

Waggoner, Daniel F.

Working Paper

Spline methods for extracting interest rate curves from coupon bond prices

Working Paper, No. 97-10

Provided in Cooperation with:

Federal Reserve Bank of Atlanta

Suggested Citation: Waggoner, Daniel F. (1997) : Spline methods for extracting interest rate curves from coupon bond prices, Working Paper, No. 97-10, Federal Reserve Bank of Atlanta, Atlanta, GA

This Version is available at:

<https://hdl.handle.net/10419/100789>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Spline Methods for Extracting Interest Rate Curves from Coupon Bond Prices

Daniel F. Waggoner

Federal Reserve Bank of Atlanta

Working Paper 97-10

November 1997

Abstract: Cubic splines have long been used to extract the discount, yield, and forward rate curves from coupon bond data. McCulloch used regression splines to estimate the discount function, and, more recently, Fisher, Nychka, and Zervos used smoothed splines, with the roughness penalty selected by generalized cross-validation, to estimate the forward rate curve. I propose using a smoothed spline but with a roughness penalty that can vary across maturities, to estimate the forward rate curve. This method is tested against the methods of McCulloch and Fisher, Nychka, and Zervos using monthly bond data from 1970 through 1995.

JEL classification: G12, C13

Key words: term structure, smoothing splines, generalized cross-validation

The author thanks Rob Bliss for his many helpful suggestions concerning variable roughness penalties and Mark Fisher for providing code and help in implementing the smoothed spline approach. The views expressed here are those of the author and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the author's responsibility.

Please address questions regarding content to Daniel F. Waggoner, Senior Quantitative Analyst, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713, 404/521-8278, 404/521-8956 (fax), daniel.f.waggoner@atl.frb.org.

Questions regarding subscriptions to the Federal Reserve Bank of Atlanta working paper series should be addressed to the Public Affairs Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713, 404/521-8020. The full text of this paper may be downloaded (in PDF format) from the Atlanta Fed's World-Wide Web site at http://www.frbatlanta.org/publica/work_papers/.

Spline Methods for Extracting Interest Rate Curves from Coupon Bond Prices

Section 1: Introduction

Spline methods have long been used to extract discount, yield, and forward rate curves from prices of coupon bonds. McCulloch (1975) proposed using regression cubic splines to extract the discount function. This method works well, in the sense that, both in-sample and out-of-sample, it accurately prices bonds and is stable (Bliss 1997, Waggoner 1996). However, the forward rate curves produced by McCulloch's method often tend to oscillate. Though there is no theoretical result that implies that the forward rate curve can not oscillate, most practitioners would prefer one that did not exhibit this behavior, particularly if it were stable and priced bonds well. Fisher, Nychka, and Zervos (1995), proposed using a cubic spline with a roughness penalty to extract the forward rate curve. The roughness penalty stiffens the spline, which reduces the oscillatory behavior, but also reduces the fit. The roughness penalty is chosen by a generalized cross-validation method (GCV) to regulate the trade-off between goodness of fit and the stiffness of the spline. Bliss (1997) has found that the method of Fisher tends to misprice short maturity securities. We propose using Fisher's method, but with a variable roughness penalty. Using a small roughness penalty on the short end of the term structure and a larger penalty on the long end allows the flexibility to price short term securities well without giving up the desirable oscillation damping on the long end.

This paper is organized as follows. In Section 2, we specify the notation to be used throughout the paper and develop the bond pricing model. In Section 3, we discuss interest rate curve extraction techniques. In Section 4, we present the findings of our empirical tests. Section 5 contains an explanation of our particular choice of variable roughness penalty. Section 6

explores the generalized cross-validation method used by Fisher to select the level of the roughness penalty. Our conclusions are presented in Section 7.

Section 2: Interest Rate Curves and Bond Prices

We quickly review the definition of the discount function, the yield curve and the forward rate curve and their relation to the price of a coupon bond. The discount function, denoted by $\delta(t)$, is the current price of a risk-free, zero-coupon bond paying one dollar at time t . We shall use $y(t)$ to denote the zero-coupon yield curve, and $f(t)$ to denote the instantaneous forward rate curve.

These are related to the discount function via the equations

$$(2.1) \quad \delta(t) = \exp(-ty(t))$$

and

$$(2.2) \quad \delta(t) = \exp\left(-\int_0^t f(s)ds\right).$$

The term interest rate curve can be used to generically refer to any one of these three related curves.

In a world with complete markets and no taxes or transaction costs, absence of arbitrage implies that the price of any coupon bond can be computed from an interest rate curve. In particular, if the principal and interest payment of a bond is c_j dollars at time t_j , for $1 \leq j \leq K$, then the pricing equation for the bond is

$$(2.3) \quad \sum_{j=1}^K c_j \delta(t_j) = \sum_{j=1}^K c_j \exp(-t_j y(t_j)) = \sum_{j=1}^K c_j \exp\left(-\int_0^{t_j} f(s)ds\right).$$

In the real world of taxes and transaction costs, we would expect the price of a coupon bond to be only approximated by (2.3).

We are not interested in pricing a bond, given an interest rate curve, but in estimating an interest rate curve, given a set of bond prices. We develop the model and notation that we will use in the remainder of this paper. Let $\{B_i\}_{1 \leq i \leq N}$ be a set of bonds, let $\tau_1 < \tau_2 < \dots < \tau_K$ be the set of dates on which principal and interest payments occur, let $c_{i,j}$ be the principal and interest payment of the i^{th} bond on date τ_j , and let P_i be the observed price¹ of the i^{th} bond. The pricing equation is

$$(2.4) \quad P_i = \hat{P}_i + \varepsilon_i,$$

where \hat{P}_i is defined by

$$(2.5) \quad \hat{P}_i = \sum_{j=1}^K c_{i,j} \delta(\tau_j) = \sum_{j=1}^K c_{i,j} \exp(-\tau_j y(\tau_j)) = \sum_{j=1}^K c_{i,j} \exp\left(-\int_0^{\tau_j} f(s) ds\right).$$

Since our simple model omits such obvious factors as taxes and liquidity, the error term, ε_i , will contain both systematic and random factors. When we wish to make explicit \hat{P}_i 's dependence on either the discount function, the yield curve, or the forward rate curve, we write $\hat{P}_i^\delta(\cdot)$, $\hat{P}_i^y(\cdot)$, or $\hat{P}_i^f(\cdot)$, respectively. We use the notation $\hat{P}_i^*(\cdot)$ when we do not wish to specify the particular interest rate curve. In the next section, we explore techniques for using this model to obtain estimates of interest rate curves.

Section 3: Curve Extraction and Cubic Splines

All of our extraction methods use cubic splines as the functional form for either the discount or forward rate curve. A function g , defined on the interval $[t_1, t_k]$, is a cubic spline with node

¹ In practice, we do not observe a single price, but a bid and an asked quote. We define the observed price to be the average of the bid and asked quotes.

points $t_1 < t_2 < \dots < t_k$, if

(1) g is a cubic polynomial on each of the subintervals $[t_{j-1}, t_j]$, for $1 < j \leq k$.

(2) g is twice continuously differentiable over the entire interval $[t_1, t_k]$ ².

A natural choice for the node points would be a subset of $\{\tau_1, \tau_2, \dots, \tau_k\}$, the set of cash flow dates³. If we were to use all of these points, then cubic spline interest rate curves would be able to price bonds as well as any other functional form. However, we would also like an estimating technique to produce “reasonable” interest rate curves. Cubic splines, particularly ones with a large number of node points, tend to oscillate. We view excessive oscillations, particularly at longer maturities, as unreasonable behavior. In a risk-neutral world, interest rate curves contain information concerning both current and expected prices of zero-coupon bonds. Large oscillations in an interest rate curve can imply oscillations in expected prices. Though it is perfectly reasonable for the current price of a six-month bill to be \$95 with the expectation that in a year a six-month bill will sell for \$96 and that in two years its price will be \$94. It is more of a stretch to expect that in 20 years the price of a six-month bill will be \$95 dollars, in 21 years it will sell for \$96, and that in 22 years its price will be \$94. For this reason, we prefer interest rate curve extraction techniques that produce curves that are less likely to oscillate, particularly at longer maturities. There are several methods which one could use to reduce the oscillations and increase the smoothness of a cubic spline. McCulloch used regression splines for this purpose, and Fisher, Nychka, and Zervos used smoothed splines. We propose using a modification of the smoothed spline.

² See Ahlberg, Nilson and Walsh (1967) for a more complete discussion of splines and their various applications.

³ In addition, we include $\tau_0 \equiv 0$ as a node point since we want the interest rate curve to be defined over the interval from zero to the longest bond in the sample.

A regression spline reduces the oscillatory behavior by reducing the number of node points. The flexibility of a cubic spline over an interval is determined by the number of node points in that interval. By controlling number and spacing of node points, one can reduce oscillations at longer maturities, while retaining flexibility at shorter maturities. Once the placement of the nodes has been determined, the interest rate curve, ψ , is chosen to be the cubic spline which minimizes the objective function

$$(3.1) \quad \sum_{i=1}^N (P_i - \hat{P}_i^*(\psi))^2 .$$

A smoothed spline controls oscillations by imposing a roughness penalty in the objective function, as opposed to reducing the number of node points. The interest rate curve, ψ , is chosen to minimize the objective function

$$(3.2) \quad \sum_{i=1}^N (P_i - \hat{P}_i^*(\psi))^2 + \lambda \int_0^{\tau_k} [\psi''(t)]^2 dt$$

over the space of all cubic splines with node points $\tau_0 < \tau_1 < \tau_2 < \dots < \tau_k$ ⁴. Minimizing this expression is a trade-off between minimizing the first term, which measures the goodness of fit, and the second term, which measures smoothness. The positive constant λ determines the tradeoff between fit and smoothness and is called the roughness penalty. If λ were zero, then we would be in the regression spline case, and as λ increases, g tends to a linear function. The flexibility of the spline is determined by both the spacing of the nodes and the magnitude of λ , but as λ increases, the spacing of the nodes becomes less important. Thus for large values of λ , the flexibility of the spline is approximately the same across all regions. This is problematic,

⁴ For the discount function or the yield curve, it can be shown that the function which minimizes (3.2) over the space of all twice continuously differentiable functions will be a cubic spline with node points $\tau_0 < \tau_1 < \dots < \tau_k$. For the forward rate curve, a quartic spline will be the minimizing function (Reinsch 1971).

since we want the spline to be more flexible on the short end than on the long end. This leads to the modified smoothed spline.

A modified smoothed splines estimates the interest rate curve, ψ , by minimizing

$$(3.3) \quad \sum_{i=1}^N (P_i - \hat{P}_i^*(\psi))^2 + \int_0^{\tau_K} \lambda(t) [\psi''(t)]^2 dt$$

over the space of all cubic splines with node points $\tau_0 < \tau_1 < \tau_2 < \dots < \tau_K$. By allowing the roughness penalty to vary across maturities, we can damp oscillations on the long end, while retaining flexibility on the short end.

*McCulloch's Method*⁵

McCulloch proposed using a regression cubic spline to approximate the discount function. The suggested number of node points is approximately the square root of the number of bonds used in the estimation and are spaced so that roughly a equal number of bonds mature between adjacent nodes. Though the number and spacing of the node points is ad hoc, this choice works well in practice (Bliss 1997). The discount function is constrained to satisfy $\delta(0) = 1$. With this choice of nodes and constraints, the discount function is chosen to be the cubic spline which minimizes

$$(3.4) \quad \sum_{i=1}^N (P_i - \hat{P}_i^\delta(\delta))^2 .$$

Because McCulloch works with the discount function, the minimizing function can be easily found using least squares.

⁵ See McCulloch (1975) for a more complete description of this method.

*Fisher's Method*⁶

Fisher proposed using a smoothed cubic spline to approximate the forward rate curve. The recommended number of nodes is approximately one third the number of bonds used in the estimation and the nodes should be spaced so that roughly an equal number of bonds mature between adjacent nodes⁷. The forward rate curve is chosen to be the cubic spline which minimizes the expression

$$(3.5) \quad \sum_{i=1}^N (P_i - \hat{P}_i^f(f))^2 + \lambda \int_0^{\tau_k} [f''(s)]^2 ds.$$

The value of λ is determined by generalized cross-validation (GCV)⁸. In particular, λ is chosen to minimize the expression

$$(3.6) \quad \gamma(\lambda) = \frac{\text{RSS}(\lambda)}{(N - \text{ep}(\lambda))^2},$$

where

- N is the number of bonds,
- $\text{RSS}(\lambda)$ is the residual sum of squares. More formally, if f_λ is the forward rate curve which minimizes (3.5), then $\text{RSS}(\lambda) = \sum_{i=1}^N (P_i - \hat{P}_i^f(f_\lambda))^2$,
- $\text{ep}(\lambda)$ is the effective number of parameters⁹,
- θ is the cost or tuning parameter.

⁶ See Fisher, Nychka, and Zervos (1995) for a more complete description of this method.

⁷ Theoretically, the set of nodes should be the set of all dates on which a cash flow occurs. Using this much smaller set of nodes does not substantially change the resulting forward rate curve, but does reduce the computations needed to implement this method.

⁸ See Wahba (1990) for a discussion of generalized cross-validation and its properties.

⁹ In some sense, by imposing a roughness penalty we are reducing the number of parameters in our model. When λ is zero, the effective number of parameters is two more than the number of node points. As λ tends to infinity, the effective number of parameters approaches two. See Fisher, Nychka, and Zervos (1995) for a rigorous definition of the effective number of parameters.

In plain-vanilla GCV, θ is equal to one. In general, a larger value of θ tends to produce a stiffer spline. In our work, we follow Fisher, Nychka, and Zervos (1995), and set θ equal to two. Because Fisher's method works with the forward rate function, non-linear techniques must be used to find the minimizing function. For a fixed λ , (3.5) can be minimized using non-linear least squares. GCV can be implemented by using any of the standard line searches, such as the golden section search or Brent's method, however, since $\gamma(\lambda)$ can have multiple local minima, it is important to do a grid search in conjunction with the line search.

Variable Roughness Penalty (VRP) Method

As in Fisher Nychka, and Zervos (1995), we propose using a cubic spline to approximate the forward rate function with the number of nodes chosen to be approximately one third the number of bonds used in the estimation and spaced so that roughly a equal number of bonds mature between adjacent nodes. However, the cubic spline forward rate curve is chosen to minimize the function

$$(3.7) \quad \sum_{i=1}^N (P_i - \hat{P}_i(f))^2 + \int_0^{\tau_k} \lambda(s) [f''(s)]^2 ds.$$

We take $\lambda(t)$ to be

$$(3.8) \quad \lambda(t) = \begin{cases} 0.1 & 0 \leq t \leq 1 \\ 100 & 1 \leq t \leq 10, \\ 100,000 & 10 \leq t \end{cases}$$

for t measured in years¹⁰. This choice of penalty function will be more fully explored in Section 5. The VRP method is non-linear and can be implemented using non-linear least squares.

¹⁰ Fisher, Nychka, and Zervos (1995), use daily interest rates, while we use annual rates. To compare roughness penalties, ours must be multiplied by a factor of $365^5 \approx 6.5 \times 10^{12}$.

Section 4: Empirical Results

Following Bliss (1997), we tested the three methods described in Section 3 by comparing their in-sample and out-of-sample performance in correctly pricing bonds. We used the monthly CRSP bond data from 1970 through 1995 in our test, excluding bills with less than 30 days to maturity, notes and bonds with less than a year to maturity, callable bonds, and flower bonds. For each month, we divided the securities into two groups by putting every other security, ordered by maturity, in the same group. One group was used to extract the interest rate curve and perform in-sample tests, and the other group was used to perform out-of-sample tests. We also ensured that the bond of longest maturity was in the in-sample group. We used two measures of pricing error, weighted mean absolute error and hit rate. The weighted mean absolute error (WMAE) is the average distance between the midpoint of the bid and ask and the computed price, with the weighting by the inverse of duration. The hit rate is the percentage of computed prices that lie between the bid and asked quotes.

Table 1a, gives the in-sample results. In-sample, Fisher performs slightly better than McCulloch in pricing securities with more than one year to maturity, but performs much worse in pricing securities with less than a year to maturity. Since there are many more short term securities than long term securities, the nodes tend to be more concentrated on the short end. In the absence of a roughness penalty, the concentration of the nodes determines the flexibility of the spline. This allows McCulloch's spline to be more flexible on the short end. If a roughness penalty is imposed, the magnitude of the penalty influences the flexibility of the spline more than the concentration of nodes. Since Fisher uses a roughness penalty that is constant across maturities, it is not significantly more flexible on the short end than on the long end. Thus it is difficult to both dampen oscillations at the long end and price short term securities well.

By moving to a variable roughness penalty, we retain flexibility on the short end, while damping oscillations on the long end, and thus are better able to price short term securities. Again referring to Table 1a, we see that in-sample the VRP method performs slightly better than McCulloch across all maturities. It performs much better than Fisher for securities with less than a year to maturity, slightly better for securities with one to five years to maturity, and slightly worse for securities with more than five years to maturity.

In-sample, one can always increase performance by increasing the flexibility of the spline. However, increasing the flexibility of the spline too much will generally cause out-of-sample performance to decline. In fact, one could determine the optimum level of flexibility by maximizing the out-of-sample fit. Thus it is always important to consider out-of-sample tests. Table 1b presents the out-of-sample performance of the three methods. The results of out-of-sample comparisons similar are to the in-sample ones. McCulloch and the variable roughness penalty methods had similar performance across all maturities. Fisher performed much worse for securities with less than a year to maturity and about the same for longer maturity securities. The one exception to this characterization is the weighted mean absolute error of the Fisher method for bonds with more than 10 years to maturity. We will return to this aberration in section 5.

Next, we consider the smoothness of each method as measured by the average value of the square of the second derivative of the forward rate curve. In particular, we are comparing the following integral:

$$\frac{1}{t_2 - t_1} \int_{t_1}^{t_2} [f''(t)]^2 dt .$$

Smaller values of this integral indicate a smoother forward rate curve. The results are presented in Tables 2a and 2b. From these tables, we see that Fisher varies over a much wider range than

either McCulloch or the VRP method. The smoothness of Fisher's method is determined by the magnitude of the roughness penalty, which is chosen by generalized cross-validation. Because the roughness penalty varies over a wide range, the smoothness of Fisher's forward rate curves vary over a wide range. Also, Fisher tends to be smoother than either McCulloch or VRP on the short end. This supports our contention that Fisher's method was unable to price short term securities well because it was too stiff on the short end. Finally, McCulloch is often smoother than the VRP method, though the amount by which it was smoother was small. So the VRP method priced bonds slightly better, but McCulloch was slightly smoother.

Section 5: Choice of Roughness Penalty

Our choice for the roughness penalty is a three-tiered step function, which is constant on the intervals from 0 to 1 year, from 1 to 10 years, and 10 to 30 years. The divisions correspond to the difference between Bills, Notes, and Bonds. Though this choice is ad hoc, we have found that the behavior of the interest rate curve is not sensitive to the particular shape of the roughness penalty. For instance, the three roughness penalties pictured in figure 3 all produced similar results in terms of in-sample and out-of-sample fit. Figures 4 through 6 show the response of out-of-sample weighted mean absolute error to changes in the roughness penalty. These 3-D plots were produced by varying the value of the roughness penalty on the short end (less than 1 year) and the long end (more than 10 years), with the intermediate values taken to be the geometric average of the short and long values. These plots clearly show that if the roughness penalty is too small or too large, then the out-of-sample performance deteriorates. We chose the value of the short and long penalties so that the weighted mean absolute error would be small, but the spline would be as stiff as possible.

Section 6: Generalized Cross-Validation

We saw in section 4 that the out-of-sample performance of Fisher's method for bonds with more than 10 years to maturity was poor in terms of weighted mean absolute error. In this section we investigate the cause of this. For a few months, the maturities of the last three bonds in the sample are widely spaced and the GCV method selects a relatively small value for the roughness penalty - less than 10. In those months, an extremely large pricing error can occur for the last out-of-sample bond. An example of this type of behavior can be seen in February of 1977. The plot of the yield curve for this month is given in Figure 7. As one can see, the last three bonds in the sample are widely spaced, which, because of the relatively small value of the roughness penalty, allows the yield curve to become negative. This causes the last out-of-sample bond to be hugely mispriced. The yield curves for March and April of that year are almost identical and there are another 5 months in 1976 and 1977 with similar shape to the one pictured, though in these cases the yield curve does not become negative. When these eight months are omitted from the sample, we find that the out-of-sample weighted mean absolute error for bonds with maturity greater than 10 years is similar to the other methods.

We also compared GCV against a fixed roughness penalty level across all months (both were constant across maturities in each month). We chose the median of the roughness penalties selected by GCV as the fixed value, which was approximately 1,000. Table 8 gives the in-sample and out-of-sample performance of the fixed roughness penalty. In comparing this with table 1a and 1b, we see that overall the fixed roughness penalty performed about as well as GCV. Furthermore, from the out-of-sample WMAE for maturities greater than 10 years, we see that the fixed roughness penalty avoided the problems described above. We feel that GCV is valuable

tool for determining the appropriate range for the level of the roughness penalty, but it can produce extreme values for the roughness penalty, which can lead to catastrophic behavior.

Section 7: Conclusions

Both in-sample and out-of-sample, using a roughness penalty which varies across maturities improves the ability of Fisher's method to price short term securities. We strongly recommend using this technique when a smoothed spline is to be employed. Though the generalized cross-validation technique is valuable for determining the appropriate range for the level of the roughness penalty, we feel that it often does not produce an optimal value for the roughness penalty. Thus we do not recommend it as an automatic procedure for determining the level of the roughness penalty.

The results produced by McCulloch's method and the VRP method, with our choice of roughness penalty, are very similar both in fit and smoothness. Since McCulloch is a linear procedure, it is easier to implement than the non-linear VRP. When one is looking for a quick, easy to use method for extracting interest rate curves, McCulloch is a good choice. However, the VRP method allows us to explicitly control the amount of smoothing applied. We have specified a choice of roughness penalty that yields a somewhat better fit than McCulloch, but is also less smooth. Figures 4 through 6 indicate how out-of-sample fit varies with the choice of the roughness penalty, and can be used to choose the roughness penalty when one wants either a smoother forward rate curve or better out-of-sample pricing of bonds.

References

- Ahlberg, J., E. Nilson, and J. Walsh, *The Theory of Splines and Their Applications*, Academic Press, 1967.
- Bliss, Robert R. "Testing Term Structure Estimation Methods," *Advances in Futures and Options Research*, 9 (1997), 197-231.
- Fisher, M., D. Nychka, and D. Zervos, "Fitting the Term Structure of Interest Rates with Smoothing Splines," Working Paper 95-1, Finance and Economics Discussion Series, Federal Reserve Board, 1995.
- McCulloch, J. H. "The Tax Adjusted Yield Curve." *Journal of Finance*, 30 (June 1975), 811-29.
- Reinsch, C. H. "Smoothing by Spline Functions II," *Numerische Mathematik*, 16 (1971), 451-54.
- Wahba, G. *Spline Models for Observational Data*, SIAM, Philadelphia, 1990.
- Waggoner, D. "The Robustness of Yield and Forward Rate Extraction Methods," Unpublished working paper, September 1996.

Table 1a: In-Sample Weighted Mean Absolute Error and Hit Rates
Maturity Range

	< 1	1-3	3-5	5-10	> 10	All
McCulloch						
WMAE	0.019	0.108	0.179	0.375	0.376	0.055
Hit Rate	47.3%	42.0%	34.2%	22.4%	27.8%	37.0%
Fisher						
WMAE	0.063	0.104	0.177	0.324	0.265	0.085
Hit Rate	27.7%	45.8%	34.7%	28.2%	36.1%	35.3%
VRP						
WMAE	0.017	0.093	0.170	0.338	0.303	0.049
Hit Rate	51.2%	49.4%	36.7%	26.6%	30.4%	41.6%

Table 1b: Out-of-Sample Weighted Mean Absolute Error and Hit Rates
Maturity Range

	< 1	1-3	3-5	5-10	> 10	All
McCulloch						
WMAE	0.020	0.109	0.202	0.405	0.424	0.056
Hit Rate	45.2%	41.9%	31.3%	18.9%	24.0%	35.1%
Fisher						
WMAE	0.062	0.105	0.198	0.422	0.697	0.091
Hit Rate	26.4%	44.6%	31.4%	19.2%	28.3%	31.6%
VRP						
WMAE	0.019	0.096	0.199	0.413	0.356	0.052
Hit Rate	49.1%	48.5%	32.1%	20.0%	26.0%	38.6%

Table 2a: Minimum, Median, and Maximum of the Average Value of the Square of the Second Derivative of the Forward Rate Curve Over All Maturities

	Fisher	VRP	McCulloch
Minimum	4.3×10^{-22}	1.2×10^{-5}	1.1×10^{-6}
Median	9.5×10^{-6}	6.7×10^{-4}	1.2×10^{-4}
Maximum	0.091	0.016	0.014

Table 2b: Monthly Comparison of the Average Value of the Square of the Second Derivative of the Forward Rate Curve

	Percentage of Months McCulloch was Smoother Than Fisher	Percentage of Months McCulloch was Smoother Than VRP	Percentage of Months Fisher was Smoother Than VRP
Maturities Less Than a Year	4%	79%	99%
Maturities Greater Than a Year	50%	80%	60%
All Maturities	24%	79%	86%

Figure 3: Sample Roughness Penalty Shapes




Figure 4: Out-of-sample weighted mean absolute error of securities with less than 1 year to maturity

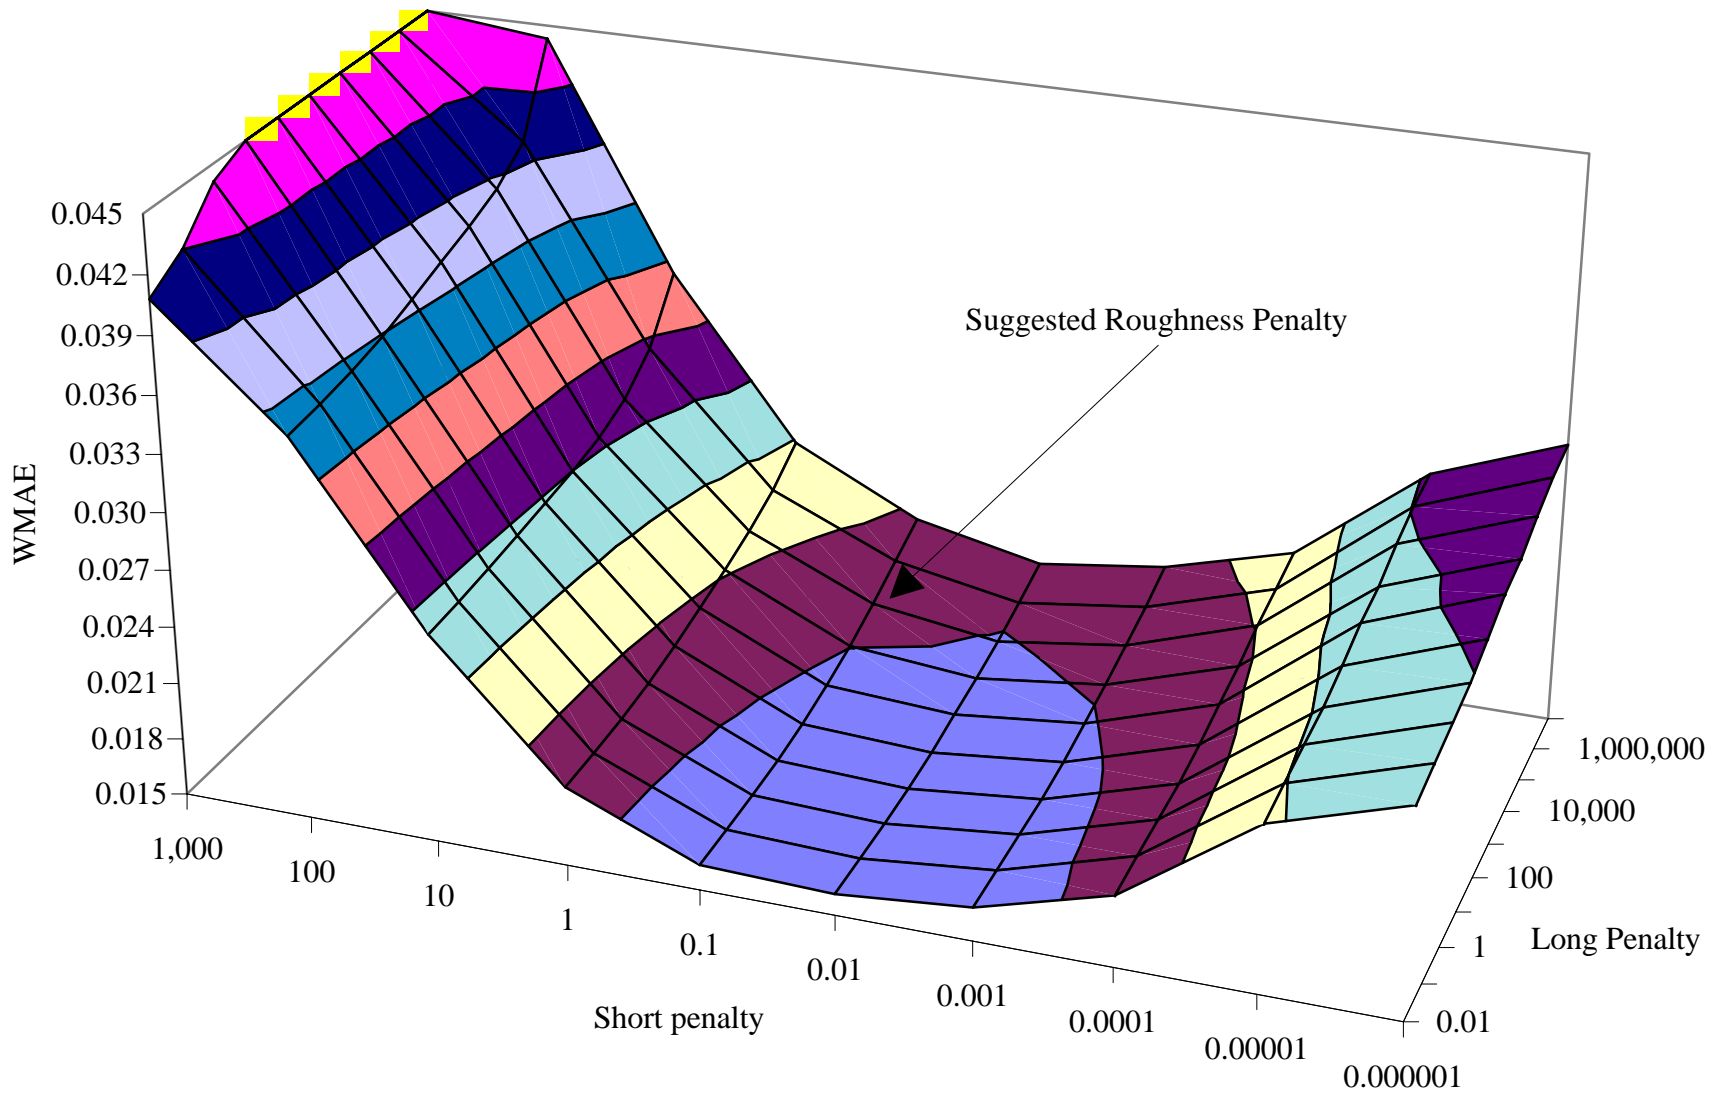


Figure 5: Out-of-sample weighted mean absolute error of securities between 1 and 10 years to maturity

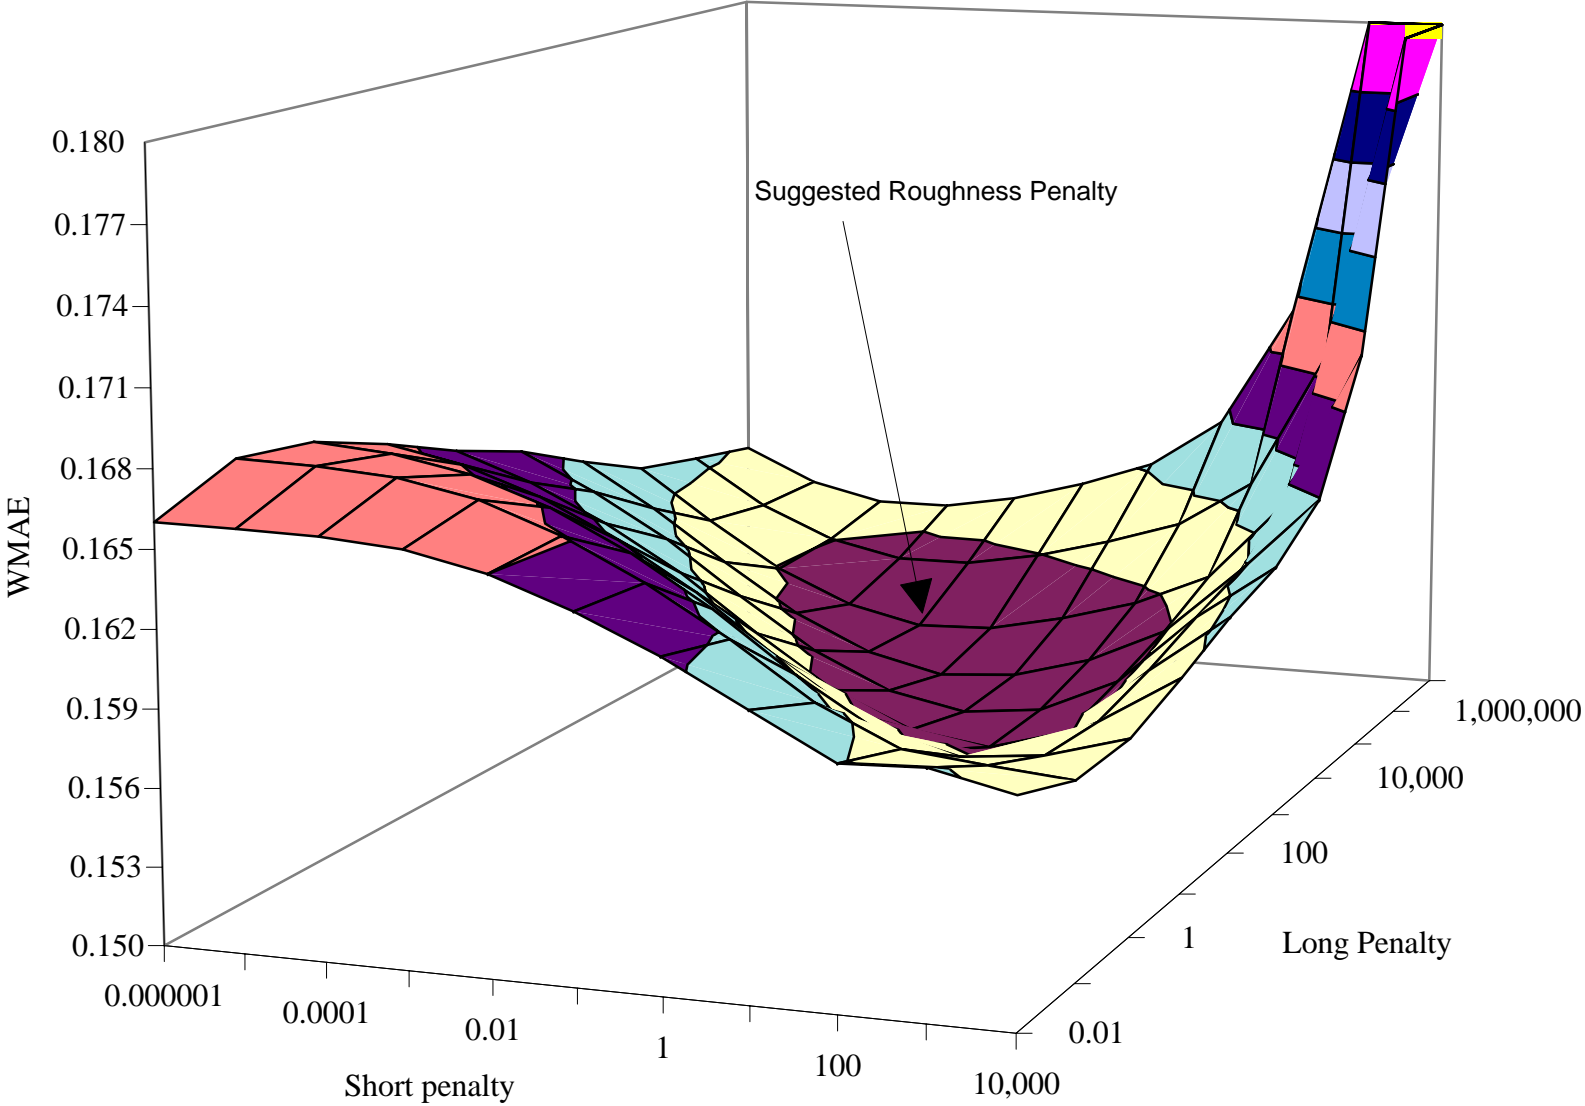


Figure 6: Out-of-sample weighted mean absolute error of securities with more than 10 years to maturity

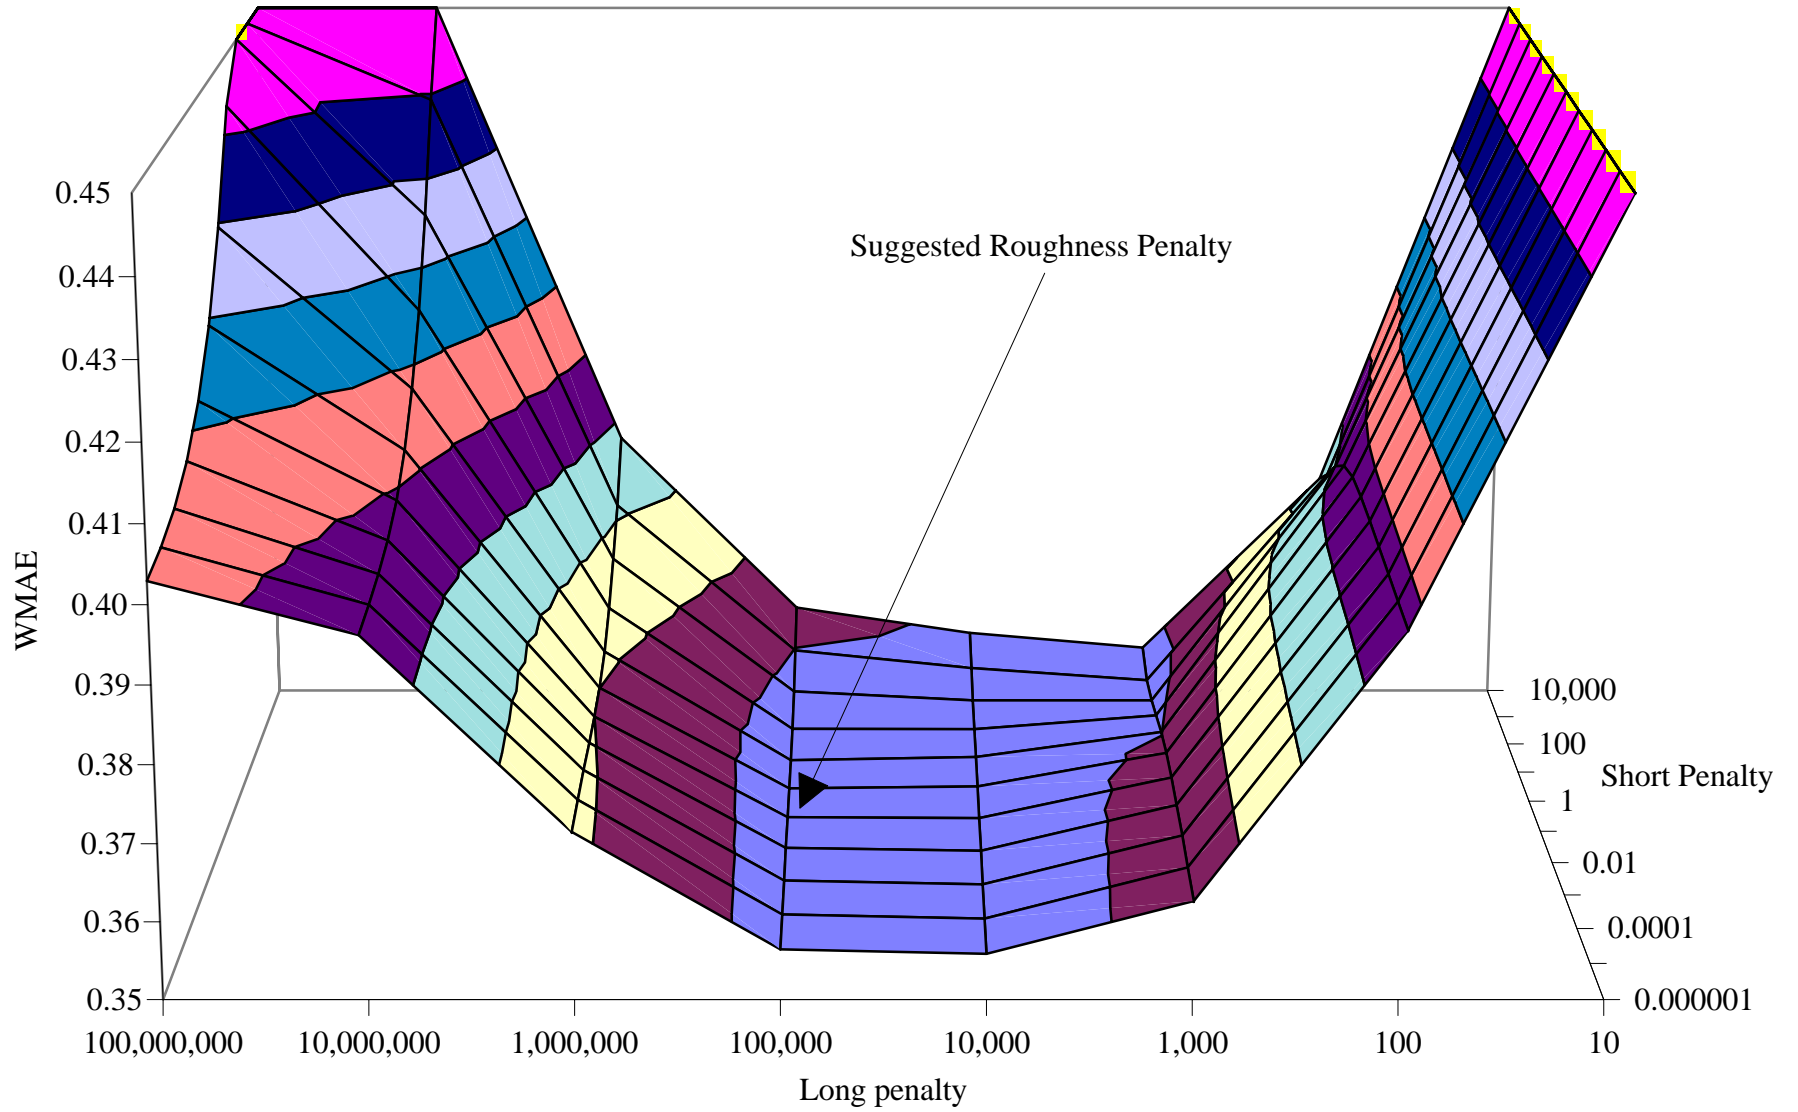


Figure 7: Fisher Yield Curve for February 28, 1977

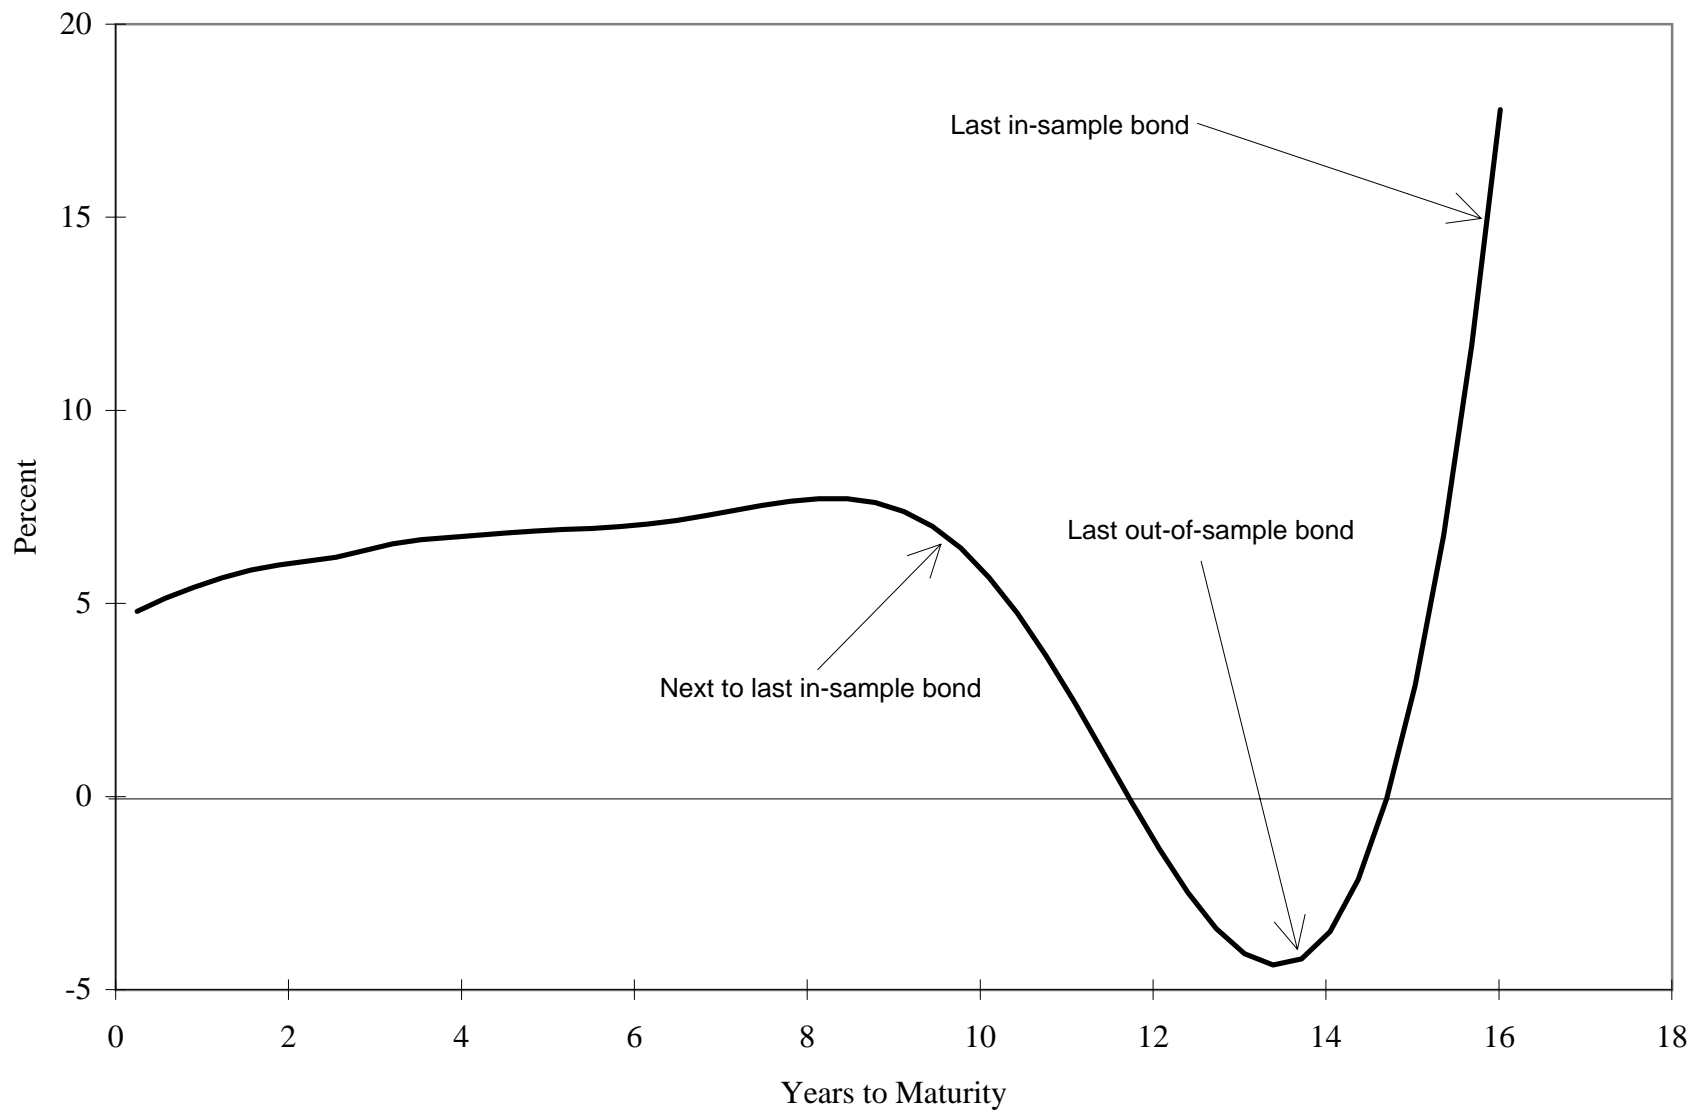


Table 8: Weighted Mean Absolute Error and Hit Rates for a Roughness Penalty Level of 1000

Maturity Range

	< 1	1-3	3-5	5-10	> 10	All
In-Sample						
WMAE	0.056	0.102	0.175	0.337	0.248	0.079
Hit Rate	18.3%	45.0%	33.6%	26.2%	36.7%	32.4%
Out-of-Sample						
WMAE	0.056	0.103	0.197	0.412	0.355	0.081
Hit Rate	18.8%	44.2%	31.1%	20.7%	29.9%	30.0%