

Bank of Canada



Banque du Canada

Working Paper 2004-48 / Document de travail 2004-48

An Empirical Analysis of the Canadian Term Structure of Zero-Coupon Interest Rates

by

David J. Bolder, Grahame Johnson, and Adam Metzler

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2004-48

December 2004

An Empirical Analysis of the Canadian Term Structure of Zero-Coupon Interest Rates

by

David J. Bolder, Grahame Johnson, and Adam Metzler

Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
gjohnson@bankofcanada.ca

The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements	iv
Abstract/Résumé	v
1. Introduction	1
2. The Data and Estimation Model	2
2.1 The estimation algorithm	2
2.2 Data filtering	5
2.3 Sample and subsample periods	5
2.4 Model fit	8
3. Summary of Descriptive Statistics	9
3.1 Average yield curves	9
3.2 Descriptive statistics—yield-curve levels	10
3.3 Descriptive statistics—first differences	12
4. Principal-Components Analysis	14
4.1 Introduction to principal-components analysis	14
4.2 Results	17
4.3 Potential applications	22
5. Holding-Period Returns	23
5.1 Definitions and notation	24
5.2 Summary results for holding-period yields—full sample	26
5.3 Summary results—subsamples	29
5.4 Risk-adjusted returns	31
6. Conclusions	32
Bibliography	34
Appendix: Data-Filtering Algorithm	40

Acknowledgements

The authors would like to thank Greg Bauer, Scott Gusba, Scott Hendry, Marianne Johnson, and Ron Morrow for suggestions, guidance, and discussion surrounding earlier drafts, as well as Glen Keenleyside for editorial assistance. Comments received during a departmental presentation of the preliminary results are also gratefully acknowledged.

Abstract

Zero-coupon interest rates are the fundamental building block of fixed-income mathematics, and as such have an extensive number of applications in both finance and economics. The risk-free government zero-coupon term structure is, however, not directly observable and needs to be generated from the prices of marketable, coupon-bearing bonds. The authors introduce the first public-domain database of constant-maturity zero-coupon yield curves for the Government of Canada bond market. They first outline the mechanics of the curve-fitting algorithm that underlie the model, and then perform some preliminary statistical analysis on the resulting yield curves. The full sample period extends from January 1986 to May 2003; it is broken down into two subsamples, reflecting the structural and macroeconomic changes that impacted the Canadian fixed-income markets over that time. The authors examine the evolution of a number of key interest rates and yield-curve measures over the period, perform a principal-components analysis of the common factors that have influenced yield changes over time, and compare holding-period returns over the sample for assets of various maturities.

JEL classification: C0, C6, E4, G1

Bank classification: Financial markets; Interest rates; Econometric and statistical methods

Résumé

Pierre angulaire du calcul pour les titres à revenu fixe, les taux de rendement coupon zéro trouvent un nombre imposant d'applications en finance et en économie. Toutefois, comme la courbe de rendement coupon zéro sans risque (celle des obligations d'État) n'est pas directement observable, elle doit être établie à partir des prix d'obligations négociables à coupons. Les auteurs présentent ici la première banque de données publique sur les courbes de rendement coupon zéro (échéance constante) des obligations du gouvernement canadien. Ils décrivent d'abord le fonctionnement de l'algorithme d'ajustement à la base de leur modèle, puis effectuent une analyse statistique préliminaire des courbes de rendement obtenues. Leur période d'estimation va de janvier 1986 à mai 2003 et a été divisée en deux afin de tenir compte des changements structurels et macroéconomiques qui ont influé sur les marchés canadiens de titres à revenu fixe durant cet intervalle. Les auteurs examinent ensuite l'évolution de plusieurs taux d'intérêt clés et de certaines caractéristiques des courbes de rendement estimées avant de faire une analyse en composantes principales des facteurs communs ayant façonné le comportement des rendements au fil du temps. Ils comparent aussi les rendements, sur la durée de détention, d'actifs de diverses échéances au cours de l'ensemble de la période considérée.

Classification JEL : C0, C6, E4, G1

Classification de la Banque : Marchés financiers; Taux d'intérêt; Méthodes économétriques et statistiques

1. Introduction

In this paper, we introduce a comprehensive database of constant-maturity zero-coupon yield curves for the Government of Canada bond market. The database provides best-fit zero-coupon curves based on historical bond closes, beginning in January 1986. It will be kept current and will be publicly available on the Bank of Canada's website. The first part of the paper reviews the underlying model used to generate these yield curves. The second part of the paper provides a preliminary statistical analysis of the data. The sample period extends from January 1986 to May 2003, a period of almost 17.5 years. We seek to analyze how the yield curve has evolved over the period under examination. Specifically, we

- examine the evolution of the level of key interest rates and yield-curve measures over time, including the distributional properties of those levels;
- examine the first differences (or daily changes) of these key interest rates and yield-curve measures, again including the distributional properties;
- perform a principal-components analysis of the common factors that have influenced the shape of the yield curve over time; and,
- examine the holding-period returns for bonds of various maturities.

Furthermore, given the significant changes that have occurred in both the macroeconomic environment and the fixed-income market over the full term of the database, we break down the full sample period into two subsample periods.

The motivation behind this work is straightforward. Zero-coupon interest rates (or spot rates) are the fundamental building block of fixed-income mathematics. These rates are used in a tremendous number of applications in both finance and economics, including bond pricing, the construction and pricing of derivative products, the generation of forward curves, estimations of inflation premiums, and modelling of the business cycle. While most financial engineering is done using zero-coupon rates generated from deposit contracts and interest rate swap rates, these rates contain a time-varying credit component that complicates matters somewhat. For any applications that require the use of risk-free interest rates, it is necessary to use a zero-coupon curve that has been constructed from government bond yields.

While the potential uses of a government zero-coupon yield curve are extensive, the estimation of such a curve is much less straightforward than it is for the interest rate swap market. By definition, the interest rate swap yield curve is a current-coupon, constant-maturity curve. The interest rate swap market has a large number of liquid nodes that maintain a constant maturity (swap spreads are quoted for 1-year intervals using a constant-maturity basis). Each date has only one specific

interest rate associated with it. This makes the derivation of a zero-coupon curve relatively straightforward.¹ Estimation of a zero curve using the yields on government bonds is, by comparison, a much more difficult problem. The Canadian government bond market contains a large number of issues (80 or more, depending upon the time) of varying maturity, coupon rate, and yield. Of these bonds, however, only about seven or eight actively trade in the secondary market with any significant frequency. Furthermore, it is not unusual for cash flows that occur on the same date to have different yields, depending on whether these flows represent an interest payment (a coupon) or a principal repayment (a residual). These problems necessitate the use of numerical curve-fitting techniques to extract zero-coupon rates. These techniques require that several assumptions be made, depending upon the final use of the interest rates produced. The details of the model used, the justification for it, and the assumptions behind it are provided in section 2.

These problems with estimating zero-coupon yield curves using Government of Canada bond yields are significant, and, as a result, historical databases of zero-coupon yield curves have not been readily available in Canada. While historical term-structure databases exist in the U.S. Treasury market (such as McCulloch and Kwon),² to the best of our knowledge this work represents the first historical database of Canadian risk-free zero-coupon rates in the public domain.

2. The Data and Estimation Model

2.1 The estimation algorithm

A number of estimation algorithms can be used to derive a zero-coupon yield curve based on observed market prices of a set of coupon-bearing bonds. The algorithms can, however, be broadly classified as either *spline-based* or *function-based*. Bolder and Gusba (2002) provide an extensive review and comparison of a number of estimation algorithms using Canadian government bond data. They conclude that, when evaluated against the criteria of goodness of fit, composition of pricing errors, and computational efficiency, the Merrill Lynch exponential spline (MLES) model, as described by Li et al. (2001), is the most desirable term-structure estimate model tested. The MLES model, therefore, was selected as the estimation algorithm used to build the historical database of zero-coupon yields. This model, as with the others that were evaluated

1. For a detailed description of the construction of swap yield curves, see Ron (2000).
2. The U.S. Treasury bond term-structure database is available on J.H. McCulloch's website at <<http://www.econ.ohio-state.edu/jhm/jhm.html>>.

by Bolder and Gusba (2002), is strictly based on curve-fitting techniques. That is, it is a strictly mathematical process defined as fitting a continuous function to a set of discretely observed data points. The process of generating the yield curve makes no underlying economic assumptions, nor does it impose any functional form to the yield curve.

The MLES is used to model the discount function, $d(t)$, as a linear combination of exponential basis functions. It does not, contrary to its name, utilize splines at all. The original paper models the discount function as a single-piece exponential spline, which is simply equivalent to fitting a curve on a single interval. The discount function is given as

$$d(t) = \sum_{k=1}^N \zeta_k e^{-k\alpha t}. \quad (1)$$

The ζ_k are unknown parameters for $k = 1, \dots, N$ that must be estimated. The parameter α , while also unknown, can loosely be interpreted as the long-term instantaneous forward rate. The larger the number of basis functions used, the more accurate the fit that is realized. For our purposes, we use nine basis functions (that is, $N = 9$). We find that, for values of N higher than nine, there is not a substantial improvement in the residual error.

Given the above theoretical form for the discount functions, the next step is to compute the theoretical bond prices. The theoretical price of any bond is simply the sum of the discounted values of its component cash flows, including principal and interest payments. This can be expressed as

$$\hat{P}_i = \sum_{j=1}^{m_i} c_{ij} d(\tau_{ij}), \quad (2)$$

where m_i represents the number of cash flows associated with the i th bond in the sample, c_{ij} is the specific cash flow associated with time τ_{ij} , and d represents the appropriate discount factor. If we denote each of the basis functions as $f_k(t)$, where $k = 1, 2, \dots, N$, we can then solve for matrix H , defined by

$$H_{ik} = \sum_{j=1}^{m_i} c_{ij} f_k(\tau_{ij}). \quad (3)$$

The matrix H is an $N \times D$ matrix, where N is the number of bonds and D is the number of basis functions used. Following this methodology, the column vector of theoretical prices, \hat{P} , can be expressed as $\hat{P} = HZ$, where $Z = (\zeta_1, \dots, \zeta_D)^T$ is the column vector of unknown parameters.

We next construct a diagonal matrix, W , that incorporates the weights associated with each bond. It is necessary to weight the various bonds because we are ultimately generating a yield curve by solving for theoretical bond prices, and we are trying to minimize the pricing error across a full sample of bonds. Given the higher price sensitivity per unit of yield for longer-term bonds (higher duration), if we did not weight the results, the model would treat a given price error on a 1-year bond the same as if it occurred on a 30-year bond. In actuality, the yield error associated with the errors is much greater for the shorter-term instrument.³ To compensate for the greater price/yield sensitivity of longer-term instruments, the bonds are weighted by the reciprocal of their modified duration. This places less weight on pricing errors of longer-term bonds, essentially equalizing the weighting in yield space across bonds of different maturities.

The final step in deriving the discount function is to estimate the parameters ζ_1, \dots, ζ_D . We assume that the pricing errors $\hat{P}_j - P_j$ are normally distributed with a zero mean and a variance that is proportional to $1/w_j$, where w_j is the weight assigned to bond j . We next need to find the set of parameters, ζ_1, \dots, ζ_D , that maximizes the log-likelihood function:

$$l(\zeta_1, \dots, \zeta_D) = - \sum_{j=1}^N w_j (\hat{P}_j - P_j)^2. \quad (4)$$

This can be expressed in matrix form as

$$l(Z) = -\|W(HZ - P)\|^2. \quad (5)$$

Given that the theoretical prices are linear functions of the unknown parameters ζ_1, \dots, ζ_D , we can find the maximum-likelihood estimate using generalized least squares.

This leaves one parameter that is still unknown, α , and there are two options for dealing with it. First, as stated earlier, the value of α can loosely be interpreted as the long-term instantaneous forward rate. As such, we can utilize economic theory and estimate the parameter directly, rather than treat it as an unknown. Second, we can use numerical optimization techniques to solve for the value of α that minimize the residual pricing error. This is the approach that we used in estimating the yield curves in the database. Li et al. (2001) recommend a range of 5 per cent to 9 per

3. For example, a \$0.15 price error on a 30-year bond indicates a yield error of only 1 basis point. The equivalent \$0.15 error on a 1-year bond represents a yield error of 15 basis points.

cent, but we find that any economically reasonable value (given the context of the yield curve being examined) works well.

2.2 Data filtering

The database of bond prices that we use covers the period from January 1986 to May 2003. With approximately 250 days of data for each year in the analysis horizon, this provides more than 4,300 observations. Unfortunately, a small number of dates in any given year (typically, between 10 and 15) appear to be problematic. These dates are characterized by highly non-standard term-structure shapes and/or large yield errors. There are two possible explanations for these anomalies. First, there is one (or more) data entry error(s). Second, the data were entered correctly (that is, they are correct zero-coupon curves), but market conditions were such that the results are somewhat non-standard. These market conditions could include things such as one or more issues trading “on special”⁴ in the repo market—and therefore having market yields well below other, comparable issues—or macroeconomic shocks that result in large, sudden yield movements.⁵ Observationally, it is very difficult to distinguish between these two alternatives. Moreover, in the former case, we would desire to either fix the problem or exclude the date, to avoid the inclusion of erroneous data in our sample. In the latter case, however, these dates represent real data that should be included in the sample, to correctly describe the dynamics of the Canadian term structure. To mitigate this problem, we develop a filtering algorithm to help to objectively determine which bonds to exclude from the sample. The appendix provides details of the algorithm.

2.3 Sample and subsample periods

Substantial changes have occurred to both the structure of the government bond market and the characteristics of the Canadian economy over the horizon of this study. The period can essentially be considered to consist of two distinct subsample periods or regimes for the fixed-income markets.

The first subsample period, from January 1986 to December 1996, can be characterized as follows:

-
4. An issue is referred to as trading “on special” in the repo market if, due to large demand, the interest rate available on a loan that uses that specific issue as collateral is significantly lower than available for other acceptable collateral.
 5. Events such as the 1994 downgrade of Canada’s foreign debt are examples of shocks that could cause yield-curve distortions.

- Relatively high and volatile inflation and inflation expectations, particularly for the first half of the subsample. Over this period, the core consumer price index (CPIX) averaged 3.4 per cent, with a standard deviation of approximately 1.4 per cent. Inflation expectations (as measured by the yield spread between nominal and real return bonds) averaged 3.84 per cent, with a standard deviation of 0.55 per cent.⁶
- Large government borrowing requirements, as a consequence of consecutive federal government deficits. Gross borrowing requirements peaked at approximately \$67 billion in 1996. As a result of these large borrowing requirements, little attention was paid to ensuring an efficient issuance structure. Emphasis was put on simply meeting the government's financing needs.
- A fragmented bond market, characterized by a large amount of relatively small, illiquid issues. There was no predictable issuance pattern and there was no regular pattern of building large, liquid benchmark issues. Stripping and reconstituting individual cash flows was extremely difficult for much of the period, and cash flows with the same maturity often traded at different yields, depending upon which underlying bond they came from. Cash flows with the same maturity date but from different underlying securities were not fungible.
- Few restrictions in primary and secondary market activity, allowing for the possibility of a single participant accumulating a significant position in a specific security. This could result in the specific issue being difficult to borrow in the repo market, forcing it to trade at an artificially low yield.

The second subsample period, from January 1997 to May 2003, experienced very different conditions. It can be characterized as follows:

- Inflation and inflation expectations were low and stable. The Bank of Canada was successful in meeting its inflation targets, and, after a modest lag, the market adjusted its inflation expectations accordingly. The CPIX averaged 1.8 per cent over the period, with a standard deviation of 0.55 per cent. The nominal to real return bond spread averaged 2.1 per cent, with a standard deviation of 0.43 per cent.
- Beginning in 1996, the Government of Canada has run a sequence of budgetary surpluses. This has had a large impact on government borrowing needs. Gross government borrowing, which had peaked at approximately \$67 billion in 1996, fell to \$43 billion in 2001.
- Numerous steps were taken by the Department of Finance and the Bank of Canada that helped to make the government bond market more efficient. These included the introduction of an official benchmark program with explicit issuance targets, a regular and formal consultation with market participants to discuss potential changes to the government debt program, and the implementation of a bond buyback program. The bond buyback program allowed market participants to sell older, off-the-run issues back to the government, either on a cash basis or in trade for the new benchmark bond. The new presence of a large buyer of the illiquid, off-the-run bonds caused them to begin to trade significantly closer to their "fair" value.

6. The yield spread between nominal and real return bonds is subject to a number of distortions that makes its use as a true measure of inflation expectations problematic. As such, these inflation expectations need to be interpreted with care. For a detailed discussion of the issues surrounding the use of this yield gap as a measure of inflation expectations, see Christensen, Dion, and Reid (2004).

-
- Both the Bank of Canada and the Investment Dealers' Association (IDA) implemented measures to enhance the integrity of the primary and secondary market. These measures included auction disclosure rules and IDA Article V.⁷
 - 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 were introduced for book-entry strip bonds.⁸ This allowed 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 a similar type⁹ 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 that had the same issuer and maturity were valued identically, regardless of which underlying issue they came from.
 - Computerized trading strategies and quantitative valuation approaches gained popularity as a means of arbitraging away pricing inefficiencies in the government and swap yield curves. Hedge funds, many of which specialize in fixed-income relative-value arbitrage, also became more significant factors in the fixed-income market.

There are effectively two different regime shifts that should be captured. 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 to the issuance pattern, changes by CDS, and the growing importance of quantitative trading strategies. As such, no specific date marks a perfect break point. The selection of January 1997 as a break 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 of 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. If those changes did indeed make the government bond market more efficient, then the theoretical model should produce a better fit (fewer pricing errors) in the latter period. As well, other differences in the mechanics and behaviour of the zero-coupon yield curves between the two periods may be evident.

7. See <http://www.bankofcanada.ca/en/financial_markets/index.htm> for details of both the auction terms of participation and IDA Article V.

8. CUSIP stands for Committee on Uniform Securities Identification Procedures. A CUSIP number identifies most securities. The CUSIP system facilitates the clearing and settlement process of securities.

9. Fungible cash flows had to be interest payments or principal payments. Interest and principal payments are not yet fungible with each other.

2.4 Model fit

Table 1 summarizes some of the high-level details of the estimation results on an annual basis.

Table 1: Estimation Details

Year	Mean RMSE ¹ (bps) ²	Mean MAE ³ (bps)	Mean available instruments	Mean filtered instruments
1986	17.26	12.11	131	61
1987	15.74	11.12	140	71
1988	13.01	9.47	136	72
1989	11.52	8.39	130	64
1990	12.66	8.84	122	63
1991	10.20	7.62	115	56
1992	11.09	8.48	113	38
1993	9.98	7.56	105	24
1994	5.86	4.02	97	52
1995	5.87	3.21	90	47
1996	8.64	4.87	88	42
Period 1 average	11.08	7.79		
1997	7.68	4.39	76	34
1998	5.50	3.40	84	38
1999	5.36	3.42	82	29
2000	6.45	3.69	76	32
2001	2.83	3.04	69	28
2002	4.52	2.75	65	28
2003	3.84	2.30	63	27
Period 2 average	5.17	3.28		

¹RMSE: Root-mean-square error

²bps: Basis points

³MAE: Mean absolute error

As the table shows, the model provides a significantly better fit in the second subsample period. Both error measures (RMSE and MAE) are smaller. This can be interpreted as being indicative of a more efficient bond market, in the sense that there is more consistency in valuation across different specific issues. Cash flows with similar maturities trade much more consistently in the second subsample than they did in the first.

Brousseau (2002) finds similar results in several European bond markets during the 1990s. Specifically, he tests U.S., French, German, and Spanish interest rate markets (both government and swap yield curves) and evaluates how well market yields correspond to a theoretical yield curve over the period from 1994 to 2000. His findings show that U.S. and French curves were consistently well-fitted over the whole period, while German curves were poorly fitted at the beginning of the period, but improved to the point of matching the U.S. treasury curves by the end of the period. Spanish yield curves never exhibited an excellent fit. Brousseau partially attributes this behaviour to the fact that, while computerized trading techniques and quantitative pricing models were already the rule in the United States and France by 1994, they did not gain popularity in Germany until later in the 1990s. This process accelerated as the German yield curve became the reference curve for the European economy. The Spanish market, which did not adopt quantitative trading strategies to the same degree, did not experience the same shift.

3. Summary of Descriptive Statistics

Section 3.1 shows what an average zero-coupon curve looked like over both the full sample period and the two subsample periods. It also provides some basic descriptive statistics to outline the evolution of four key yield-curve variables over the term of the database. The key yield-curve variables selected for this study are the 3-month yield, the 10-year yield, the slope of the yield curve (the difference between the 3-month and 10-year rates), and the degree of curvature of the yield curve. The curvature measure, denoted as C , is calculated as follows:

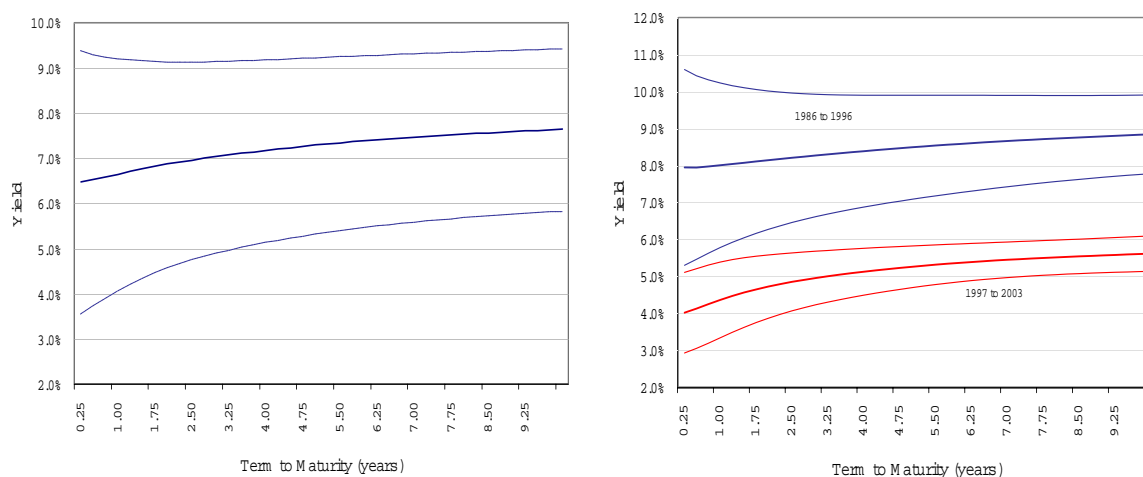
$$C = (6y) - 0.5(2y + 10y) . \quad (6)$$

That is, curvature is equal to the difference between the yield on a 6-year bond and a linear interpolation between the 2-year and 10-year yields.¹⁰

3.1 Average yield curves

As a first step in examining the results, Figure 1 depicts what an average yield curve looked like over both the full sample period and the two subsample periods. The average level of the yield curve is shown, framed by a one-standard-deviation confidence band.

10. The two portfolio choices (a 6-year bond and a combination of 2-year and 10-year bonds) also have equivalent durations. This results from the fact that the duration of a zero-coupon bond is equal to its time to maturity.

Figure 1: Average Zero-Coupon Yield Curves — Full Period and Subsample

As Figure 1 shows, the average 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 those averages, however, was extremely large. A one-standard-deviation band covered a range for the 3-month rate of approximately 3.5 per cent to 9.5 per cent, while for the 10-year rate the band ranged from 5.5 per cent to 9.5 per cent.

Figure 1 also shows the degree to which the yield behaviour differed between the two subsample periods. The average pre-1997 yield curve was upward sloping, with a 3-month rate of about 8 per cent and a 10-year rate of about 9 per cent. The dispersion of yields around the average levels was, however, extremely large. A one-standard-deviation confidence band for the 3-month rate covered a range from 5.5 per cent to 10.5 per cent. For 10-year yields, the range was between 7.5 per cent and 10 per cent. The post-1997 average curve had yields that were so much lower that the upper confidence band of the second subsample was well below the lower confidence band of the first. Furthermore, the dispersion of yields around these average levels was much narrower.

3.2 Descriptive statistics—yield-curve levels

Whereas Figure 1 graphically depicts the general shape of the yield curve over the horizon of the database, this section provides a more detailed statistical description of the various yield-curve measures and their evolution over the full period. It also further highlights the differences between the two subsample periods. Figure 7 illustrates the evolution of the levels of the four key

components of the term structure that were examined, including the raw data, the trend, and summary statistics. Table 2 shows highlights of the results.

Table 2: Summary Yield-Curve Statistics—Full Sample Period

Yield-curve measure	Mean	Max	Min	Std dev	Skew	Kurtosis	Jarque-Bera probability^a
3-month yield	6.46%	13.57%	1.78%	2.90%	0.61	2.42	0.00
10-year yield	7.62%	11.32%	4.53%	1.80%	0.00	1.60	0.00
Slope	1.16%	4.07%	-3.21%	1.66%	-0.61	3.03	0.00
Curvature	12.9 bps	82.3 bps	-46.7 bps	19.9 bps	0.19	3.09	0.00

a. This represents the probability that the series is normally distributed.

As is obvious from Figure 7, the overall level of yields had a significant downward trend over the term of the sample, while the levels of steepness and curvature had a modest upward trend. The distribution of levels was clearly non-normal for all measures.

Figures 8 and 9, combined with Tables 3 and 4, show the same information broken down for the two subsample periods.

Table 3: Summary Yield-Curve Statistics—1986 to 1996

Yield-curve measure	Mean	Max	Min	Std dev	Skew	Kurtosis	Jarque-Bera probability
3-month yield	7.94%	13.57%	2.76%	2.65%	0.22	2.13	0.00
10-year yield	8.84%	11.32%	6.21%	1.06%	-0.24	2.40	0.00
Slope	0.90%	3.93%	-3.21%	1.83%	-0.55	2.46	0.00
Curvature	9.0 bps	82.3 bps	-46.7 bps	19.2 bps	-0.15	2.61	0.00

Table 4: Summary Yield-Curve Statistics—1997 to 2003

Yield-curve measure	Mean	Max	Min	Std dev	Skew	Kurtosis	Jarque-Bera probability
3-month yield	4.01%	5.76%	1.78%	1.09%	-0.23	1.72	0.00
10-year yield	5.61%	7.03%	4.53%	0.48%	0.80	3.45	0.00
Slope	1.60%	4.07%	-0.35%	1.20%	0.39	1.90	0.00
Curvature	19.5 bps	72.6 bps	-19.9 bps	19.3 bps	0.77	2.65	0.00

Yields in general appeared to move lower and become less volatile in the second subsample period, with both the mean and standard deviations of the 3-month and 10-year yield measures significantly lower in the later sample. The slope of the yield curve, on the other hand, increased

significantly in the second period, and negative slopes (or yield curve inversions) went from being a fairly common occurrence in the earlier period to an extremely rare event in the later. The degree of curvature also increased significantly in the second period, although there was no material change in variance. In no case did the distribution of any of the key variables take on a normal shape.¹¹

3.3 Descriptive statistics—first differences

Whereas section 3.2 examined the behaviour of the levels of certain key yield-curve measures over time, an examination of the first differences of these levels is likely of more interest. The first differences (or daily changes) in the level and shape of the yield curve drive the short-term risk and return behaviour for government bonds. Since a zero-coupon bond has no interest payments, its return is entirely driven by price changes. Over very short time periods (such as daily), these price changes are almost entirely driven by changes in the yield level.¹²

The behaviour of these short-term returns is of particular interest. Almost all derivative pricing algorithms, portfolio management tools, and risk-measurement models make some underlying assumptions about the distributional properties of returns over a given time horizon, with the most common assumption being that returns are normally distributed. Since a zero-coupon bond makes no coupon payments, its return is entirely driven by changes in price. These price changes can arise from two sources. The first is the simple accretion of price towards the maturity value that happens over time (zero-coupon bonds are issued at a discount and mature at par). The second source is a change in yield. Over relatively short time horizons, the second source is by far the most significant. It follows, then, that to assume 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 models that make the assumption of normality could be producing results that provide inaccurate prices or risk measures.

11. This is not surprising: since nominal yields are bound at zero, it would be impossible for the distribution of nominal yield levels to be normal.

12. The price of a zero-coupon bond is also impacted by the simple passage of time. Since these instruments have no coupon payments, they trade at a discount to par. The price then steadily converges to par over the life of the bond. While over the long term this effect dominates, its impact on a daily basis is much smaller than that of price changes driven by yield movements. This convergence to par is not an issue for these descriptive statistics, since the first differences were calculated using constant-maturity data.

Figure 10 illustrates the behaviour of the daily changes in the four main yield-curve measures over the full sample period. Table 5 provides the statistical details. Figure 11 and Tables 6 and 7 show the same information broken down for the two subsample periods.

Table 5: Yield-Curve Measure First Differences—Full Sample Period^a

Yield-curve measure	Mean	Max	Min	Std dev	Skew	Kurtosis	Jarque-Bera probability
3-month yield	-0.1	188.3	-120.7	14.5	0.7	14.9	0.00
10-year yield	-0.1	62.1	-92.1	7.3	-0.4	14.9	0.00
Slope	0.0	93.1	-176.3	15.0	-0.6	11.5	0.00
Curvature	0.0	66.7	-50.8	4.8	0.6	29.5	0.00

a. All measures are expressed in basis points.

Table 6: Yield-Curve Measure First Differences—1986 to 1996

Yield-curve measure	Mean	Max	Min	Std dev	Skew	Kurtosis	Jarque-Bera probability
3-month yield	-0.3	188.3	-120.7	17.3	0.6	11.5	0.00
10-year yield	-0.1	62.1	-92.1	8.3	-0.4	13.8	0.00
Slope	0.1	93.1	-176.3	17.7	-0.6	9.2	0.00
Curvature	0.0	66.7	-50.8	5.8	0.5	21.0	0.00

Table 7: Yield-Curve Measure First Differences—1997 to 2003

Yield-curve measure	Mean	Max	Min	Std dev	Skew	Kurtosis	Jarque-Bera probability
3-month yield	0.0	70.8	-51.5	7.9	0.9	12.2	0.00
10-year yield	-0.1	23.5	-22.4	5.2	0.2	4.3	0.00
Slope	-0.2	53.4	-76.8	8.9	-0.6	8.8	0.00
Curvature	0.0	32.7	-33.3	2.4	0.2	78.9	0.00

Three key observations can be made. First, not surprisingly, the average change in the various yield-curve measures was very small, essentially zero for all measures in both subsample periods. Given that these represent daily changes, this small size is to be expected.¹³ Second, the uncertainty surrounding the average measure was very high, with standard deviations that were

13. Realistically, it would be impossible for the average daily change to be significantly different from zero over any reasonable length of time, as this would result in interest rates either falling below zero (if the average change was negative), or reaching very high levels (if the average change was positive).

very large relative to the mean value. Both the standard deviations and the range of values, however, did become significantly lower in the second subsample. Third, and potentially most importantly, the distribution of yield changes is clearly not normal. Rather, the distributions appear to have two distinct properties: (i) they are all highly leptokurtic, with a much larger proportion of the observations close to the mean than would be expected under a normal distribution, and (ii) they are subject to extreme outliers, with every measure having several observations that were up to 12 to 13 standard deviations away from the mean (effectively, a statistical impossibility under a normal distribution). Although 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.

Short-term changes in the yield curve, therefore, clearly were not normally distributed (as indicated by the Jarque-Bera probabilities). As a result, the short-term returns on the underlying zero-coupon bonds were not normally distributed. The historical characteristics of these distributions (both highly leptokurtic and subject to extreme outliers) has some interesting repercussions for the pricing algorithms, portfolio management models, and risk measures that rely on the underlying assumption of normally distributed returns. These models would have systematically underpredicted the probability of a very small change in yields, while at the same time also underpredicting the probability of a very large change in yields. 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 implied volatility than do options with strike prices 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 either move very little or very much (relative to a normal distribution) would have been abnormally profitable.

4. Principal-Components Analysis

4.1 Introduction to principal-components analysis

Principal-components analysis attempts to describe the behaviour of correlated random variables in terms of a small number of uncorrelated “principal components” (PCs). The main idea is that this behaviour, or co-movement, can usually be described by a small number of principal components. Thus, we aim to describe the interrelationships between a large number of correlated

random variables in terms of a much smaller number of uncorrelated random variables. This allows us to determine the main factors that drive the behaviour of the original, correlated random variables.

To begin, consider a random (column) vector $X = (X_1, X_2, \dots, X_M)^T$, where T denotes the matrix transpose operator, with covariance matrix Σ . As long as none of the X_i is an exact linear combination of the other components of the random vector X , Σ will be positive definite. If Σ is a positive definite matrix of dimension m , it has a complete set of m distinct and strictly positive eigenvalues, and there exists an orthogonal matrix A , consisting of the unit eigenvectors of Σ , such that

$$A^T \Sigma A = D \quad , \quad (7)$$

where D is a diagonal matrix with the eigenvalues of Σ along the diagonal.

Consider the random vector defined by

$$Z = A^T X \quad . \quad (8)$$

The covariance matrix of Z is given by

$$\begin{aligned} E[ZZ^T] &= E[A^T X X^T A] \\ &= A^T E[XX^T] A \\ &= A^T \Sigma A \\ &= D. \end{aligned}$$

Thus, by making the transformation $Z = A^T X$, we have constructed a set of uncorrelated random variables,

$$Z^*_i = a_i^T X,$$

where a^i is the i^{th} unit eigenvector of Σ , corresponding to the eigenvalue λ_i . Using the fact that A is an orthogonal matrix, so that $a_i^T a_j = 0$ for $i \neq j$, it follows that

$$\Sigma = \lambda_1 a_1 a_1^T + \lambda_2 a_2 a_2^T + \dots + \lambda_m a_m a_m^T \quad . \quad (9)$$

To get an intuition for the objective of principal-components analysis, consider m linear combinations of the first $l < m$ Z_i . That is, define the m -dimensional random vector Y via

$$Y = A_l Z_l,$$

where

$$A_l = [a_1 \ a_2 \ \dots a_l],$$

$$Z_l = [Z_1 \ Z_2 \ \dots Z_l]^T.$$

The covariance matrix of Y is given by

$$\begin{aligned} E[YY^T] &= E[A_l Z_l Z_l^T A_l^T] \\ &= A_l E[Z_l Z_l^T] A_l^T \\ &= A_l D_l A_l^T. \end{aligned} \tag{10}$$

Expanding (9) as in (8), the covariance matrix of Y is given by

$$\Sigma_Y = \lambda_1 a_1 a_1^T + \lambda_2 a_2 a_2^T + \dots + \lambda_l a_l a_l^T. \tag{11}$$

From (10) we see that, if the last $m-l$ eigenvalues of Σ are small, then Σ_Y will be a good approximation of Σ . Intuitively, this means that only the first l principal components are needed to adequately describe the correlation and co-movement of the original random variables, X_i . In other words, there are really only l “driving forces” that govern the co-movement of the original variables. In our context, it turns out that we will be able to describe a very large portion of the correlation between zero-coupon spot rates across the maturity spectrum with only three PCs.

It is preferable to work with standardized data when the objective is to describe the correlation matrix. If we define

$$X^*_i = X_i - E \frac{[X_i]}{\sigma_i},$$

then we have random variables with zero mean and unit variance. In addition, the covariance matrix of X^* will be identical to the correlation matrix of X . Hence, we can find the eigenvalues and eigenvectors of Σ^* and apply the above results, eliminating small eigenvalues. In addition, since we are dealing with sample data, we will replace $E[X_i]$ with \bar{X}_i and σ_i with S_i , the sample standard deviation for the i^{th} variable (zero-coupon rate, in this case).

The usefulness of identifying a small number of factors that drive the returns of fixed-income securities has been recognized by market participants; Litterman and Scheinkman (1991) were the first to use principal-components analysis to accomplish this task. They find that over 98 per cent of the variation in the returns on government fixed-income securities can be explained in terms of three factors, which they call *level*, *steepness*, and *curvature*. Subsequent analysis of other sovereign debt markets, such as Switzerland and Germany (Buhler and Zimmermann 1996), and the short-term money markets (Knez, Litterman, and Scheinkman 1994) provides similar results, with the same three factors explaining a large percentage of the variation in bond returns. Section 4.2 applies this analysis to zero-coupon Government of Canada bond yields over the period, to determine whether the variation in Canadian yields can be explained by similar factors. As well, the relative importance of the factors, and how that relative ranking evolved over time, will be examined.

4.2 Results

For each year in our sample, we construct the sample correlation matrix from the centred data, $X_{i,t}^c = X_{i,t} - \bar{X}_i$. This is done simply for convenience—when we express the zero-coupon rate as linear combinations of the principal components, there will be no constant term. Table 8 shows the results from our principal-components analysis. We define the percentage variation explained by the i^{th} PC as

$$\lambda_i = \frac{\lambda_i}{\sum_{j=1}^m \lambda_j}.$$

As such, the percentage explained indicates how large a given eigenvalue is relative to the rest.

The results shown in Table 8 are interesting, because they indicate that, similar to Litterman and Scheinkman's results, an average of 99.6 per cent of the correlation between zero-coupon rates can be explained in terms of only three uncorrelated PCs, and that this total explanatory power was stable across the two subsample periods.

Table 8: Percentage Variation Explained

Year	Component 1	Component 2	Component 3	Total
1986	0.910	0.066	0.017	0.993
1987	0.970	0.022	0.006	0.998
1988	0.902	0.080	0.013	0.995
1989	0.709	0.250	0.031	0.990
1990	0.831	0.118	0.046	0.996
1991	0.919	0.075	0.003	0.997
1992	0.877	0.109	0.010	0.997
1993	0.937	0.042	0.012	0.991
1994	0.964	0.033	0.001	0.999
1995	0.931	0.064	0.004	0.999
1996	0.925	0.068	0.005	0.998
Period 1 mean	0.898	0.084	0.013	0.996
1997	0.547	0.428	0.019	0.994
1998	0.763	0.213	0.019	0.995
1999	0.834	0.148	0.015	0.997
2000	0.751	0.234	0.013	0.998
2001	0.814	0.179	0.006	0.999
2002	0.690	0.293	0.015	0.998
2003	0.679	0.307	0.012	0.998
Period 2 mean	0.726	0.258	0.014	0.997
Full sample mean	0.831	0.152	0.014	0.996

It is possible to add some context to how these three principal components impact the shape of the yield curve. Recall that the PCs can be expressed as

$$Z = A^T X^*,$$

where the columns of A constitute the unit eigenvectors of the sample correlation matrix for our selected zero-coupon rates. Since A is orthogonal, it is invertible, with $A^{-1} = A^T$. Therefore, we can express the standardized data, X^* , as a linear combination of the PCs:

$$X^* = AZ.$$

Thus, the i^{th} standardized zero-coupon rate can be written as

$$X^*_i = a_{i,1}Z_1 + a_{i,2}Z_2 + \dots + a_{i,m}Z_m,$$

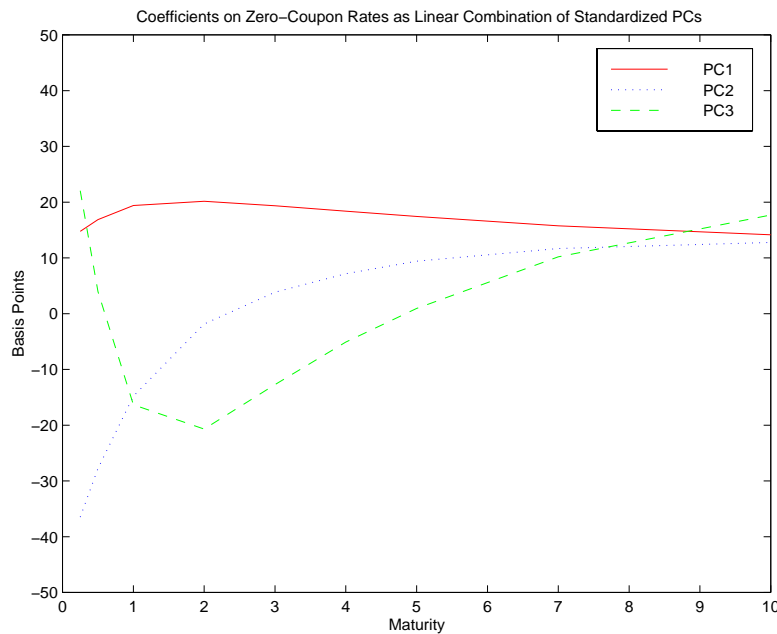
where $a_{i,j}$ is the i^{th} component of the j^{th} eigenvector (the unit eigenvector corresponding to the eigenvalue λ_j). Since $X^*_i = \frac{(X_i - \mu_i)}{\sigma_i}$, we can express the i^{th} zero-coupon rate as

$$X_i = \mu_i + \sigma_i a_{i,1} Z_1 + \sigma_2 a_{i,2} Z_2 + \dots + \sigma_i a_{i,m} Z_m. \quad (12)$$

Because we are dealing with sample data, we replace μ_i with $\hat{\mu}_i$, the sample mean of the i^{th} rate, and σ_i with S_i , the sample standard deviation for the i^{th} rate. The coefficients $a_{i,j}$ are obtained from the eigenvectors of the estimated correlation matrix of the standardized data. As well, since the first three PCs explain over 99.5 per cent of the variation, we can shorten equation (12) to use only $m = 1,2,3$.

Figure 2 shows the sensitivities of each rate versus their maturity for the first three PCs. In other words, a given curve in the graph plots the (rescaled) components of the eigenvectors corresponding to the first three factors. This helps to facilitate the interpretation of these PCs.

Figure 2: Sensitivities of Zero-Coupon Rates to First Three Factors



For the first PC, we see that the sensitivities of the rates to this factor are roughly constant across maturities. Thus, if this PC increased by a given amount, we would observe a (approximately) parallel shift in the zero-coupon term structure. This PC corresponds to the *level* factor of Litterman and Scheinkman. As Table 8 indicates, this PC is the most important determinant in the movements of the term structure, accounting for an average of 83 per cent of the total variation. The explanatory power of this first PC is not stable over the two subsample periods, however. The *level* is significantly less important in explaining total variation in the zero-coupon yield curve in

the latter subsample period than it is in the first (although it is still by far the most important of the three).¹⁴

The second PC tends to have an effect on short-term rates that is opposite to its effect on long-term rates. An increase in this PC causes the short end of the yield curve to fall and the long end of the yield curve to rise. This is the *steepness* factor—a change in this factor will cause the yield curve to steepen (positive change) or flatten (negative change). Table 8 shows that the relative importance of this factor changed materially over the two subsample periods, accounting for roughly twice the variation in the second subsample period as it did in the first. Changes in the *steepness* of the yield curve explain significantly more of the total variation in the curve in the post-1997 sample than they do in the pre-1997 sample. In both cases, though, this factor is second to the *level* PC in terms of the amount of variation in the yield curve it explains.

The third PC corresponds to the *curvature* factor, because it causes the short and long ends to increase, while decreasing medium-term rates. This gives the shape of the zero-coupon yield curve more or less curvature. This PC seems to be the least significant of the three, accounting for an average of less than 2 per cent of the total variation in term-structure movements. As well, the amount of variation explained by this PC does not vary significantly over the two subsample periods.

To further facilitate the interpretation of these principal components, Figure 3 illustrates what happens to a sample yield curve when there is a shock to one PC. The figure illustrates the actual yield curve on 12 March 2003 and the recalculated yield curve (via equation (12)) after increasing one PC by five units.

14. The p -values for heteroscedastic t -tests for equality of means between the two subsample periods for both *level* and *steepness* are 0.00.

Figure 3: Sensitivities of Zero-Coupon Rates to a Principal-Component Shock

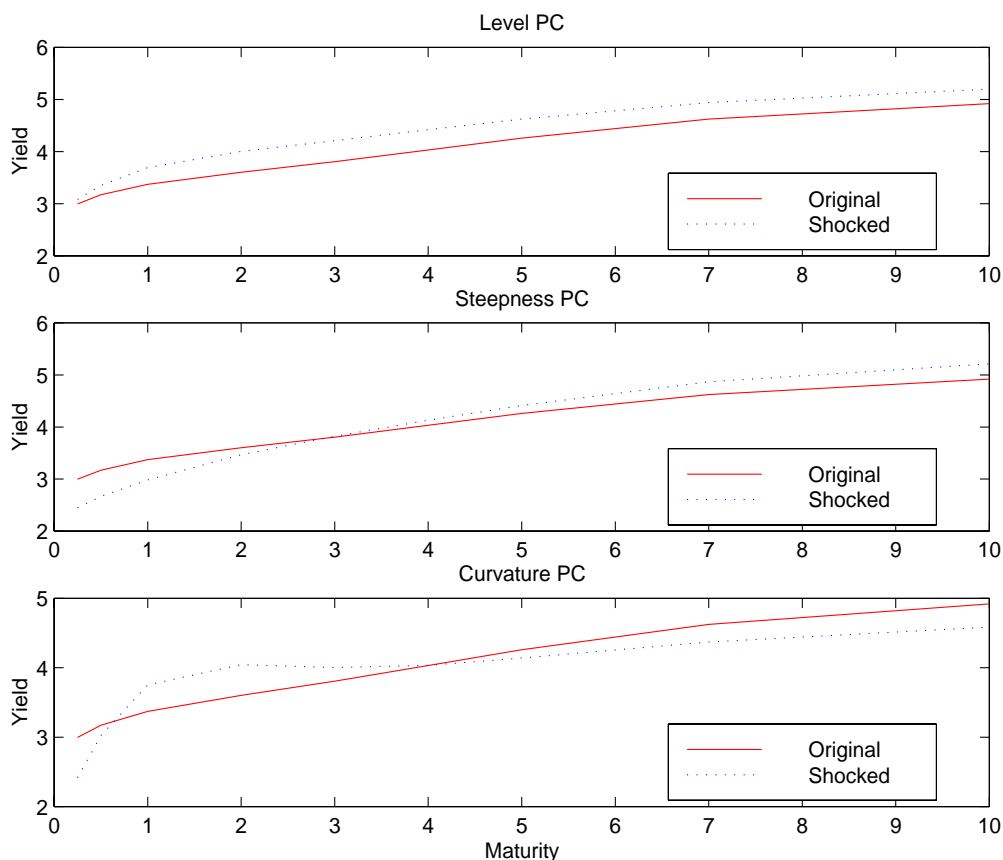


Figure 3 shows that, although the first PC does not induce a perfectly parallel shift in the term structure (the shift higher in yields is accompanied by a modest steepening of the yield curve), it affects all rates by roughly the same amount. The second PC makes the yield curve become steeper, whereas the third PC introduces more curvature into the shape of the term structure.

As Table 8 shows, the first three PCs account for an average of 99.6 per cent of the total variation in the yield curve, and this proportion is stable over the two subsamples. Since we have used nine maturity points to describe a specific zero-coupon curve ($m = 9$), there are, by definition, nine principal components. The fourth through ninth PCs, however, have very little effect on the term structure, accounting for a total of only 0.4 per cent of the variation. Not surprisingly (given the linkages between the U.S. and Canadian government bond markets), these results are broadly similar to those of the earlier studies. The same three factors (level, steepness, and curvature) are found to explain almost all of the variation in Canadian bond yields, and the relative ranking of the importance of these factors is the same.

4.3 Potential applications

The principal-components analysis essentially shows that any daily shift in the zero-coupon yield curve can be separated into three uncorrelated components. These components, in order of importance, are a more or less parallel movement, a change in steepness, and a change in curvature. Traditional interest rate risk management emphasizes duration, assuming that only the first PC is important. This analysis clearly shows that this type of hedging ignores a substantial amount of risk, because in the post-1997 period a parallel shift represented only 73 per cent of the total variation in yields. A portfolio can be duration-neutral (i.e., the assets have the same duration as the liabilities), but still be exposed to changes in slope and curvature (the second and third PCs).

Principal-components analysis allows the creation of PC durations, with each duration measuring the sensitivity of the portfolio to one of the three components. A more complete hedging strategy would then entail having a portfolio of assets that offset the three key durations (level, slope, and curvature) of the liability portfolio. The analysis above shows that, properly constructed, this hedge would protect against more than 99 per cent over the variability in the term structure. The use of principal-components analysis in hedging fixed-income portfolios has been the subject of a relatively large amount of research; additional information is provided by Barber and Copper (1996), Golub and Tilman (1997), and Lardic, Priaulet, and Priaulet (2001).

Principal-components analysis is also useful in the construction of functional forms of the yield curve. Yield curves can be modelled in a number of ways, depending on the motives for the use of the final curve. Practitioners, who focus on pricing accuracy, generally favour straight curve-fitting algorithms (such as the MLES) and no-arbitrage models, which are primarily concerned with accurately fitting the term structure at a specific point in time. At the other end of the spectrum are the economic-based models, which employ expectations regarding inflation, future economic growth, and the dynamics of the short rate over time. Falling somewhere in between these two approaches is the functional specification of the yield curve. Functional representations of the yield curve use a small number of time-varying latent factors to model its evolution. In one of the earliest and best-known papers on this approach, Nelson and Siegel (1987) construct a parametrically parsimonious model of the yield curve based on three time-varying parameters. Although these parameters are not specifically interpreted as factors in the original paper, a number of subsequent authors have interpreted them as the specific factors level, slope, and curvature.¹⁵ While the results of the principal-components analysis confirm that the dynamics of

15. See, for example, Diebold and Li (2003), Diebold, Rudebusch, and Aruoba (2003), and Duffie and Kan (1996).

the yield curve can be reasonably fully modelled with only these three factors, these functional forms can further be developed by the addition of macroeconomic factors (as in Diebold, Rudebusch, and Aruoba 2003).

Another potential application of principal-components analysis lies in Monte Carlo simulations of interest rate paths. One way to simulate the yield-curve changes necessary for a Monte Carlo simulation would be to generate random vectors from the joint (multivariate) distribution of zero-coupon rates. We would like to use as many rates as possible, however, and we cannot forget the fact that yields of differing maturities are highly correlated. Thus, we could not simply generate 3-month changes independently of 10-year changes, since this would ignore the patterns in the co-movements of these rates. To generate yield-curve changes via the joint distribution of yields, therefore, would be cumbersome and very difficult—it would involve generating them from a high-dimension distribution (we would most likely want to use at least 10 rates). Further, this joint distribution could potentially be very complicated in the presence of our empirically observed correlations between rates. Simulation might be straightforward if we assume that yield changes are multivariate normal; however, as we have shown in sections 3 and 4, the normal distribution does a very poor job of describing yield changes.

An alternative would be to simulate the principal components themselves and transform them back into yields. There are two main reasons why this makes the simulation process easier. First, the principal components are uncorrelated, so we can generate them independently. This means that we do not have to take into account the likely effect of a change in the 3-month rate on the 10-year rate, which simplifies tremendously the target distribution in the Monte Carlo simulation. Second, we find that there are only three principal components that are significant drivers of the shape of the term structure. Thus, instead of generating samples from the 10-dimensional distribution of yields, we can simply generate from the 3-dimensional distribution of the first three components.

5. Holding-Period Returns

The expectations hypothesis of the term structure of interest rates maintains that longer-term interest rates are simply the geometric average of future short-term rates, plus a term premium. If the value of the term premium is set to zero, it follows that the expected returns on bonds of all maturities are equal over a given time horizon (e.g., buying a 10-year bond and selling it in one year provides the same expected return as buying a 1-year bond and holding it to maturity). If the term premium is positive, longer-term bonds will have a higher expected return over a given

investment horizon than shorter-term instruments will. This extra return, however, compensates for the higher risk associated with longer-term instruments.

We will use the historical yield series to answer three questions related to this hypothesis. First, have bonds of different maturities provided equivalent returns for a given holding period, or have longer-term instruments provided some measure of excess return (i.e., is the term premium zero or positive)? Second, were the returns earned from holding longer-term instruments riskier (more variable) than they were for shorter-term bonds? Third, were the risk-adjusted returns across maturities equivalent, or did one sector tend to outperform the others on a risk-adjusted basis?

5.1 Definitions and notation

The usual notation for the price at time t of a default-free, zero-coupon bond (having a face value of \$1) maturing on date T is $P(t, T)$. This means that the price of a bond with a time-to-maturity of T years can be expressed as $P(t, t+T)$. The continuously compounded zero-coupon rate, $z(t, T)$, is defined as:

$$P(t, T) = e^{-(T-t) \cdot z(t, T)},$$

which gives:

$$z(t, T) = -\frac{1}{T-t} \log(P(t, T)).$$

Our yield-curve data consist of 4,204 daily yield curves. In the usual notation, the data consist of a set of 4,204 functions:

$$\{z(t_i, T); i = 1, 2, \dots, 4204\}.$$

We will use this notation in the following analysis, with the assumption that one year has 360 days, since our yield curves contain rates in 3-month increments (each of which is taken to be 0.25 years). As a result, for any given date t_i , we have the 90/360-year (3-month) and 180/360-year (6-month) rates, but not the 90/365-year or 180/365-year rates.

The next step is to define holding-period returns. An N -day holding-period return (HPR) beginning at time t_i on T -year bonds is defined as the net percentage return that is realized from the following hypothetical strategy:

- At a given date, t_i , purchase a risk-free zero-coupon bond maturing in T years (i.e., at date $t_i + T$). For simplicity, assume that the face value of the bond is \$1. The price of this bond at date t_i is given by $P(t_i, t_i + T)$.
- Hold the bond for N days.
- On date $t_i + N/360$, sell the bond. Note that, as of date $t_i + N/360$, the bond will have a time-to-maturity of $(T - N/360)$ years, but will still be maturing at date $t_i + T$. Also note that $t_i + N/360$ will represent a date that is N days after t_i . The price of this bond when it is sold at $t_i + T$ is $P(t_i + N/360, t_i + T)$.
- Defining $d = \frac{N}{360}$, the net percentage return to the strategy is

$$\begin{aligned} HPR(t_i, N, T) &= \frac{P(t_i + d, t_i + T) - P(t_i, t_i + T)}{P(t_i, t_i + T)} \\ &= \frac{P(t_i + d, t_i + T)}{P(t_i, t_i + T)} - 1. \end{aligned}$$

The holding-period return for a zero-coupon bond is therefore simply the difference between the price at which the bond is sold, $P(t_i + d, t_i + T)$, and the price at which the bond is bought, $P(t_i, t_i + T)$, divided by the total investment made on date t_i .¹⁶ Expressed in terms of zero-coupon yields, rather than prices, we get

$$\begin{aligned} HPR(t_i, N, T) &= \frac{\exp(-(T - d) \cdot z(t_i + d, t_i + T))}{\exp(-T \cdot z(t_i, t_i + T))} \\ &= \exp(T \cdot z(t_i, t_i + T) - (T - d) \cdot z(t_i + d, t_i + T)) - 1. \end{aligned}$$

It will be useful in the subsequent sections to consider holding-period yields (HPYs), which are denoted by $X(t_i, N, T)$ and defined as the solution to the relation

$$e^{d \cdot X(t_i, N, T)} - 1 = HPR(t_i, N, T). \quad (13)$$

One way to interpret these HPYs is as follows: if an investor had been paid interest continuously at a fixed rate of $X(t_i, N, T)$ on a \$1 investment between dates t_i and $t_i + d$, they would have earned $\$HPR(t_i, N, T)$ in interest. Note that, since it depends on the bond price at $t_i + d$, $X(t_i, N, T)$ is not known at time t_i . It is simply the rate at which an investor would have earned exactly the same

16. This allows for a much cleaner calculation of holding period returns than would be the case using coupon-bearing instruments, which would necessitate including the return earned from the re-investment of periodic cash flows.

amount in interest as if they had used the HPR strategy. We can use equation (13) to derive an explicit expression for $X(t_i, N, T)$:

$$X(t_i, N, T) = \frac{1}{d} \cdot (T \cdot z(t_i, t_i + T) - (T - d) \cdot z(t_i + d, t_i + T)).$$

In the following sections, we deal exclusively with HPYs, comparing them with yields on zero-coupon bonds that are held until maturity. For this purpose, we need to express the HPRs as some type of continuous rate, as opposed to a simple percentage return.

Our analysis focuses on the concept of excess HPYs. The excess yield is defined as the excess of the HPY compared with some risk-free reference rate. The risk-free reference rate is defined as the yield on a zero-coupon government bond with d years to maturity. This yield is risk-free in that the investor does not need to sell the bond at time $t_i + d$, but rather the bond matures with a known terminal value of \$1. As a result, the realized yield is known at time t_i with certainty (i.e., it is risk-free).

5.2 Summary results for holding-period yields—full sample

HPYs are calculated for a holding period of $N = 180$ days and using zero-coupon instruments with maturities of $T = 1, 2, 5,$ and 10 years. To calculate excess returns, these HPYs are compared with the yield on a zero-coupon instrument with a 180-day maturity. The data are as follows:

- There are 3,334 observations.
- The first observation, $X(t_1, 180, T)$, is the yield corresponding to the return realized between 31 March 1989 and 27 September 1989.¹⁷
- The last observation, $X(t_{3334}, 180, T)$, is the yield corresponding to the return realized between 29 November 2002 and 28 May 2003.

As before, we examine both the entire data set and the two subsample periods: the pre-1997 subsample and the post-1997 subsample. Table 9 and Figure 4 show the summary results for the full period. It is immediately evident that HPYs get both larger and more volatile as the maturity of the bonds held increases. A 1-year instrument has the smallest mean and median return, a standard deviation that is one-fifth as large as the 10-year instrument, and does not produce a negative return in any of the periods examined. At the other end of the spectrum, the 10-year bond has the highest average return (both mean and median), the largest standard deviation, and has produced negative returns of up to -46 per cent over a 180-day period. The results conform with

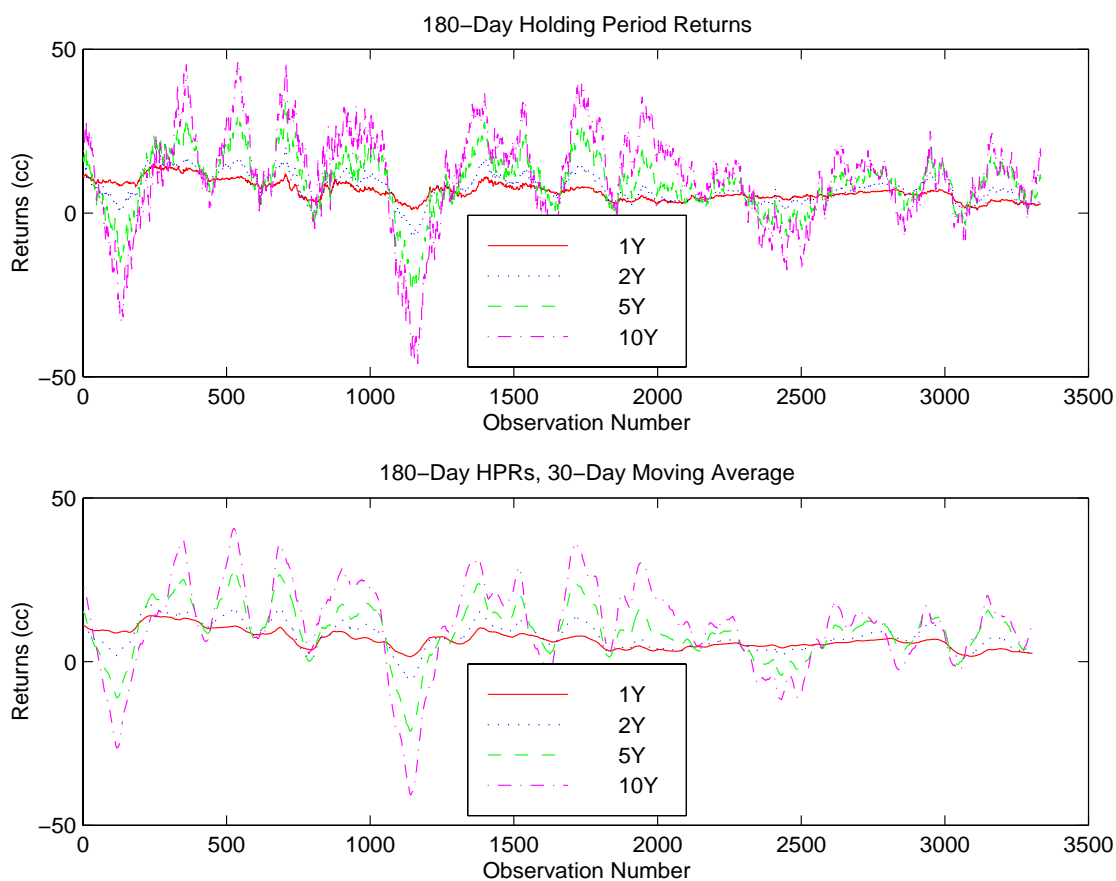
17. This represents the first 180-day period that did not have a large number of missing observations. Prior to 31 March 1989, there is a gap of 35 days.

the notion of longer-term assets being riskier, and therefore demanding a positive risk premium (higher expected return).

Table 9: Summary Statistics for 180-Day HPYs

Bonds	Mean (%)	Median (%)	SD (%)	Max (%)	Min (%)	Skewness	Kurtosis
1-year	6.66	6.29	2.89	15.08	1.01	0.58	2.86
2-year	7.58	7.05	4.39	20.43	-6.75	0.03	2.95
5-year	9.22	9.75	8.67	34.13	-24.20	-0.51	3.97
10-year	10.95	12.14	14.53	46.19	-45.98	-0.70	4.20

Figure 4: Daily Observations of 180-Day HPYs



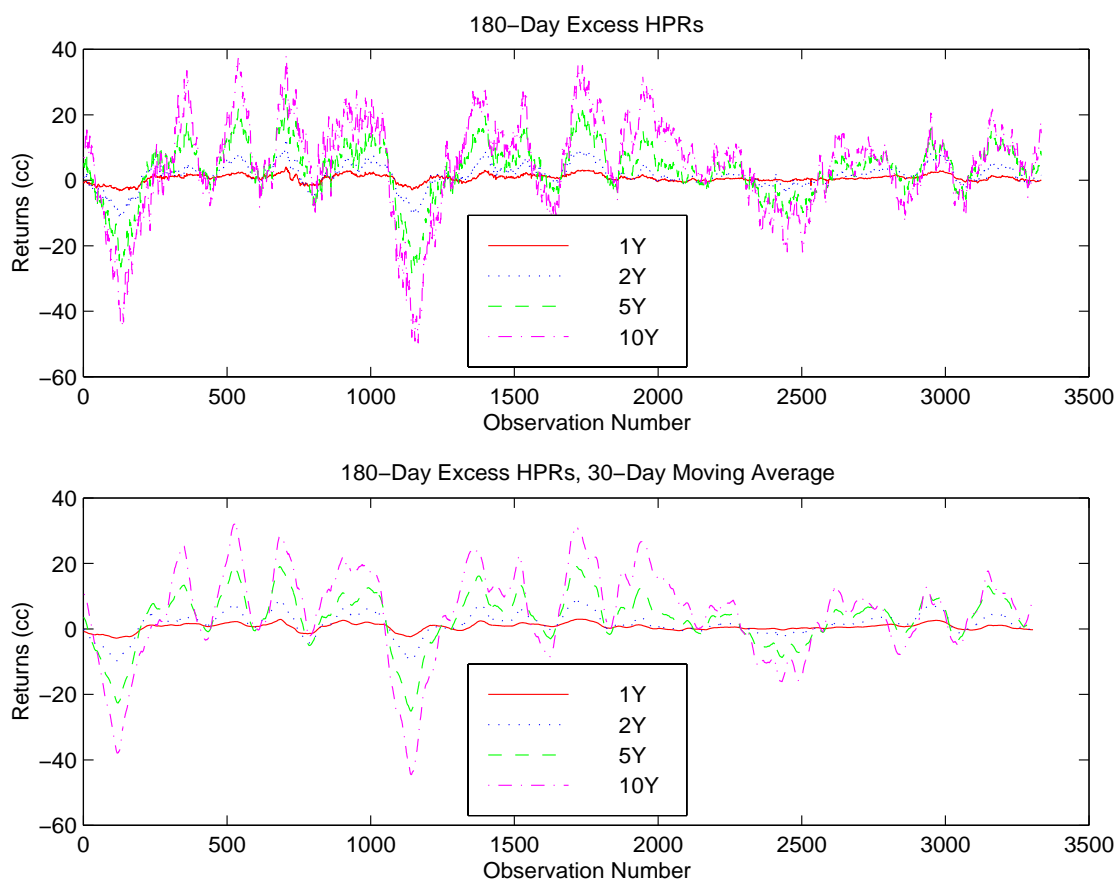
As stated earlier, excess HPYs are a more accurate way to compare the dynamics of HPYs over time, since they adjust for changes in the level of the reference risk-free rate (the 180-day rate, in

this case) over the periods examined. Table 10 and Figure 5 show the summary statistics for the excess HPYs over the full sample.

Table 10: Summary Statistics for 180-Day Excess HPYs

Bonds	Mean (%)	Median (%)	SD (%)	Max (%)	Min (%)	Skewness	Kurtosis
1-year	0.61	0.66	1.18	4.05	-3.27	-0.53	3.55
2-year	1.53	1.61	3.57	12.53	-11.10	-0.64	4.09
5-year	3.17	4.05	8.35	26.18	-28.24	-0.80	4.35
10-year	4.89	6.45	14.48	38.18	-49.66	-0.80	4.18

Figure 5: Daily Observations of 180-Day Excess HPYs



Not surprisingly, the summary statistics for excess HPYs show the same pattern as they do using absolute HPYs. The magnitude of the incremental risk and return becomes clearer, however. A 1-year instrument shows a minimal excess return, with a mean value of only 66 basis points. The 10-year bond shows a mean excess return of 489 basis points, almost 7.5 times as large. The difference in standard deviations is equally striking, with the 1-year instrument having a standard deviation of 118 basis points, whereas the 10-year bond shows a standard deviation of almost 1,450 basis points.

5.3 Summary results—subsamples

Section 3.1 demonstrated that yields are generally significantly less volatile in the second subsample (the post-1997 period). This section examines the behaviour of excess HPYs over the subsamples. Theoretically, the reduced risk of longer-dated bonds (the lower yield volatility evident in the second subsample) should result in lower excess HPYs.

Table 11: Summary Statistics for Excess HPYs: Pre-1997

Bond	Mean (%)	Median (%)	SD (%)	Max (%)	Min (%)	Skewness	Kurtosis
1-year	0.69	1.04	1.4	4.05	-3.27	-0.73	2.79
2-year	1.69	2.37	4.37	12.53	-11.10	-0.78	3.27
5-year	3.46	4.93	10.17	26.18	-28.24	-0.85	3.55
10-year	5.24	8.30	17.48	38.18	-49.66	-0.85	3.48

Table 12: Summary Statistics for Excess HPYs: Post-1997

Bond	Mean (%)	Median (%)	SD (%)	Max (%)	Min (%)	Skewness	Kurtosis
1-year	0.51	0.40	0.70	2.86	-1.61	1.09	4.33
2-year	1.32	1.05	2.20	8.13	-3.57	0.71	3.23
5-year	2.79	3.35	5.35	15.86	-11.58	-0.16	2.66
10-year	4.45	5.59	9.82	31.79	-21.92	-0.17	2.96

As Tables 11 and 12 show, both the average excess HPYs and their associated standard deviations are lower in the post-1997 period. The decline in the level of risk associated with longer-term instruments (the standard deviation of the excess HPYs) is most noticeable. While the fall in the value of the excess HPYs is not very large (ranging from 19 basis points for the 1-year asset to 69 basis points for 10-year assets),¹⁸ the decrease in the standard deviation of these returns is material, with all maturities experiencing decreases of approximately 50 per cent. These results support the hypothesis that the bond market became “safer” in the post-1997 period, offering broadly similar returns, but with a lower level of risk associated with those returns.

Another way to illustrate the levels of risk in the two subsample periods is to examine the probability of realizing negative excess returns. In other words, what were the odds of earning less than the risk-free rate by owning longer-maturity assets? To pose the question in a slightly different way, what were the odds of realizing positive excess returns? Tables 13 and 14 show the frequency and size of negative and positive excess returns over the two subsample periods.

18. These differences in excess HPYs are not statistically significant, given the small differences and relatively large associated standard deviations.

Table 13: Frequency and Size of Negative Excess Returns

Maturity	<u>Pre-1997</u>		<u>Post-1997</u>	
	Probability (%)	Mean (%)	Probability (%)	Mean (%)
1-year	25.4	-1.44	23.9	-0.23
2-year	25.8	-4.28	31.1	-0.93
5-year	29.4	-8.94	30.0	-3.67
10-year	31.6	-15.16	30.5	-7.17

Table 14: Frequency and Size of Positive Excess Returns

Maturity	<u>Pre-1997</u>		<u>Post-1997</u>	
	Probability (%)	Mean (%)	Probability (%)	Mean (%)
1-year	74.6	1.42	76.1	0.74
2-year	74.2	3.76	68.9	2.33
5-year	70.6	8.63	70.0	5.56
10-year	68.4	14.65	69.5	9.55

The probability of earning a positive excess return for a given period is, obviously, equal to one minus the probability of earning a negative excess return.

Interestingly, the chance of losing money (as defined by earning less than the risk-free rate) does not vary considerably either across bond maturities or across subsample periods. In all cases, the chances of earning a negative excess return are roughly between 25 per cent and 30 per cent. What does vary considerably, over both maturity and subsample periods, is the size of the average negative excess returns: it increases sharply as the term-to-maturity of the bond increases. As well, the mean excess negative return is significantly lower for all terms in the post-1997 period. A similar pattern occurs in the size of the mean excess positive returns: it increases as the term-to-maturity of the underlying bond increases, and it decreases significantly in the post-1997 period.

The negative and positive excess return statistics shown in Tables 13 and 14 continue to support the hypothesis that the bond market has become less risky in the post-1997 period. While the probabilities of earning either positive or negative excess returns does not materially change, the size of those excess returns does. In the pre-1997 period, the size of excess return shocks (both positive and negative) are significantly larger than they are in the post-1997 period.

5.4 Risk-adjusted returns

The excess HPY data in the previous section support two primary conclusions. First, both the expected risk and the expected return increase as the time-to-maturity of the bond examined increases. For the full sample, the mean 180-day excess HPY increases from 0.61 per cent for a 1-year bond to 4.89 per cent for a 10-year bond. The standard deviation of these returns increases from 1.18 per cent for the 1-year asset to almost 14.5 per cent for the 10-year asset. It appears that longer-dated assets carry a positive risk premium to compensate for the additional volatility of their returns. Second, while the mean excess HPYs decline slightly in the post-1997 period, the level of risk (as measured by the standard deviation of those excess HPYs) declines much more substantially.

In this section, we examine the HPYs after adjusting them for risk. The simplest way to do this is to construct Sharpe ratios for the various HPYs. Sharpe ratios are defined as the ratio of the excess return on a risky asset to its volatility. We can calculate the ex post (or historic) Sharpe ratio as follows:

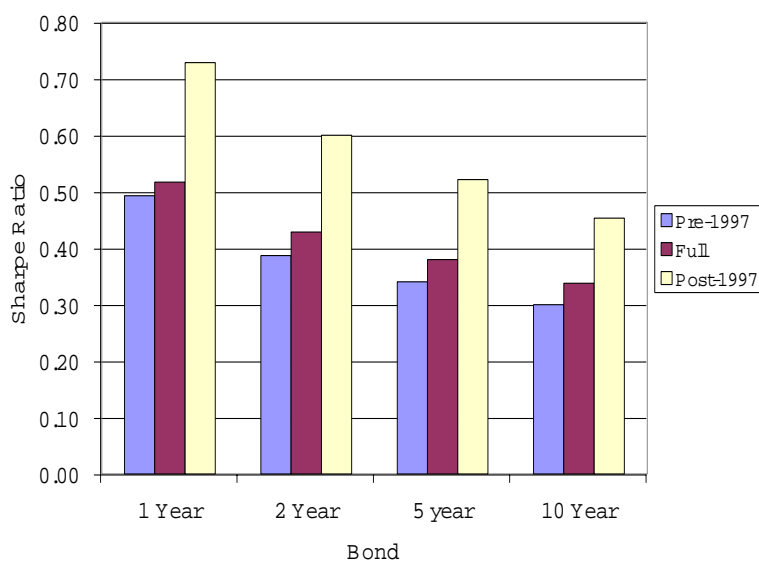
$$S = \frac{HPY}{\sigma_{HPY}}.$$

The Sharpe ratio (S) is equal to the mean excess HPY divided by the standard deviation of those HPYs. The larger the value of S , the higher the risk-adjusted return.

Table 15 and Figure 6 show the Sharpe ratios for the various maturities over the full sample, the pre-1997 subsample period, and the post-1997 subsample period.

Table 15: Sharpe Ratio Calculations—Excess HPYs

Bond	Full sample			Pre-1997 subsample			Post-1997 subsample		
	Mean Excess HPY (%)	SD (%)	Sharpe	Mean Excess HPY (%)	SD (%)	Sharpe	Mean Excess HPY (%)	SD (%)	Sharpe
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

Figure 6: Sharpe Ratios

Two main characteristics of the Sharpe ratios are apparent. First, they decrease with the time-to-maturity of the bonds for the full sample period and for both subsample periods. The incremental return earned by longer-maturity bonds does not compensate for the extra risk (according to this measure of risk adjustment). Second, Sharpe ratios for every maturity are higher in the post-1997 period than in the pre-1997 period, by approximately 50 per cent. As before, it appears that the Canadian bond market offered a superior risk-reward trade-off in the post-1997 period.

6. Conclusions

We have introduced a comprehensive database of Government of Canada zero-coupon yield curves, and provided some preliminary statistical analysis of the data to draw some general conclusions about the behaviour and evolution of Canadian government bond yields over the period. Three main conclusions can be drawn.

The first is that the behaviour of the government bond yields was significantly different in the latter part (January 1997 to May 2003) of the sample. By almost any measure, the bond market became a “safer” place in this second subsample period. Indications of this decrease in risk are numerous. The yield-curve model provided a much better fit in the latter period, indicative of less idiosyncratic pricing behaviour by individual securities. The level of volatility of the various yield-curve measures (3-month, 10-year, slope, and curvature) fell significantly in the second subsample period, both for levels and first differences. Measures of risk based on holding-period returns also showed lower risk in the second subsample. While excess HPYs were smaller in the

second subsample, the standard deviation of these returns was significantly lower, leading to better risk-adjusted performance across the yield curve.

The second conclusion results from the principal-components analysis. This analysis shows that three factors (which we refer to as level, slope, and curvature) account for over 99.5 per cent of the total variation in the yield curve over the full sample period. While this proportion has remained extremely stable over the period (ranging from 99.0 to 99.9 per cent), the breakdown among the three factors varies considerably. The *level* factor explains an average of almost 90 per cent of the variability in the first subsample period, but only 73 per cent in the second. The amount of variability explained by the *slope* factor, meanwhile, rises from 8 per cent in the first subsample period to almost 26 per cent in the second. In the period after 1996, the absolute level of yields appears to have become relatively less significant, with the shape of the yield curve accounting for an increasing amount of yield-curve variability.

The third conclusion concerns the distributional properties of the daily changes of the various yield-curve measures. The distribution of these daily changes is not normal for any other measures examined. The actual distributions are much more leptokurtic than normal, with a larger proportion of observations than would be expected occurring right around the mean. As well, the distributions are characterized by “fat tails,” with a much larger proportion of outliers observed than would be expected. These distributional properties hold in both subsample periods (although variability is much lower in the second). The daily changes in both the level (3-month and 10-year interest rates) and shape (slope and curvature) of the yield curve average zero (and, in fact, are far more likely to be zero than a normal distribution would suggest), but are prone to disproportionately large moves in either direction. The behaviour of the yield curve, in general, can be characterized as general stability punctuated by periods of extreme moves.

We have provided a relatively high-level statistical overview of the behaviour of the Government of Canada yield curve over a period of approximately 17.5 years. Our analysis is based upon what we believe to be the first constant-maturity Government of Canada yield curve to be available in the public domain. These data are a new and rich resource for further research, and will be made available on the Bank of Canada’s website.¹⁹

19. While we will make the data available in the public domain, the authors request that anyone making use of the database cite this paper.

Bibliography

- 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.
- Brousseau, V. 2002. "The Functional Form of Yield Curves." European Central Bank Working Paper No. 148.
- 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(3): 55–67.
- Christensen, I., F. Dion, and C. Reid. 2004. "Real Return Bonds, Inflation Expectations, and the Break-Even Inflation Rate." Bank of Canada Working Paper No. 2004-43.
- Diebold, F. and C. Li. 2003. "Forecasting the Term Structure of Government Bond Yields." Department of Economics, University of Pennsylvania. Manuscript.
- Diebold, F., G. Rudebusch, and S.B. Aruoba. 2003. "The Macroeconomy and the Yield Curve: A Nonstructural Analysis." Federal Reserve Bank of San Francisco Working Paper No. 2003-18.
- Duffie, D. and R. Kan. 1996. "A Yield-Factor Model of Interest Rates." *Journal of Mathematical Finance* 6(4): 379–406.
- 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* (Summer): 72–85.
- Ho, T.S.Y. 1992. "Key Rate Durations: Measures of Interest Rate Risks." *Journal of Fixed Income* (September): 29–44.
- 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. 2001. "PCA of the Yield Curve Dynamics: Questions of Methodologies." *Journal of Bond Trading & Management* 1(4): 327–49.
- Li, B., E. DeWetering, G. Lucas, R. Brenner, and A. Shapiro. 2001. "Merrill Lynch Exponential Spline Model." Merrill Lynch Working Paper.
- Litterman, R. and J. Scheinkman. 1991. "Common Factors Affecting Bond Returns." *Journal of Fixed Income* 1: 54–61.
- Nelson, C.R. and A. Siegel. 1987. "Parsimonious Modelling of Yield Curves." *Journal of Business* 60: 473–89.
- Ron, U. 2000. "A Practical Guide to Swap Curve Construction." Bank of Canada Working Paper No. 2000-17.
- Willner, R. 1996. "A New Tool for Portfolio Managers: Level, Slope and Curvature Durations." *Journal of Fixed Income* 6(1): 48–59.

Figure 7: Yield-Curve Measures — Full Sample Period

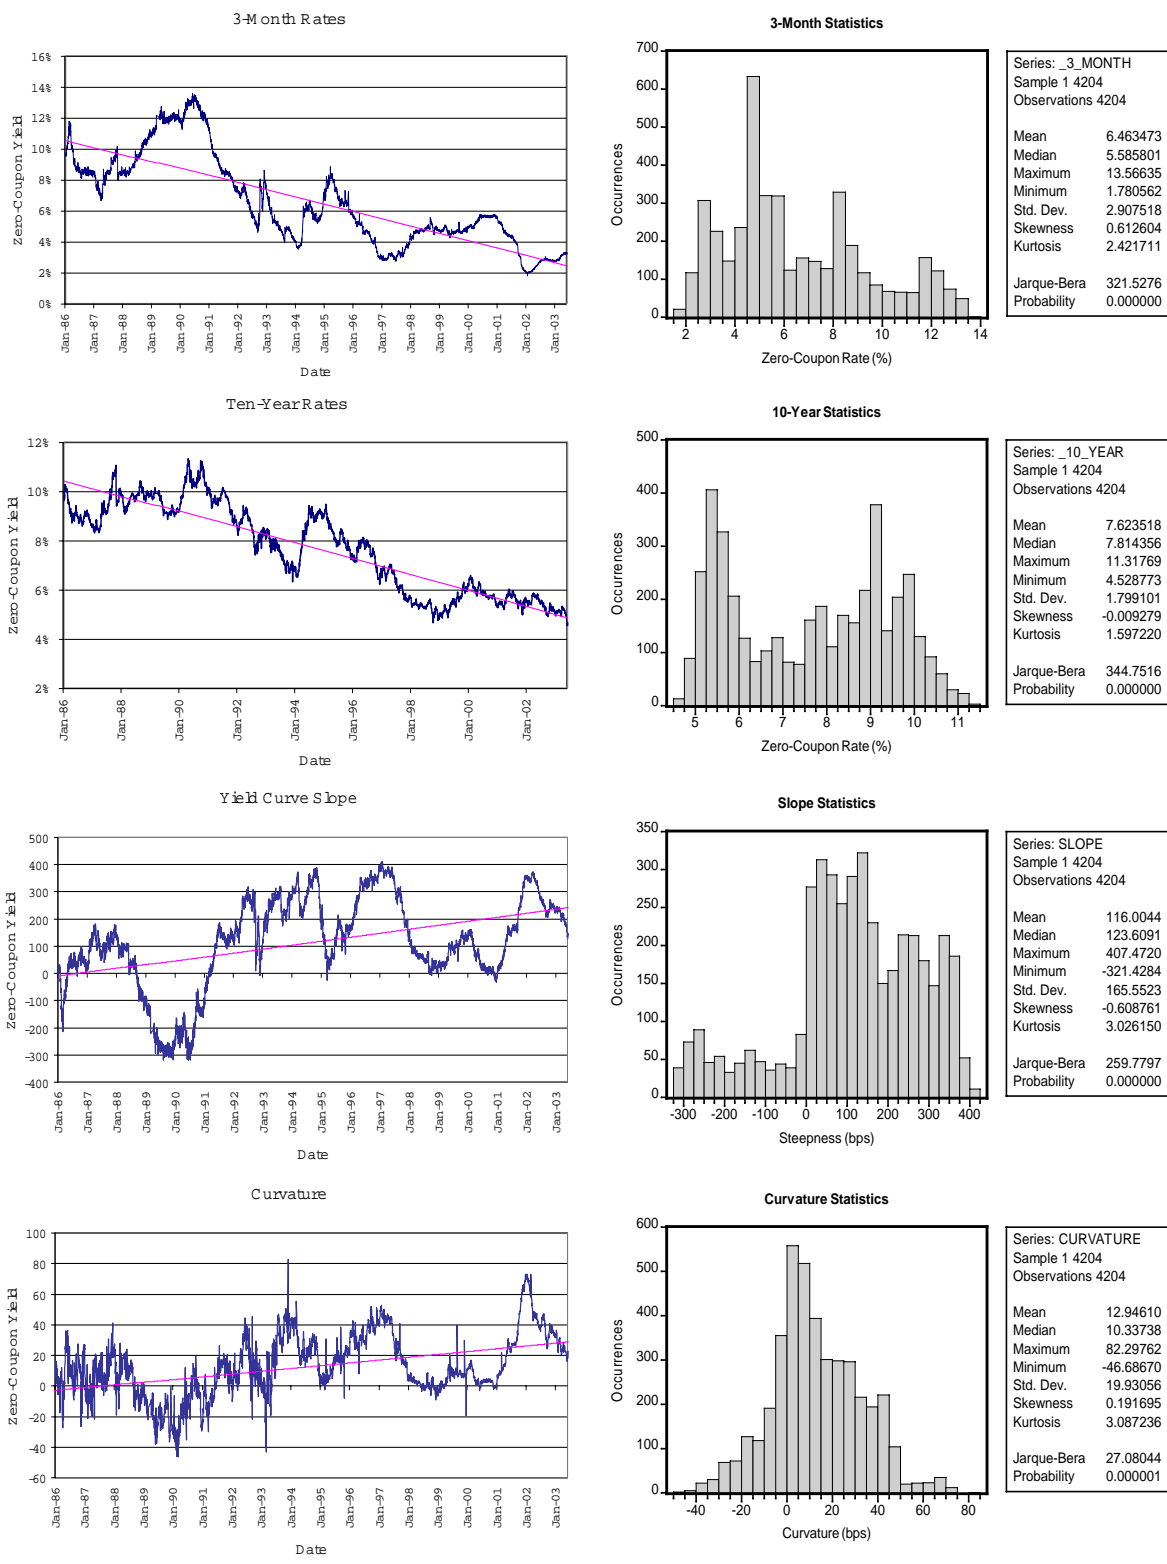


Figure 8: Yield-Curve Measures — 1986 to 1996

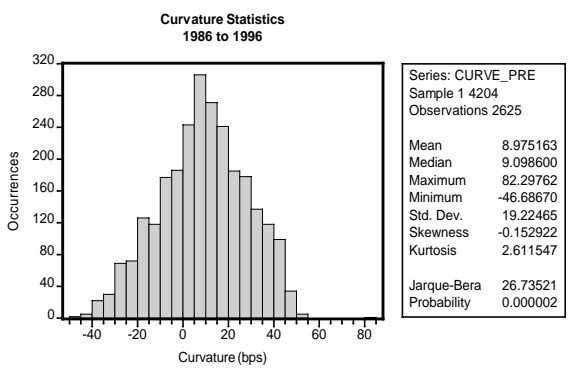
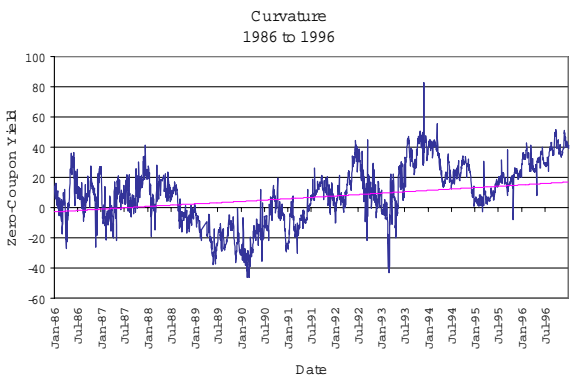
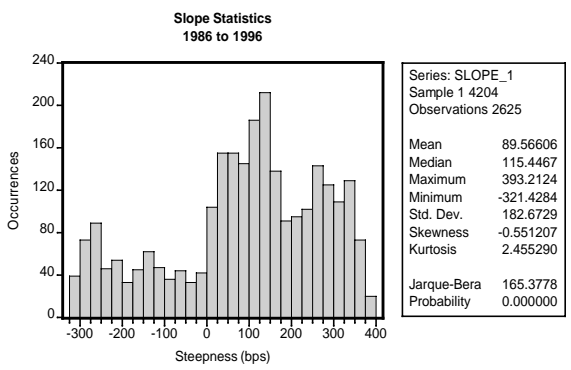
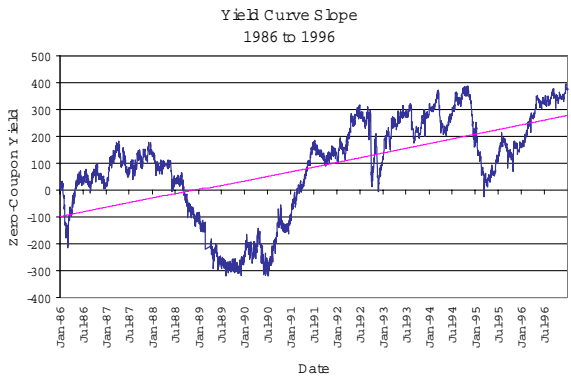
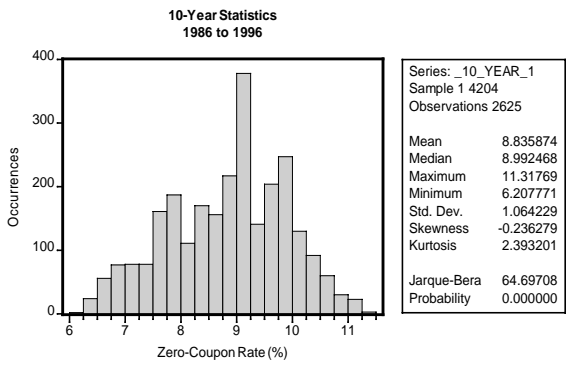
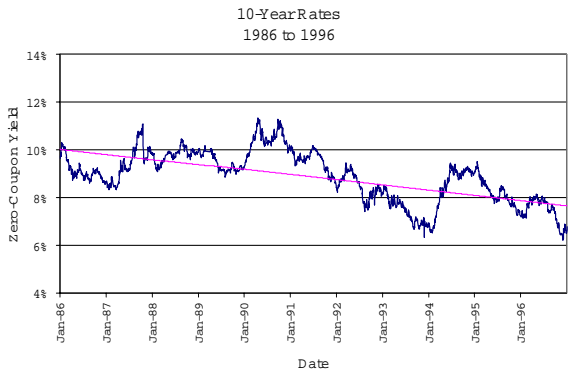
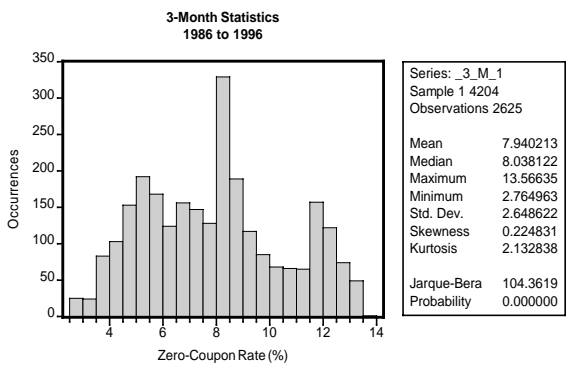
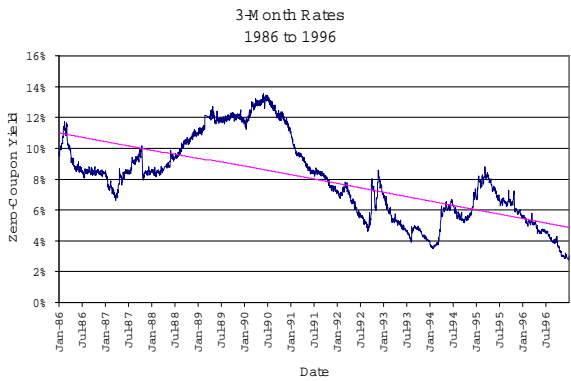


Figure 9: Yield-Curve Measures — 1997 to 2003

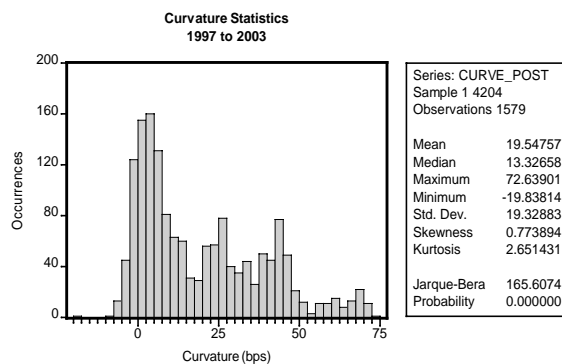
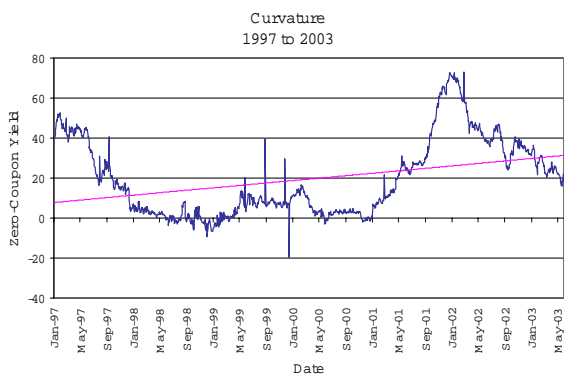
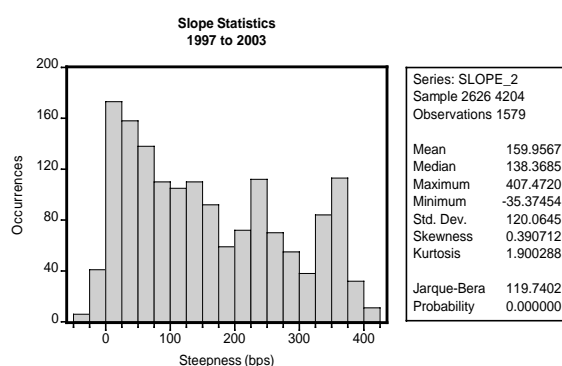
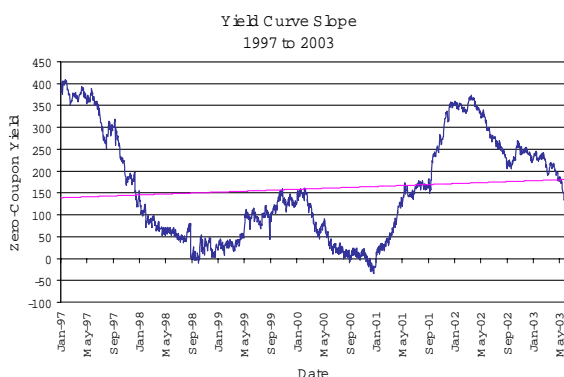
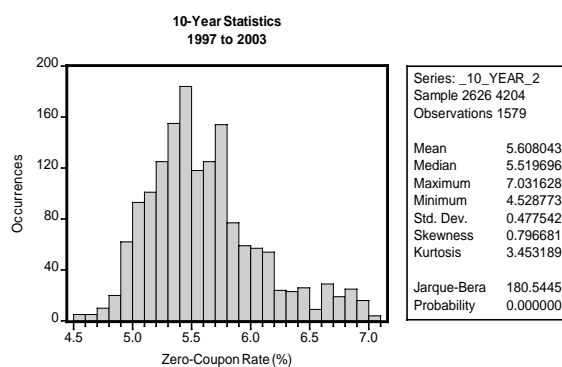
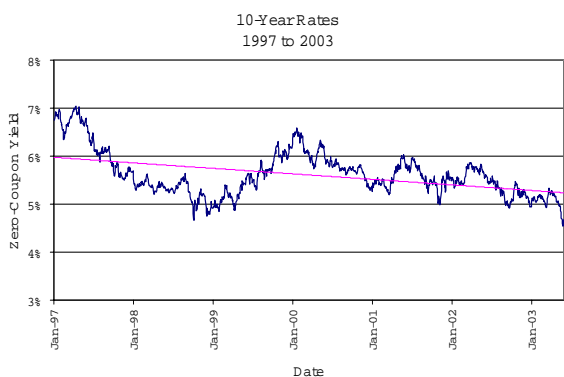
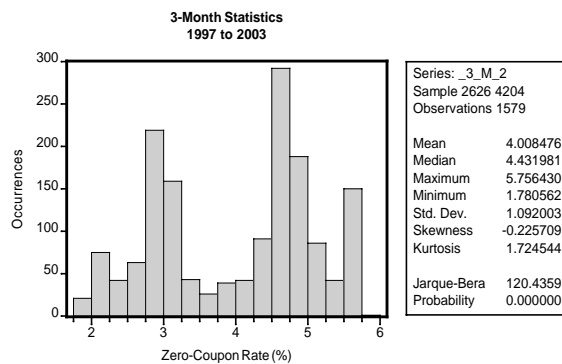
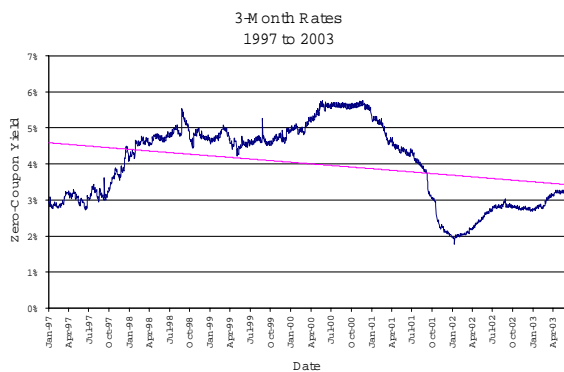


Figure 10: Yield-Curve Measure First Differences — Full Sample Period

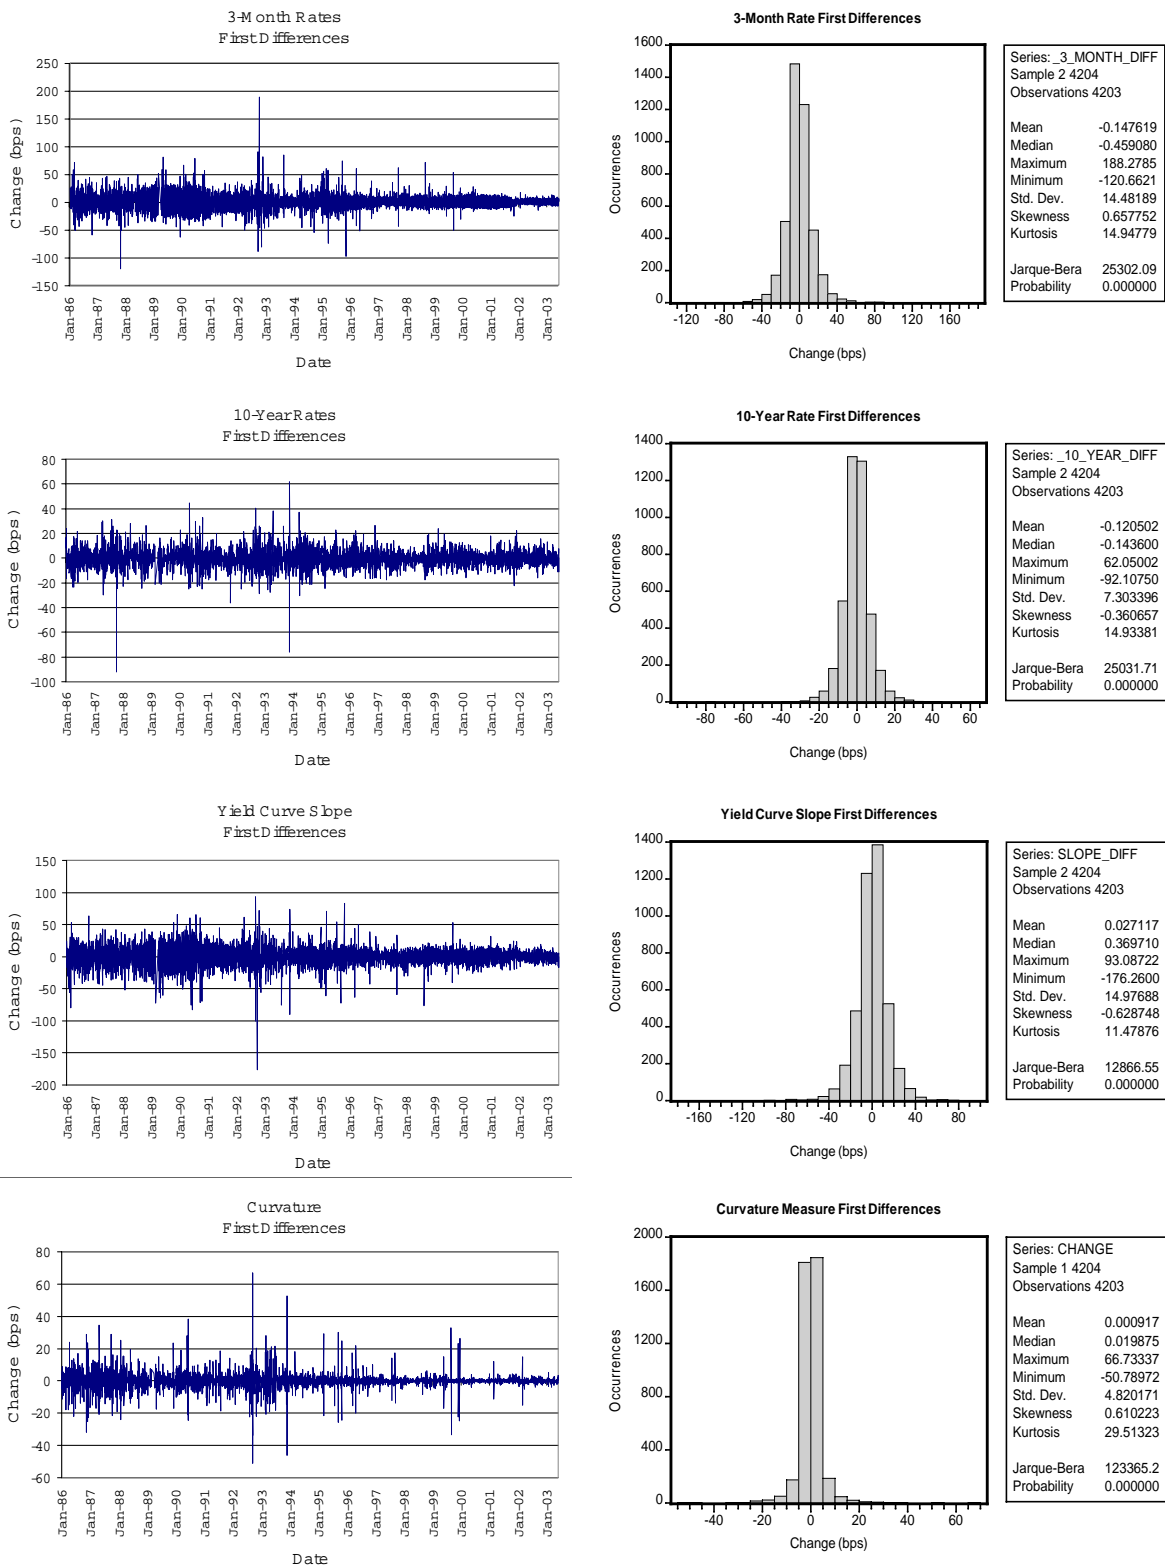
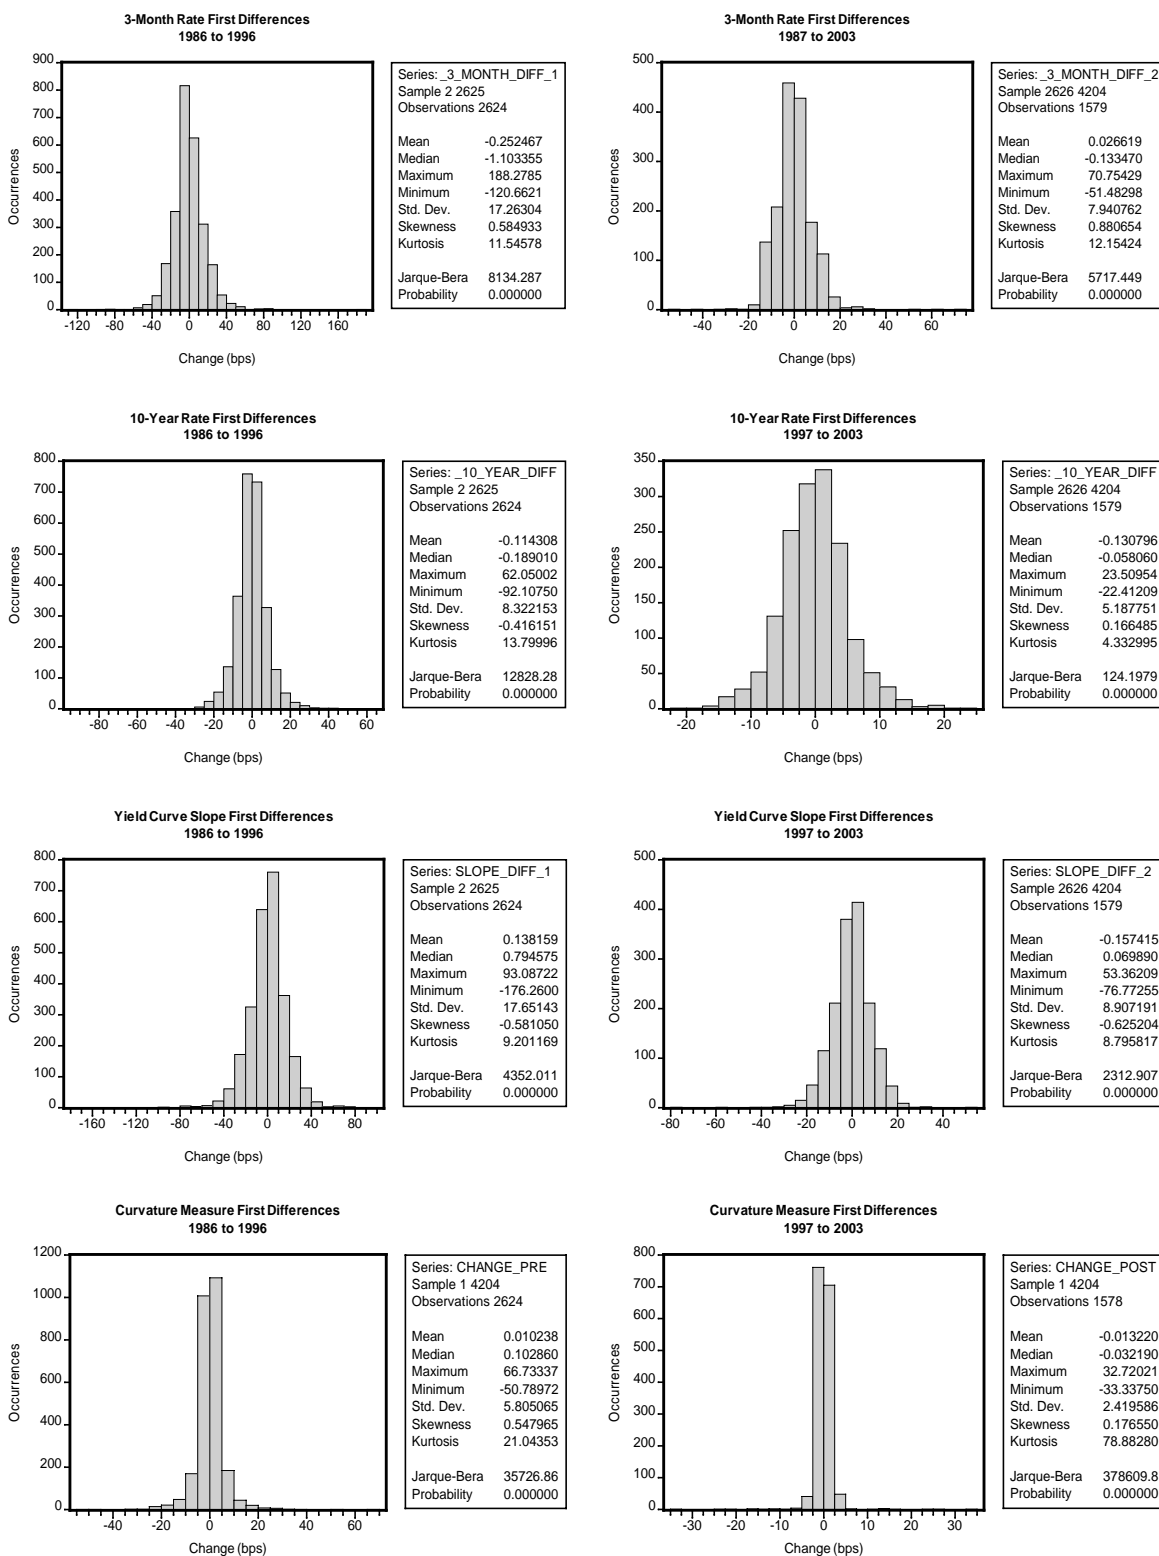


Figure 11: Yield-Curve Measures — Distributions of First Differences



Appendix: Data-Filtering Algorithm

As the yield-curve data were constructed, a small number of observations stood out as “strange” days. Specifically, this meant one of three things:

- The yield curve for the day in question did not exhibit a typical smooth, monotonic shape (with the exception of the hump in the 20- to 25-year section and inversions at the very short end).
- The curve looked very different from its immediate neighbours. That is, the curve looked very different from both the preceding two days and the following two days.
- Pricing errors were uncharacteristically large on the given day (or the model fit was very poor).

It is natural that some days in our sample have poor data. Two examples would be days on which there were data entry errors and days with stale price data for some bond issues. Either of these mistakes could result in an inaccurate yield curve exhibiting one, or both, of the problems described above. It is tempting to simply write such days off and delete them from the sample, ascribing their “strange” shape to errors of the type described above. However, it is impossible to tell in hindsight whether a given “bad” curve has an atypical shape because of raw data problems or because of unusual market conditions. For example, consider a curve with the “strange” shape, and let this be the curve corresponding to day t_i . If the curves from the previous two days, t_{i-2} and t_{i-1} , look similar to each other, and the following two days, t_{i+1} and t_{i+2} , look similar to each other, but differ from t_{i-2} and t_{i-1} , then it may simply be that the market was in an adjustment period on day t_i , and as such the curve looks “strange” due to the fact that it was in the process of undergoing a significant change. Given this type of ambiguity, it is difficult to objectively sort through the collection of yield curves and determine those days that truly should be discarded.

To make the process as objective as possible, the following filtering algorithm was developed:

- (i) For a given year, the mean and standard deviation of both the yield root mean square error and the yield mean average error were calculated. These are denoted by \overline{YRMSE} , σ_{YRMSE} , \overline{YMAE} , and σ_{YMAE} , respectively.
- (ii) We computed $\overline{YRMSE} + \sigma_{YRMSE}$ and $\overline{YMAE} + \sigma_{YMAE}$.
- (iii) We identified days, indexed by i , such that $YRMSE_i > \overline{YRMSE} + \sigma_{YRMSE}$ or $YMAE_i > \overline{YMAE} + \sigma_{YMAE}$.

- (iv) For each of these days, the bond with the largest pricing error in yield terms was eliminated. The algorithm was then rerun for that day.
- (v) The new results were evaluated based on two criteria. First, was there a large decrease in pricing errors when the bond was excluded? Second, did the curve fit its neighbours better after excluding the bond? If the answer to these questions was yes, the curve was kept. If excluding the bond made little difference, the curve was discarded. Neither of these criteria were necessary or sufficient in and of themselves, and in some cases curves were kept that satisfied one condition but not the other.

As noted earlier, these criteria are somewhat subjective. It was necessary to balance the risk of throwing out a curve that looked suspicious, but contained valid information, against keeping a day that contained inaccurate information. Ultimately, the algorithm erred on the side of keeping potentially erroneous curves. Since the number of suspect curves was very small relative to the entire sample, it was felt that the potential inclusion of a small number of erroneous curves would not significantly distort the results.

Bank of Canada Working Papers

Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. *Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.*

2004

- 2004-47 The Monetary Origins of Asymmetric Information in International Equity Markets G.H. Bauer and C. Vega
- 2004-46 Une approche éclectique d'estimation du PIB potentiel pour le Royaume-Uni C. St-Arnaud
- 2004-45 Modelling the Evolution of Credit Spreads in the United States S.M. Turnbull and J. Yang
- 2004-44 The Transmission of World Shocks to Emerging-Market Countries: An Empirical Analysis B. Desroches
- 2004-43 Real Return Bonds, Inflation Expectations, and the Break-Even Inflation Rate I. Christensen, F. Dion, and C. Reid
- 2004-42 International Equity Flows and Returns: A Quantitative Equilibrium Approach R. Albuquerque, G. Bauer, and M. Schneider
- 2004-41 Characterization of the Dynamic Effects of Fiscal Shocks in a Small Open Economy N. Rebei
- 2004-40 Prévission et analyse de la production manufacturière au Canada : comparaison de modèles linéaires et non linéaires F. Demers
- 2004-39 A Forecasting Model for Inventory Investments in Canada M. Chacra and M. Kichian
- 2004-38 Finance Constraints and Inventory Investment: Empirical Tests with Panel Data R. Cunningham
- 2004-37 The Implications of Transmission and Information Lags for the Stabilization Bias and Optimal Delegation J.-P. Lam and F. Pelgrin
- 2004-36 Optimal Taylor Rules in an Estimated Model of a Small Open Economy S. Ambler, A. Dib, and N. Rebei
- 2004-35 The U.S. New Keynesian Phillips Curve: An Empirical Assessment A. Guay and F. Pelgrin
- 2004-34 Market Valuation and Risk Assessment of Canadian Banks Y. Liu, E. Papakirykos, and M. Yuan

Copies and a complete list of working papers are available from:

Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: <http://www.bankofcanada.ca>

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : <http://www.banqueducanada.ca>