximately 6.5 per cent and a 10-year interest rate of approximately 7.5 per cent. The variation around these averages, however, was extremely large. The one standard-deviation confidence band covers a range for the 3-month yield of approximately 3.5 per cent to 9.5 per cent, while for the 10-year rate, the band is between 5.5 per cent and 9.5 per cent.

Chart 3, which shows the average yield curves for the two subperiods, gives some indication of the magnitude of the changes in the shape, slope, and level of the yield curve over the two samples.

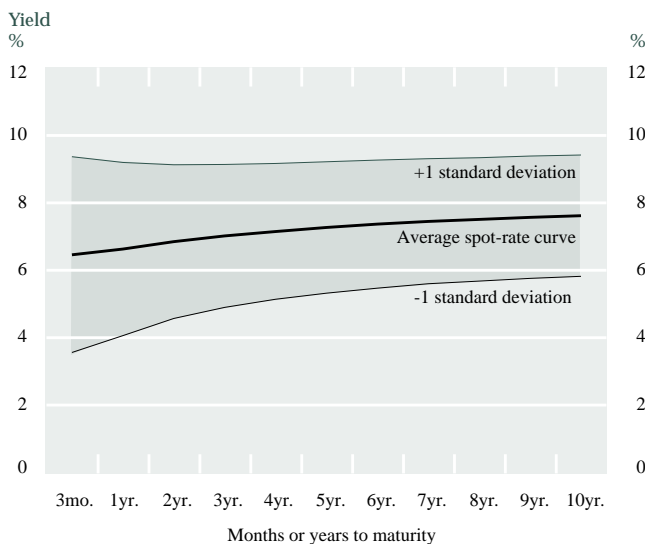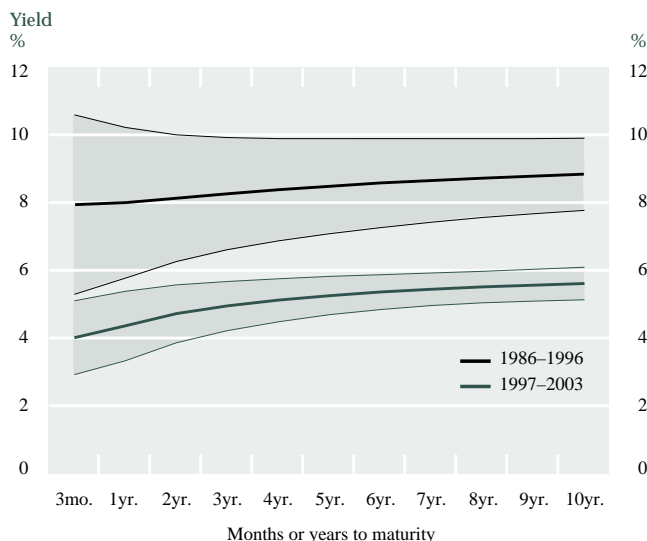**Chart 2**

Average Spot-Rate Curve, January 1986 to May 2003

Yield
%

Months or years to maturity

---

7. The graph plots the average daily root mean-square error for each year covered by the sample. The root mean-square error is equal to the square root of the sum of the squared yield differentials.

Chart 3

Average Spot-Rate Curve

Subperiods

Table 1

Summary Yield-Curve Statistics[a]

Yield-curve measure	Full sample			1986–96			1997–2003		
	Mean[b]	Min.	Max.	Mean[b]	Min.	Max.	Mean[b]	Min.	Max.
3-month yield	6.46 (2.9)	1.78	13.57	7.94 (2.8)	2.76	13.57	4.01 (1.1)	1.78	5.76
10-year yield	7.62 (1.8)	4.53	11.32	8.84 (1.1)	6.21	11.32	5.61 (0.5)	4.53	7.03
Slope	1.16 (1.7)	-3.21	4.08	0.90 (1.8)	-3.21	3.93	1.60 (1.2)	-35.37	407.47
Curvature	0.13 (0.20)	-0.47	0.82	0.09 (0.19)	-0.47	0.82	0.20 (0.1)	-0.20	0.73

a. All values are expressed as percentages.
b. Standard deviations of the measures appear in brackets below the means.

As Chart 3 shows, the two yield curves are significantly different. The pre-1997 yield curve is upward sloping, with an average 3-month rate of approximately 8 per cent and an average 10-year rate of about 9 per cent. The dispersion of yields over the period is very high, with the standard-deviation confidence band ranging from 5.5 per cent to 10.5 per cent for the 3-month rate and from 7.5 per cent to 10 per cent for the 10-year rate.

While the general shape of the post-1996 yield curve is similar to that of the curve from the first subperiod (both are upward sloping), there are two obvious differences. First, the general level of yields in the second subperiod was materially lower. So much lower, in fact, that the upper confidence band of the later subperiod is well below the lower confidence band of the first. Second, it appears that the dispersion of yields around the average curve is much narrower in the second subperiod.

While the average yield curves shown above graphically depict the differences in the term structure between the two periods, statistical tests were also conducted on a number of representative yield-curve measures to quantify the observed differences between subperiods. The specific measures examined were the 3-month yield, the 10-year yield, the slope of the yield curve, and the degree of curvature of the yield curve. The slope measure was defined as the difference between the 10-year yield and the 3-month yield, while curvature

was defined as the difference between the 6-year yield and the average of the 2-year and 10-year yields. The statistical tests confirmed what is shown in Charts 2 and 3. The level of yields, both 3-month and 10-year, were not only significantly lower in the post-1996 period, but they were also much less variable. The slope of the yield curve was found to be materially steeper in the second period, while the degree of curvature showed relatively little change. The full details are presented in Table 1.

Daily Yield Changes and Distributions

The previous section examined the appearance of average yield curves over the period, presenting some detail about the average values and variances of the levels for four different yield-curve measures. Of potentially more interest than the levels of these measures, however, is the behaviour of the daily changes in these levels, since it is these changes in yield that drive the short-term risk and return behaviour for government bonds. Almost all derivative-pricing algorithms, portfolio-management tools, and risk-measurement models make some assumptions about the distribution of returns over a given time horizon, the most common being that returns are normally distributed. Since, for a zero-coupon bond, short-term returns are almost entirely driven by yield changes, it follows that an assumption that returns are normally distributed is equivalent to assuming that, over short time horizons, yield changes also have a normal distribution. If this is, in fact, not the case, then any model that makes the assumption of normality

could be producing results that provide inaccurate prices or risk measures.

Table 2 presents the statistical details of the daily changes. Two key observations can be made from Table 2. Specifically:

- While the mean change for all measures was very close to zero, the uncertainty surrounding the changes for all measures was very large. As was evident in Table 1 (summary table of yield levels), however, this variance was significantly lower in the post-1996 period.

- All of the yield-curve measures had some extreme outliers, with every measure having several observations that were up to 12 to 13 standard deviations away from the mean. While the absolute magnitude of the outliers was much smaller in the second subsample, their distance from the mean as measured by standard deviations was very similar.

It is this presence of extreme outliers that suggests that the distribution of daily yield changes may not be normal. These extreme outliers were not one-time occurrences, but happened with some degree of regularity. Over the full sample period, each of the various measures saw between 30 and 50 observations fall outside of four standard deviations from the mean. Observations as extreme as those highlighted in Table 2 are practically statistical impossibilities under the normal distribution.[8]

Chart 4 takes this analysis a step further, plotting the distribution graphs of the daily yield changes for each yield-curve measure compared with a theoretical normal distribution. This provides further evidence that the daily yield changes do not appear to be normally distributed.

For simplicity of presentation, the various yield-curve measures are not broken up into subperiods, and only the distributions for the full samples are plotted. The characteristics of the distributions in the subperiods are consistent, however, with the distribution of the period as a whole, with no material change between the first and second time periods. As Chart 4 shows, not only do the distributions contain some extreme outliers, they also have far more observations close to the mean than would be expected under the normal distribution. These two characteristics (the presence of a relatively large number of extreme outliers and observations that are heavily centred around the mean) represent clear evidence of non-normality, and formal statistical tests firmly reject the hypothesis that the distributions are normal.[9]

These two characteristics (the presence of a relatively large number of extreme outliers and observations that are heavily centred around the mean) represent clear evidence of non-normality.

The historical shapes of these distributions have some interesting repercussions for pricing algorithms, portfolio-management models, and risk measures that rely on the underlying assumption of normally distributed yield changes. These models would have systematically underpredicted the probability of a very small change in yields (or overstated the risk), while at the same time they would also have underpredicted the probability of a very large change in yields (or understated the risk). Options markets do, however, appear to compensate for at least part of this pattern by pricing options with various strike prices using different implied volatility levels. Options with strike prices that are further away from the current price trade with a higher

Table 2

Summary Yield-Curve Measures: Daily Changes[a]

Yield-curve measure	Full sample			1986–96			1997–2003		
	Mean[b]	Min.	Max.	Mean[b]	Min.	Max.	Mean[b]	Min.	Max.
3-month yield	-0.15 (14.5)	-120.6	188.3	-0.25 (17.3)	-120.7	188.3	0.03 (7.9)	-51.5	70.8
10-year yield	-0.12 (7.3)	-92.1	62.1	-0.11 (8.3)	-92.1	62.1	-0.13 (5.2)	-22.4	23.5
Slope	0.03 (15.0)	-176.3	93.1	0.14 (17.7)	-176.3	93.1	-0.16 (8.9)	-76.8	53.4
Curvature	0.00 (4.8)	-50.8	66.7	0.01 (5.8)	-50.8	66.7	-0.01 (2.4)	-33.3	32.7

a. All values are in basis points (one-hundredth of a per cent).
b. Standard deviations for the measures appear in brackets below the means.

---

8. For example, an observation that is 12 standard deviations from the mean would only be expected to happen once every 10^{20} years under the normal distribution.

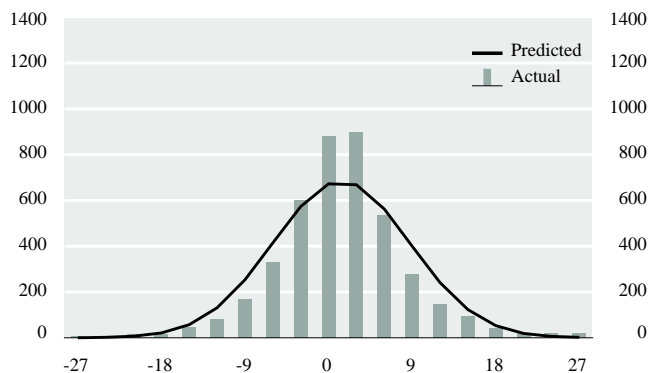9. Jarque-Bera probabilities for all distributions are 0.0000.

Chart 4

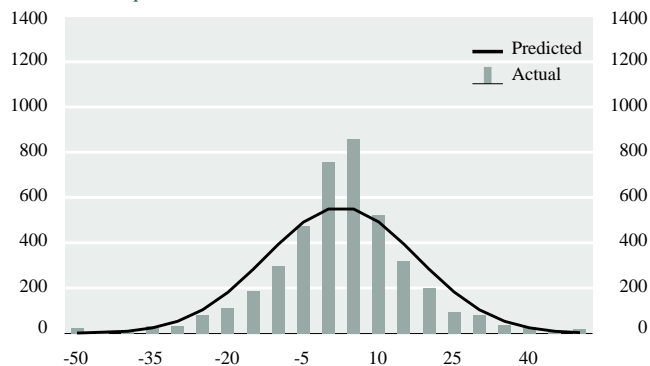**Probability Distributions of Daily Yield Changes**

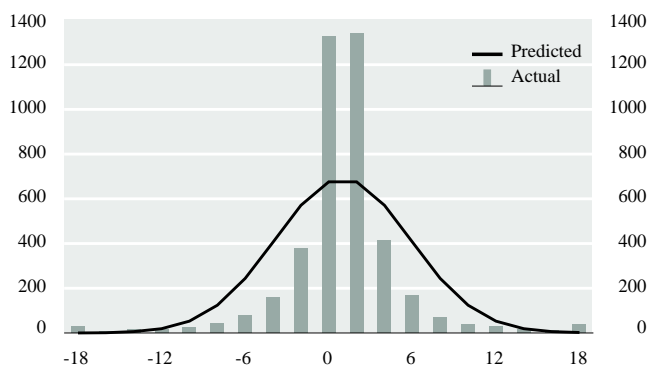3-month yields, first differences
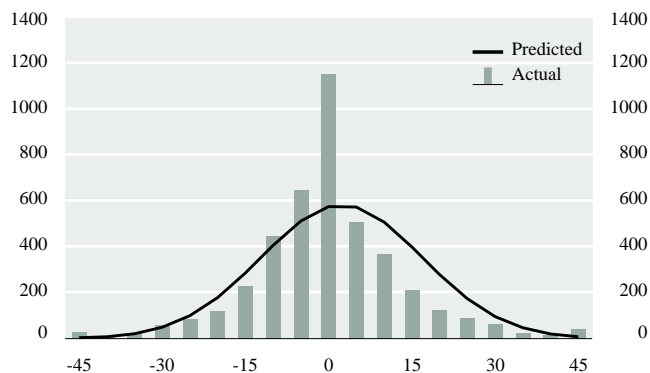
10-year yields, first differences

Yield-curve slope, first differences

Curvature measure, first differences

implied volatility than do options with strike prices very close to the current price. This, in effect, compensates for the fact that the deep out-of-the-money options are more likely to be exercised than the standard normal-distribution assumptions of some option-pricing models would indicate. Nonetheless, it remains an interesting question as to whether specific trading strategies that were structured to benefit from the tendency of yields to move either very little or very much (relative to a normal distribution) would have been abnormally profitable.[10]

Principal-Component Analysis

Principal-component analysis attempts to describe the behaviour of a range of correlated random variables (in this case, the various spot yields for different times to maturity) in terms of a small number of

uncorrelated principal components. This type of analysis makes it possible to identify a relatively small number of factors that have affected the behaviour of the entire zero-coupon curve over the period examined. While there are many practical applications for this analysis, the ability to construct more complete hedges for a bond portfolio is likely one of the most important. Once a small number of key factors that explain almost all of the risk faced by a given bond portfolio are determined, it is possible to create a replicating portfolio that immunizes the original against any shocks driven by those key factors. For anything other than a very simple bond portfolio, this would provide superior protection compared with more naive immunization approaches, such as simple duration matching.

This approach was first applied to bond yields by Litterman and Scheinkman (1991), who found three common factors that influenced the returns on all treasury bonds over the period they examined, explaining, on average, 98.4 per cent of the variance

---

10. An example of such a strategy would have been to maintain a short position in bond options with strike prices close to the current market price, while holding an offsetting long position in deep out-of-the-money options.

in yields. The first factor, which they called *level*, represented an approximately parallel shift higher or lower in the yield curve. A shock to this factor raised or lowered all yields by roughly the same amount. Level was by far the most important factor, accounting for 89.5 per cent of the total explained variance. The second factor was called *steepness*, since a positive shock to this factor lowered short-term spot rates, while raising longer-term rates. This factor was found to account for 8.5 per cent of total explained variance. A positive shock to the third factor, which they called *curvature*, lowered both short- and long-term yields, while raising mid-term yields. This had the effect of increasing the degree of curvature in the term structure. The curvature factor accounted for 2.0 per cent of the explained variance. This model has been applied to other interest rate markets with similar results,[11] and it has become standard practice in finance to refer to shifts in yield curves as being driven by three underlying factors: level, slope, and curvature.

Table 3 shows the results of a principal-component analysis of the Canadian zero-coupon curves, both over the full period and for each of the two subperiods. Similar to the results of Litterman and Scheinkman, the first three factors were found to explain essentially all of the variation in the spot rates over the period. The first factor, level, was by far the most important. The proportion of total variance explained by this factor fell fairly sharply, however, in the second subperiod. The second most important factor, slope, saw a corresponding increase in the proportion of total variance it explained, having over three times as much explanatory power in the second subperiod as it did in the first. This shift suggests an important change in the dynamics of the yield curve over the full sample. A change in the absolute level of interest rates

Table 3

Percentage Variation Explained

Period	Factor 1 Level	Factor 2 Slope	Factor 3 Curvature	Total
1986–96	89.8	8.4	1.3	99.6
1997–2003	72.6	25.8	1.4	99.7
Full sample	83.1	15.2	1.4	99.6

11. For example, Buhler and Zimmermann (1996) find similar results for the Swiss and German markets, while Knez, Litterman, and Scheinkman (1994) proposed a three- and four-factor model to explain U.S. money market returns.

became a significantly less important risk factor in the second subperiod than in the first, while changes in the slope of the curve became significantly more important. While outright changes in the level of yields still constituted the dominant risk factor, the risk associated with changes in the slope of the curve tripled in the second subperiod. Finally, the curvature factor explained a relatively small, but constant, degree of variation in both subperiods.

The first three factors were found to explain essentially all of the variation in the spot rates over the period.

Chart 5, which shows the impact of a shock to one of the factors on a sample yield curve, helps to explain the interpretation of the various factors and demonstrates why they were named level, slope, and curvature.

All three factors behave very similarly to those described in Litterman and Scheinkman. The level factor corresponds to a roughly parallel shift in yields across the term structure; a positive shock to the slope factor lowers short-term rates, while increasing longer-term rates; while a positive shock to the curvature factor lowers very short- and long-term rates but increases mid-term yields.

Traditional interest rate risk management emphasizes duration matching, immunizing a portfolio against parallel shifts in yields. This type of hedging strategy effectively only immunizes the portfolio against shifts in the first factor (level). While this clearly represents the largest source of variation in yields, it ignores a substantial amount of risk (only about 73 per cent of the total variation in yields would have been immunized against in the second subperiod). A more complete hedging strategy would be to calculate the sensitivity of a bond portfolio to all three factors. Once these sensitivities are known, it is possible to construct a hedging portfolio that offsets the exposure to these factors. Properly constructed, a hedging strategy based on all three factors would protect against over 99 per cent of the variability in the term structure. This type of hedging strategy has been the subject of a relatively large amount of research, and additional information can be found in Barber and Copper
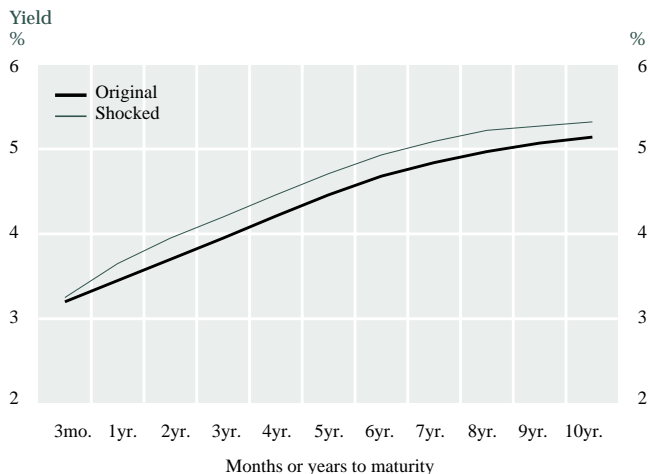
Chart 5

Effect of Shocks to Individual Factors

Level

Yield
%

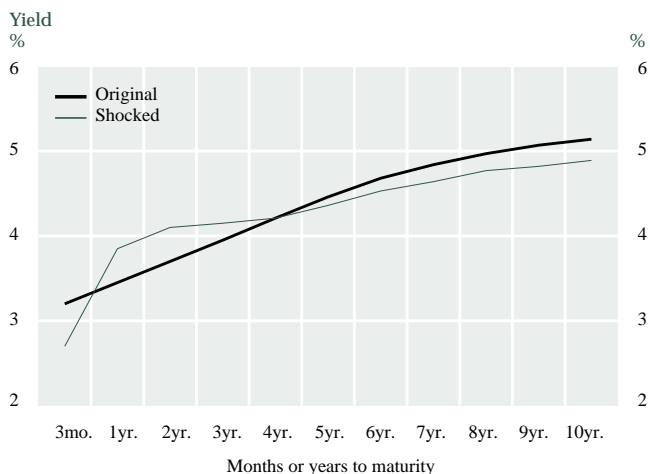Steepness

Yield
%

Curvature

Yield
%

(1996), Golub and Tilman (1997), and Lardic, Priaulet, and Priaulet (2003).

Holding-Period Returns

The expectations hypothesis on the term structure of interest rates maintains that longer-term interest rates are simply the average of expected future short-term rates plus a risk premium. If the value of this risk premium is set to zero, it follows that the expected returns on zero-coupon bonds of all maturities will be equal over a given time horizon (e.g., purchasing a 10-year bond and selling it within one year provides the same expected return as holding a 1-year bond to maturity).[12] If the term premium is positive, longer-term bonds will have a higher expected return over a given horizon than shorter-term bonds. This extra expected return represents compensation for the additional risk associated with the longer-term bonds, since its actual realized return is uncertain.

The historical yield-curve series can be used to examine a number of questions related to this hypothesis.[13] First, the data series was evaluated to see if bonds of different maturities provided equivalent returns, on average, for a given holding period over the sample, or if certain maturity assets earned some measure of excess returns. Second, were the returns earned from holding longer-term instruments riskier (or more variable) than they were for shorter-term bonds? Third, if the risk level did vary across maturities, did some maturities consistently produce better risk-adjusted returns than others? Finally, results were compared across the subperiods to see if the risk-return profile across the yield curve changed.

Holding-period returns (HPRs) are defined as the total return earned by purchasing an asset at the beginning of the period, holding it for a given time horizon, and either redeeming it at maturity for a known amount, or selling it in the secondary market. HPRs were calculated using zero-coupon bonds with maturities of one, two, five, and 10 years for a holding period of 180 days. The difference between these returns and the risk-

---

12. The expected capital gain or loss from selling the longer-term instrument at the end of the horizon would exactly offset the interest differential between the two maturities.

13. Any conclusions would only represent the behaviour of the yield curve over the time period examined. The sample size may not be large enough to draw broad-based conclusions.

free return that could have been earned by simply purchasing a 180-day bond and holding it to maturity is referred to as the excess HPR, and it is this result that is of most interest. The use of excess HPRs allows for an adjustment for changes in the level of the reference risk-free rate over the period. This is an important point in a sample that saw short-term yields range from 2 per cent to 14 per cent, as an HPR of 10 per cent (for example) could represent either a good or bad outcome, depending on the prevailing level of the risk-free rate.

Table 4 shows the summary results for HPRs across both the full period and the subsamples. Two main conclusions emerge from the results:

- Excess HPRs get both larger and more volatile as the term to maturity of the underlying bond increases. Longer-term bonds had higher levels of risk, but also provided higher levels of return. This observation holds for both the first and second subsamples.

- While excess HPRs appeared to be lower in the 1997 to 2003 period, the difference was not statistically significant. Volatility of returns, however, was significantly lower in the second subsample, with the standard deviation of excess returns generally only half the level it was for the first subsample.

It appears, therefore, that longer-maturity bonds carried a risk premium to compensate for their higher levels of risk (or variability of return), and that both return and return variability fell in the second subsample.

This analysis cannot tell, however, whether the incremental returns earned on the longer-dated assets were

sufficient to compensate for the additional variability of those returns. For this, it is necessary to examine the risk-adjusted returns for the various bond maturities.

One of the simplest methods for calculating risk-adjusted returns for different assets is to construct a ratio of the excess HPRs to the volatility of those returns, commonly known as a Sharpe ratio (Sharpe 1966, 1975). These ratios were calculated for both the full sample period and both subperiods, and the results appear in Table 5.

Table 5

Sharpe Ratios

Bond	Full sample			1986–96			1997–2003		
	Mean excess (%)	Std. dev. (%)	Sharpe ratio	Mean excess (%)	Std. dev. (%)	Sharpe ratio	Mean excess (%)	Std. dev. (%)	Sharpe ratio
1-year	0.61	1.18	0.52	0.69	1.40	0.49	0.51	0.70	0.73
2-year	1.53	3.57	0.43	1.69	4.37	0.39	1.32	2.20	0.60
5-year	3.17	8.35	0.38	3.46	10.17	0.34	2.79	5.35	0.52
10-year	4.89	14.48	0.34	5.24	17.48	0.30	4.45	9.82	0.45

Two main characteristics of the Sharpe ratios are apparent. First, the ratios decrease with the time to maturity of the bonds for all the sample periods. It would appear, therefore, that the incremental return earned by extending maturity did not compensate for the increase in risk. Second, Sharpe ratios for all maturities examined were significantly higher in the 1997 to 2003 sample, indicating a superior risk-reward trade-off in the second subperiod.

There is one overriding caveat to any conclusions that may be drawn from the analysis of excess HPRs, and that is that they are all based on *ex post* observations. The majority of the shocks that took place over the period resulted in yields falling further than could have been reasonably expected *ex ante* (as witnessed by the significantly lower yield levels in the second subperiod), resulting in large positive returns for longer-maturity fixed-income assets. As such, the *ex post* excess HPRs for the sample period are likely not indicative of what was expected *ex ante*, nor should they be seen as indicative of what should be expected in the future. In general, during periods of regime shifts, *ex post* observations are not good measures of what was (or should be) expected *ex ante*.

Table 4

Summary Statistics for 180-Day Excess Holding-Period Returns

Bond	Full sample (%)			1986–96 (%)			1997–2003 (%)		
	Mean[a]	Min.	Max.	Mean[a]	Min.	Max.	Mean[a]	Min.	Max.
1-year	0.61 (1.18)	-3.27	4.05	0.69 (1.40)	-3.27	4.05	0.51 (0.70)	-1.61	2.86
2-year	1.53 (3.57)	-11.10	12.53	1.69 (4.37)	-11.10	12.53	1.32 (2.20)	-3.57	8.13
5-year	3.17 (8.35)	-28.24	26.18	3.46 (10.17)	-28.24	26.18	2.79 (5.35)	-11.58	15.86
10-year	4.89 (14.48)	-49.66	38.18	5.24 (17.48)	-49.66	38.18	4.45 (9.82)	-21.92	31.79

a. Standard deviations for the measures appear in brackets below the means.

Conclusions

This article has presented a relatively high-level overview of the behaviour of the Government of Canada zero-coupon yield curve over a period of approximately 17.5 years. The analysis is based on what we believe to be the first historical constant-maturity Government of Canada zero-coupon yield-curve series to be publicly available. Four predominant conclusions can be drawn from the results presented here.

First, the differential between the actual market prices of government bonds and the price predicted by the yield-curve model decreased sharply over the period covered by the database. This is indicative of less idiosyncratic, more consistent pricing across different bond issues—that is, cash flows of similar maturity are priced at a similar yield, regardless of which specific bond they originated from.

Second, by almost any measure, the government bond market became a "safer" place during the latter part of the sample (1997 to 2003). While it is not possible to assign direct causality, the numerous changes in the fiscal and economic environment outlined earlier in this article coincided with a marked decrease in both the level and volatility of interest rates. Furthermore, while the level of excess returns earned for various bonds was slightly lower in the second subperiod, the volatility of those returns fell even more, resulting in superior risk-adjusted returns.

The third conclusion is that, similar to the other major bond markets, variations in the Government of Canada yield curve over the sample period could be almost totally explained by three factors—level, slope, and curvature. While the total proportion of variance explained remained very stable over the entire period (ranging from 99.0 per cent to 99.9 per cent), the breakdown of the three factors varied considerably.

Finally, none of the yield-curve measures examined had daily changes that fit a normal distribution. All of the distributions were characterized by both a much larger number of observations clustered around the mean and a much larger number of extreme outliers than would be expected under an assumption of normality. The behaviour of the yield curve over the period in question could be characterized as general stability punctuated by periods of extreme moves. This has implications for the large number of portfolio-management, risk-measurement, and derivative-pricing models that rely on an underlying assumption of normality in bond returns. That assumption has clearly not held up over time.

> The database of historical daily constant-maturity Government of Canada zero-coupon yield curves is available on the Bank of Canda website at <http://www.bankofcanada.ca/en/yield_curve.html>.

Literature Cited

Barber, J. and M. Copper. 1996. "Immunization Using Principal Component Analysis." *Journal of Portfolio Management* 23: 99–105.

Bolder, D. and S. Gusba. 2002. "Exponentials, Polynomials, and Fourier Series: More Yield Curve Modelling at the Bank of Canada." Bank of Canada Working Paper No. 2002–29.

Bolder, D., G. Johnson, and A. Metzler. 2004. "An Empirical Analysis of the Canadian Term Structure of Zero-Coupon Rates." Bank of Canada Working Paper No. 2004–48.

Buhler, A. and H. Zimmermann. 1996. "A Statistical Analysis of the Term Structure of Interest Rates in Switzerland and Germany." *Journal of Fixed Income* 6: 55–67.

Golub, B. and L. Tilman. 1997. "Measuring Yield Curve Risk Using Principal Components Analysis, Value at Risk, and Key Rate Durations." *Journal of Portfolio Management* 23: 72–84.

Knez, P., R. Litterman, and J. Scheinkman. 1994. "Explorations into Factors Explaining Money Market Returns." *Journal of Finance* (December): 1861–82.

Lardic, S., P. Priaulet, and S. Priaulet. 2003. "PCA of Yield Curve Dynamics: Questions of Methodologies." *Journal of Bond Trading & Management.* 1 : 327–49.

Li, B., E. DeWetering, G. Lucas, R. Brenner, and A. Shapiro. 2001. "Merrill Lynch Exponential Spline Model." Merrill Lynch Working Paper.

Literature Cited (cont'd)

Litterman, R. and J. Scheinkman. 1991. "Common Factors Affecting Bond Returns." *Journal of Fixed Income* 1: 54–61.

Ron, U. 2000. "A Practical Guide to Swap Curve Construction." Bank of Canada Working Paper No. 2000–17.

Sharpe, W. 1966. "Mutual Fund Performance." *Journal of Business* 39: 119–38.

———. 1975. "Adjusting for Risk in Portfolio Performance Measurement." *Journal of Portfolio Management* 1: 29–34.