JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS VOL. 40, NO. 2, JUNE 2005
COPYRIGHT 2005, SCHOOL OF BUSINESS ADMINISTRATION, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

Term Structure Forecasts of Long-Term Consumption Growth

Andrea Berardi and Walter Torous*

Abstract

Relying on a simple general equilibrium model of the term structure, we show that both nominal yields and real consumption growth rates can be affine in the unobservable state variables. We can then express real consumption growth rates in terms of nominal yields rather than the unobservable state variables with the coefficients of the resultant forecasting relation being endogenously determined by the term structure model. Using term structure data over the 1985 to 2000 sample period, the empirical evidence is consistent with our model more accurately predicting real consumption growth rates than a regression model based on the term spread.

I. Introduction

This paper investigates the accuracy of using nominal interest rates to forecast real consumption growth rates. We demonstrate that a simple general equilibrium model of the term structure provides more accurate forecasts than the simple term spread at horizons of one year and longer.

Beginning with Kessel (1965), many researchers including Harvey (1988), (1989), (1991), and (1993), Estrella and Hardouvelis (1991), Plosser and Rouwenhorst (1994), Chapman (1997), Kamara (1997), Roma and Torous (1997), and Hamilton and Kim (2002) have demonstrated that the term spread, that is, the difference between yields of long- and short-term bonds, provides valuable predictive information about future economic growth. In particular, a positive term spread implies a subsequent increase in economic activity, while a negative term spread is consistent with a subsequent recession. The intuition for this result is based on the desire of investors to smooth consumption. For example, if a recession is expected individuals sell short-term bonds and buy long-term bonds to receive payoffs when their consumption level drops. As a result, short-term

*Berardi, andrea.berardi@mail.univr.it, Università di Verona, SAFE Center, via Giardino Giusti 2, Verona 37129, Italy, and Torous, wtorous@anderson.ucla.edu, University of California, Anderson Graduate School of Management, 110 Westwood Plaza, Los Angeles, CA 90095. We thank Stephen Schaefer for many helpful comments on an earlier draft of this paper. The comments of an anonymous referee and Jonathan Karpoff (the editor) significantly improved the paper. We are, however, responsible for any remaining errors.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

yields increase while long-term yields decrease thereby inverting the yield curve in anticipation of a downturn in economic activity.

Distinct from previous research, this paper investigates the link between interest rates and economic growth within a simple general equilibrium framework in which the behavior of both interest rates and real consumption growth are simultaneously modeled. Following Cox, Ingersoll, and Ross (1985), we construct a general equilibrium term structure model that imposes cross-equation restrictions that endogenously link the term structure of interest rates to the dynamics of real consumption growth. Since both nominal yields and real consumption growth rates are affine in the posited but unobservable state variables, we can then express consumption growth rates in terms of nominal yields as opposed to the unobservable state variables.

As a result, we imply real consumption growth forecasts from the current nominal term structure and do not rely on consumption data to make our forecasts.[1] Since today's yields on bonds that mature at different times in the future are set by investors who take into account the levels of consumption expected at those times, the forecasts we provide are forward looking. This is in contrast to the spread model in which a historically estimated regression, which relies on past consumption growth rates, specifies the forecasting relation. Unfortunately, the longer the forecast horizon, the more dated is the estimated regression and so the less accurate are the spread model's forecasts when compared to our forecasts.

The plan of this paper is as follows. Section II details the general equilibrium model of the term structure and derives the endogenous relation between real consumption growth and nominal interest rates. In Section III, relying on term structure data over the 1985 to 2000 sample period, we compare our forecasting model to one based on the term spread and we provide statistically reliable evidence that we more accurately forecast real consumption growth rates at horizons of one year and longer. We also statistically identify the sources of these forecast gains and investigate the robustness of our results to alternative specifications and forecasting periods. Section IV concludes.

II. The Model

Our theoretical framework is based on the standard general equilibrium economy of the Cox, Ingersoll, and Ross (1985) type. The main underlying assumptions are: i) a fixed number of identical individuals with rational expectations maximizing a time-additive logarithmic utility function; ii) a competitive economy with continuous trading and no transactions costs; iii) the existence of markets for contingent claims and for instantaneous borrowing and lending at the riskless interest rate; iv) production can be allocated to consumption or investment; and v) investment opportunities consist of a stochastic production process, a set of contingent claims, and a risk-free asset.

[1] Similar to this paper, Ang, Piazzesi, and Wei (2002) impose cross-equation restrictions on the dynamics of bond yields and GDP growth rates by imposing no-arbitrage restrictions in a three-factor vector autoregressive model. More importantly, they assume that one of the state variables is lagged GDP, and they forecast GDP using past GDP values. By contrast, we do not use consumption data to predict future consumption.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

A. State Variables

We assume that the economy is characterized by two latent state variables. Litterman and Scheinkman (1991) and Brown and Schaefer (1994) empirically document that the majority of the movement in the term structure of interest rates can be explained by two factors. In our case, these factors, x, follow risk-adjusted uncorrelated Gaussian processes,[2]

$$dx = (\phi + \Gamma x)\, dt + \Sigma dz, \tag{1}$$

where

$$x \equiv \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \quad \phi \equiv \begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix}, \quad \Gamma \equiv \begin{pmatrix} \gamma_1 & 0 \\ 0 & \gamma_2 \end{pmatrix},$$

$$\Sigma \equiv \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}.$$

The solution to this stochastic differential equation for $\tau > 0$ gives

$$E_t\{x(t+\tau)\} = a(\tau) + B(\tau)x(t), \tag{2}$$

$$\mathrm{cov}_t\{x(t+\tau), x(t+\tau)'\} = F(\tau), \tag{3}$$

where $a(\tau) \equiv \Gamma^{-1}(B(\tau) - \mathrm{I})\phi, B(\tau) \equiv \exp(\Gamma\tau), F(\tau) \equiv -(\Psi - B(\tau)\Psi B(\tau)')$, and the matrix Ψ is a function of the coefficients in matrices Γ and Σ.[3]

B. Output and Consumption

We assume that a single physical good is produced, which may be allocated to consumption or investment, and that a single technology exists allowing capital to be transformed into output. Let Q denote the nominal amount of the good invested in the production process and assume that it depends on both state variables. The following stochastic differential equation then describes the dynamics of nominal output in the economy, $dQ/Q = (x_1 + x_2)dt + \sigma_Q dz_Q$. We assume that inflation in this economy is non-stochastic. Let p denote the price level and assume that it evolves according to the following deterministic process, $dp/p = \pi dt$, where π is the non-stochastic instantaneous expected inflation rate.[4] Notice here that as in Harvey (1988) we also assume a flat term structure of inflation expectations.

By applying Ito's lemma to the expression for real output, $q = Q/p$, we derive a corresponding stochastic process for q. In equilibrium, all wealth is invested in the production process and real consumption, c, must be proportional to optimally

[2] We directly specify the process for x under the risk-adjusted probability measure by assuming that the parameter ϕ includes a risk adjustment for the market price of risk. Doing so avoids the problem of identifying the market price of risk parameter when estimating contingent claims models based on Gaussian processes (see Dai and Singleton (2000)).

[3] See Langetieg (1980), footnote 22 for further details on the calculation of Ψ.

[4] Assume that a stochastic inflation rate would result in a more complicated model that would also require us to hypothesize the nature of the interaction between inflation and output growth. Furthermore, not all of the additional parameters can be separately identified when estimating the resultant term structure model.

invested wealth, $c = \delta W$.[5] Therefore, a stochastic process of the following form holds for real consumption, $dc/c = (x_1 + x_2 - \pi - \delta)dt + \sigma_c dz_c$ where $\sigma_c = \sigma_Q$ and $dz_c = dz_Q$.

Similarly, we can derive a stochastic differential equation for $\ln c$ and by integrating this expression from t to $t + \tau$ we obtain the following expression for the growth rate in consumption over the time interval $[t, t + \tau]$,

$$E_t\left\{\ln\frac{c(t+\tau)}{c(t)}\right\} = h(\tau\,;\zeta) + J(\tau\,;\zeta)x(t), \tag{4}$$

$$\begin{aligned} \text{with} \quad h(\tau\,;\zeta) &= -\left(\delta + \pi + \tfrac{\sigma_c^2}{2}\right)\tau - \tau\iota'\Gamma^{-1}\phi + J(\tau\,;\zeta)\Gamma^{-1}\phi, \\ J(\tau\,;\zeta) &= \iota'\Gamma^{-1}\left(B(\tau) - \mathrm{I}\right), \quad \text{and} \\ \iota' &= (1\ \ 1)\,, \end{aligned}$$

where we now explicitly note the dependence upon all of the model's parameters $\zeta = (\phi_1, \phi_2, \gamma_1, \gamma_2, \sigma_1, \sigma_2, \sigma_c, \delta)$, which must be estimated.

C. Term Structure of Interest Rates

In equilibrium, the current time t price of a nominal unit discount bond with maturity date $T = t + \tau$ is given by

$$G(t\,;T) = E_t\left\{\frac{Q(t)}{Q(t+\tau)}\right\}.$$

Using standard results,[6] we derive the following closed-form solution for nominal bond prices, $G(t\,;T) = \exp[g_0(\tau\,;\zeta) - g(\tau\,;\zeta)x(t)]$,

$$\begin{aligned} \text{where} \quad g_0(\tau\,;\zeta) &\equiv \tau\sigma_c^2 - \iota'\Gamma^{-1}\left[\Gamma^{-1}\left(B(\tau) - \mathrm{I}\right) - \tau\mathrm{I}\right]\phi \\ &\quad + \frac{1}{2}g(\tau\,;\zeta)\Psi g'(\tau\,;\zeta) + \frac{1}{2}\iota'\Gamma^{-1}\Sigma\Sigma(\Gamma^{-1})'\iota\tau \\ &\quad - \frac{1}{2}\iota'\left[\Gamma^{-1}\Gamma^{-1}\left(B(\tau) - \mathrm{I}\right)\Psi\right. \\ &\quad \left. + \Psi\left(B(\tau) - \mathrm{I}\right)'(\Gamma^{-1})'(\Gamma^{-1})'\right]\iota \\ g(\tau\,;\zeta) &\equiv \iota'\Gamma^{-1}\left(B(\tau) - \mathrm{I}\right). \end{aligned}$$

Therefore, zero coupon yields can be expressed as

$$Y(t\,;T) \equiv -\frac{\ln G(t\,;T)}{T-t} = \kappa_0(\tau\,;\zeta) + \kappa(\tau\,;\zeta)x(t), \tag{5}$$

$$\begin{aligned} \text{where} \quad \kappa_0(\tau\,;\zeta) &\equiv -g_0(\tau\,;\zeta)/\tau \\ \kappa(\tau\,;\zeta) &= (\kappa_1(\tau\,;\zeta), \kappa_2(\tau\,;\zeta)) \equiv g(\tau\,;\zeta)/\tau. \end{aligned}$$

[5] See Cox, Ingersoll, and Ross (1985).
[6] See, for example, Duffie (2001).

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

D. Implicit Relation between Yields and Consumption

Real consumption growth rates, expression (4), and nominal yields, expression (5), are both affine in the state variables. The closed-form nature of these expressions implies that we can express consumption growth rates in terms of yields rather than in terms of the unobservable latent factors. Consequently, we provide an endogenous means of exploiting the nominal term structure to forecast real consumption growth rates. In other words, once we have estimated the nominal term structure model, we can simply imply the resultant real consumption growth rate forecasts.

To fix matters, we express the two posited state variables in terms of two distinct yields, for instance, the yield on a short-term bond, $Y_S \equiv Y(\tau_S)$, and the yield on a long-term bond, $Y_L \equiv Y(\tau_L)$, $\tau_L > \tau_S$. Equivalently, to make our results comparable to previous forecasting models that rely on the spread (e.g., Estrella and Hardouvelis (1991) and Harvey (1989), (1991), and (1993)), we express the state variables in terms of the yield on a short-term bond, Y_S, and the spread between long- and short-term yields, $SP \equiv Y_L - Y_S$. From (5), we have

$$Z(t) \equiv \begin{pmatrix} Y_S(t) \\ SP(t) \end{pmatrix} = \bar{\kappa}_0(\zeta) + \bar{\kappa}(\zeta)x(t),$$

$$\text{where} \quad \bar{\kappa}_0(\zeta) \equiv \begin{pmatrix} \kappa_0(\tau_S\,;\zeta) \\ \kappa_0(\tau_L\,;\zeta) - \kappa_0(\tau_S\,;\zeta) \end{pmatrix}, \qquad \text{and}$$

$$\bar{\kappa}(\zeta) \equiv \begin{pmatrix} \kappa_1(\tau_S\,;\zeta) & \kappa_2(\tau_S\,;\zeta) \\ \kappa_1(\tau_L\,;\zeta) - \kappa_1(\tau_S\,;\zeta) & \kappa_2(\tau_L\,;\zeta) - \kappa_2(\tau_S\,;\zeta) \end{pmatrix}.$$

This system can be inverted with respect to the two latent state variables. Substituting the resulting expressions into expression (4), we derive the endogenous relation between expected consumption growth, the spread, and the short-term yield,

$$(6) \qquad E_t\left\{\ln\frac{c(t+\tau)}{c(t)}\right\} = \alpha(\tau\,;\zeta) + \beta(\tau\,;\zeta)Z(t)$$
$$= \alpha(\tau\,;\zeta) + \beta_1(\tau\,;\zeta)Y_S(t) + \beta_2(\tau\,;\zeta)SP(t),$$

where $\alpha(\tau\,;\zeta) \equiv h(\tau\,;\zeta) - \beta(\tau\,;\zeta)\cdot\bar{\kappa}_0(\zeta)$, $\beta(\tau\,;\zeta) = (\beta_1(\tau\,;\zeta), \beta_2(\tau\,;\zeta)) \equiv J(\tau\,;\zeta)\cdot\bar{\kappa}^{-1}(\zeta)$.

Expression (6) gives the consumption growth forecasting relation implied by the general equilibrium term structure model. By construction, this forecasting model depends on both the short-term yield and the spread. More importantly, the coefficients on the short-term yield, $\beta_1(\tau\,;\zeta)$, and on the spread, $\beta_2(\tau\,;\zeta)$, are endogenously determined and depend explicitly on the model's parameters ζ. As a consequence, we do not use consumption data to construct our forecasts as all relevant information about future consumption growth is captured in general equilibrium by the term structure.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

By comparison, forecasting models that rely on the spread are implemented as

$$E_t\left\{\ln\frac{c(t+\tau)}{c(t)}\right\} = \hat{a}(\tau)+\hat{b}(\tau)SP(t), \tag{7}$$

where the parameters $\hat{a}(\tau)$ and $\hat{b}(\tau)$ are typically estimated from an in-sample regression of realized τ-period growth rates onto past spreads. Unlike our two-factor model, these parameters are not endogenously determined, but rather depend on the historically estimated relation between real consumption growth rates and spreads.

III. Empirical Results

In this section, we compare the two-factor model, expression (6),[7] to the spread model, expression (7), in terms of their predictive accuracy in forecasting real consumption growth rates. Implementing the two-factor model requires that we fit the general equilibrium term structure model to prevailing yields and then construct forecasts at various horizons τ. Alternatively, the spread model uses the historically estimated relation between τ period real consumption growth rates and spreads to form τ period-ahead forecasts according to (7).

Our empirical analysis relies on U.S. data drawn from the post-Volcker experiment era, 1985–2002. By doing so, we attempt to ensure that data are not sampled from differing macroeconomic regimes so that we do not erroneously attribute to a forecast error a result due entirely to a change in regimes.[8]

A. Data

The term structure data that we use to estimate the model are monthly observations over the sample period 1985 to 2000 on the annualized zero coupon yields (the average of bid and ask yields) of U.S. Treasuries for six distinct maturities: three months and from one to five years. The three-month data are taken from CRSP's Fama file, while the one- to five-year data are taken from CRSP's Fama-Bliss file.

Our consumption data are monthly observations from 1985 to 2002 on seasonally adjusted real (1996 dollars) personal expenditures on services plus nondurables from the U.S. Department of Commerce's Bureau of Economic Analysis. The corresponding deflator of personal expenditures on services plus nondurables measures the price level used to estimate the expected rate of inflation π. In particular, for each month we use the previous 10 years of monthly observations on the logarithmic change in this deflator to fit an $ARIMA(1,0,1)$ model and take the resultant one step ahead forecast as our estimate of π.

[7] Or, equivalently, expression (4).

[8] See Chapman (1997) for empirical evidence consistent with the choice of 1985 as the break point delineating the beginning of the post-Volcker experiment era. However, to assess the robustness of our results, our later analysis also considers data drawn from the pre-Volcker experiment era.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

B. Term Structure Model Estimation

We cast the estimation of the general equilibrium term structure model in a linear state-space framework.[9] Consistent with the model, the underlying state variables, $x(t)$, are explicitly recognized as unobserved while the observed bond yields are assumed to be a linear function of $x(t)$. While the term structure model is derived in continuous time, its estimation is carried out in discrete time as yield data are only available at discrete-time intervals of length $\Delta \equiv$ one month.

Suppose that at each date t we observe yields of bonds, Y, with M distinct maturity dates $T_1, T_2, \ldots, T_M$ or, equivalently, M distinct terms to maturity, $\tau_1, \tau_2, \ldots, \tau_M$, $Y_t = (Y(t\,;T_1), Y(t\,;T_2), \ldots, Y(t\,;T_M))'$. Each observed yield can be expressed as the corresponding yield given by the model plus an independent, normally distributed measurement error, e_{t,τ_i}. Measurement errors in the observed bond yields reflect noise arising from the bid-ask spread or possible quotation errors. This gives the following set of measurement equations,

$$Y_t = K_0 + Kx_t + e_t, \tag{8}$$

where $K_0 \equiv \begin{pmatrix} \kappa_0(\tau_1\,;\zeta) \\ \cdots \\ \kappa_0(\tau_M\,;\zeta) \end{pmatrix}$, $K \equiv \begin{pmatrix} \kappa(\tau_1\,;\zeta) \\ \cdots \\ \kappa(\tau_M\,;\zeta) \end{pmatrix}$, and $e_t \equiv \begin{pmatrix} e_{t,\tau_1} \\ \cdots \\ e_{t,\tau_M} \end{pmatrix}$.

In addition, the state variables' transition equation in discrete time can be written as

$$x_{t+\Delta} = a(\Delta) + B(\Delta)x_t + \upsilon_t, \tag{9}$$

where the transition errors υ_t are assumed to be independently bivariate normally distributed with the mean equal to the zero vector and covariance matrix given by $F(\Delta)$ from expression (3), which imposes cross-equation restrictions on the variance and covariance properties of the state variables. To complete the specification, the measurement errors e_t and the transition errors υ_t are assumed to be uncorrelated at all lags and to be uncorrelated with the initial state vector.

With these assumptions, we use the Kalman filter to optimally predict the underlying state variables, x_t, as well as to efficiently evaluate the corresponding likelihood function. Numerical optimization of this likelihood function over ζ gives the maximum likelihood estimator $\hat{\zeta}$ of the parameters of the general equilibrium term structure model.[10]

[9] See Duffee (1999) and references therein on estimating term structure models in a state-space framework.

[10] We do not estimate the parameter δ, the rate of patience, because this parameter does not enter the closed-form solution for yields given by expression (5). However, since other studies (e.g., Dunn and Singleton (1986) and Ferson and Constantinides (1991)) find that the intertemporal coefficient $\beta = e^{-\delta}$, a one-to-one transformation of δ, is statistically indistinguishable from one, we estimate the remaining parameters under the equivalent restriction that δ equals zero. Our results do not change qualitatively if we set β to be less than but close to one or, equivalently, if we set δ to be greater than but close to zero.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

Term Structure Model Estimation Results

We recursively estimate the general equilibrium term structure model using a fixed 10-year window of monthly data beginning in February 1985. Using 120 months of yield data from February 1985 to January 1995, we fit the term structure model as of January 1995, obtain the corresponding maximum likelihood parameter estimates, and measure the errors in pricing the sampled Treasury securities through January 1995. Subsequently, moving forward one month, the 120 months of yield data ending in February 1995 allow us to update the maximum likelihood estimates and measure the errors in fitting the term structure through February 1995. Proceeding recursively in this fashion, we obtain maximum likelihood estimates of the term structure model's parameters at monthly intervals from January 1995 to December 2000 as well as corresponding Treasury pricing errors.

Table 1 summarizes the term structure model's maximum likelihood parameter estimates from January 1995 through December 2000.[11] Notice that the estimated mean reversion coefficients, $\hat{\gamma}_1$ and $\hat{\gamma}_2$, are, on average, consistent with the first factor behaving like a random walk, $\hat{\gamma}_1 \approx 0$, while the second factor is more stationary in its behavior, $\hat{\gamma}_2 < 0$. Despite this difference, the corresponding estimated volatility coefficients, $\hat{\sigma}_1$ and $\hat{\sigma}_2$, are quite similar on average.

TABLE 1

The Term Structure Model's Parameter Estimates

Parameter Estimate	$\hat{\gamma}_1$	$\hat{\gamma}_2$	$\hat{\phi}_1$	$\hat{\phi}_2$	$\hat{\sigma}_1$	$\hat{\sigma}_2$	$\hat{\sigma}_c$
Mean	−0.0746	−0.8990	0.0059	0.0140	0.0159	0.0153	0.0031
Std. Dev.	0.0701	0.2000	0.0057	0.0147	0.0041	0.0046	0.0097

Table 1 summarizes the term structure model's parameter estimates obtained by fitting zero coupon Treasury yields. We estimate the model using maximum likelihood by casting it in a discrete-time, state-space framework and evaluating the likelihood function using the Kalman filter. Proceeding recursively, we estimate the term structure model at monthly intervals from February 1995 to December 2000.

Summary statistics for the resultant errors in pricing the sampled zero coupon Treasury yields are provided in Table 2. We define an error here as the fitted yield minus the actual yield measured in basis points. To interpret these statistics, recall that in fitting the term structure model we measure these errors by maturity for each of the preceeding 120 months of sampled yield data. For each estimation date, we can then calculate the resultant mean errors and root mean squared errors. Table 2 provides the average of these errors across all of the estimation dates.[12] Similar to Duffee (1999), the results of Table 2 indicate that the model provides an adequate fit to the term structure although it does not fit the short end as well.

[11] Given the recursive nature of our estimation procedure, these parameter estimates are not independent. Because our focus is on investigating the predictive accuracy of consumption growth forecasts, we do not provide a statistical analysis of these parameter estimates that takes their overlapping nature into account. Any serial dependence in the resultant consumption growth forecasts, however, will be explicitly taken into account.

[12] Once again, given the recursive nature of our estimation procedure, these errors are not independent and the summary statistics are provided for illustrative purposes.

TABLE 2

Summary Statistics of the Term Structure Model's Pricing Errors

Yield Maturity	Avg. of Mean Errors (bps)	Std. Dev. of Mean Errors (bps)	Avg. of Root Mean Squared Errors (bps)
3 months	2.70	6.07	17.83
1 year	−3.68	2.39	15.80
2 years	1.15	2.28	12.41
3 years	2.00	2.52	7.19
4 years	−1.72	2.60	8.09
5 years	−0.45	3.58	12.25

Table 2 summarizes the term structure model's error properties in fitting zero coupon Treasury yields over the sample period February 1985 to December 2000. An error is defined as a fitted yield minus an observed yield and is measured in basis points. Yield data, the average of bid and ask, are obtained from CRSP's Fama file (three-month maturity) and CRSP's Fama-Bliss file (one–five years). We estimate the model using maximum likelihood by casting it in a discrete-time, state-space framework and evaluating the likelihood function using a Kalman filter. Proceeding recursively, we estimate the term structure model at monthly intervals from January 1995 to December 2000 and at each date measure the resultant errors in fitting the preceeding 10 years of monthly yields.

C. Forecasting Results

To forecast real consumption growth using the two-factor model, we use the preceeding 10 years of yield data to estimate the parameters of the general equilibrium model needed to forecast three and six months ahead as well as one through five years ahead. Proceeding recursively in this fashion from January 1995 through December 2000, we compute consumption growth forecasts that we then compare to realized consumption growth rates.

Alternatively, to forecast real consumption growth using the spread model, we use the preceeding 10 years of yield *and* real consumption data to fit linear regressions of realized consumption growth rates against the spread observed between five-year and three-month yields. Linear regressions are separately fit for each of the forecast horizons. We then use the corresponding estimated coefficients to forecast consumption growth rates over that particular horizon. Proceeding recursively from January 1995 through December 2000 gives competing forecasts to those produced using our general equilibrium model.

Table 3 tabulates average differences between the spread model's and the two-factor model's squared errors in forecasting realized consumption growth rates. A positive difference here is consistent with the two-factor model being more accurate.[13] We also present the results of the Diebold and Mariano (1995) test of the null hypothesis of no difference in the accuracy of these competing forecasts.[14]

[13]Qualitatively similar results obtain for the corresponding absolute errors and are not reported here.

[14]Let d_t denote the difference in squared errors between the competing forecasts at t or, in other words, the loss differential at t. Under the null hypothesis that the population mean of the loss differential series $\{d_t\}$ is zero, the statistic $\bar{d}/\sqrt{2\pi\hat{f}_d(0)T}$ is asymptotically standard normal distributed where $\hat{f}_d(0)$ is a consistent estimate of the spectral density of the loss differential at frequency $\omega = 0$. Following Newey and West (1987), a consistent estimate of $2\pi f_d(0)$ is obtained by using a Bartlett kernel with lag selected according to Newey and West's (1994) automatic bandwidth selection procedure. As Diebold (2001) notes, the Diebold-Mariano statistic is simply a t-statistic for the hypothesis of a zero population mean loss differential, adjusted to reflect the fact that the loss differential series is not necessarily white noise. In practice, we compute the Diebold-Mariano statistic by simply regressing the appropriate loss differential series against an intercept and correcting this equation for serial correlation.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

The clear message that emerges from Table 3 is that the two-factor model provides more accurate forecasts of real consumption growth rates at forecasting horizons of one year and longer. At short horizons, three or six months, we cannot reject the null hypothesis that there is no difference in the accuracy of the competing forecasts. At longer horizons, however, we see reliable evidence that the two-factor model's forecasts are more accurate than the spread model's forecasts.[15]

TABLE 3

Comparing Predictive Accuracy in Forecasting Real Consumption Growth Rates: The Two-Factor vs. the Spread Model using the Spread between Five-Year and Three-Month Yields at all Horizons

Time Horizon	Avg. of Diff. in Squared Errors	*p*-Value of Diebold-Mariano Statistic
3 months	−0.06	0.85
6 months	0.19	0.49
1 year	0.59	0.02
2 years	1.21	< 0.01
3 years	1.22	< 0.01
4 years	0.97	< 0.01
5 years	0.65	< 0.01

Table 3 compares the predictive accuracy of the two-factor model vs. the spread model in forecasting real consumption growth rates. The spread model uses the spread between five-year and three-month yields throughout. Predictive accuracy is measured by a model's corresponding squared errors. We tabulate by forecast horizon the average across estimation dates of the differences between the spread model's and the two-factor model's squared errors. The asymptotic *p*-values of the Diebold-Mariano (1995) statistic, which tests the null hypothesis of no difference in the accuracy of these competing forecasts, are also provided.

Figure 1 graphically reinforces these conclusions. Both forecasting models are least accurate at short horizons where real consumption growth rates are extremely noisy. As expected, as the forecast horizon lengthens, the accuracy of both models tends to improve. The improvement in the forecast accuracy of the spread model, however, is less dramatic. For example, at three-, four-, and five-year horizons, the two-factor model's forecasts are within approximately 50 basis points of corresponding actual real consumption growth rates whereas the spread model's are within 100 or more basis points. To understand this, recall that as the forecast horizon lengthens, the spread model relies on more dated information. For example, using the spread model to forecast five-year growth rates implies that the latest observation available to estimate the model is the spread prevailing five years ago together with the subsequent five-year growth rate. The use of anachronous data is not a problem with our two-factor model as we do not rely on consumption data whatsoever; rather we imply this forecast from today's term structure.

To further investigate the two-factor model, Figure 2 graphically displays the time-series properties of the model's forecasting errors, defined as predicted

[15]Term spread models typically rely on the five-year yield to proxy the long rate. Our conclusions do not appreciably change if, alternatively, we compare the forecast accuracy of the two-factor model vs. a spread model based on the spread between one-year and three-month yields. In particular, as before, relying on the Diebold and Mariano test, we cannot reject the null hypothesis of no difference in the accuracy of these competing forecasts at the three- ($p = 0.66$) and six-month ($p = 0.78$) horizons and reject in favor of the two-factor model at the two- ($p = 0.03$), three- ($p < 0.01$), four- ($p < 0.01$), and five-year ($p < 0.01$) horizons. Now, however, we cannot reject the null hypothesis at the 12-month horizon ($p = 0.53$).

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

FIGURE 1

Average Root Mean Squared Errors in Forecasting Consumption Growth Rates

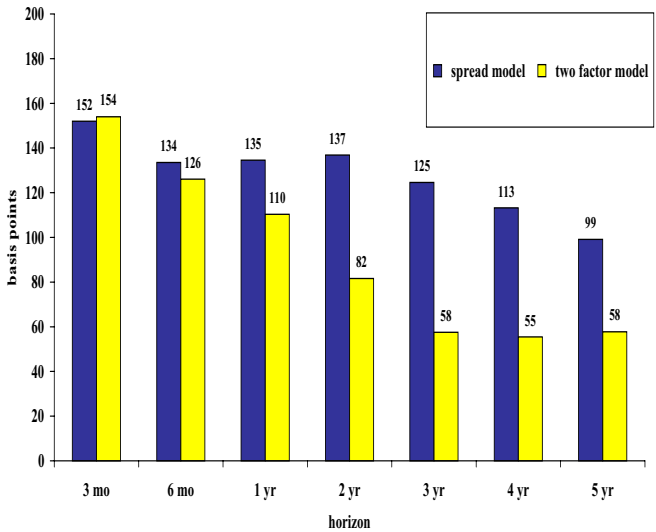


Figure 1 compares the predictive accuracy of the two-factor term structure model with the spread model in forecasting real consumption growth rates. Predictive accuracy is measured in basis points by a model's corresponding root mean squared forecasting error. The two-factor model forecasts are obtained by using the preceeding 10 years of yield data to estimate the parameters needed to forecast subsequent real consumption growth rates. Alternatively, to forecast real consumption growth using the term spread model, we use the preceeding 10 years of yield and real consumption data to fit linear regressions of realized consumption growth rates against the spread observed between five-year and three-month yields. The corresponding estimated coefficients are then used to forecast subsequent consumption growth rates over the different horizons. Proceeding recursively, we compute real consumption growth forecasts for both competing models from January 1995 through December 2000 and compare them to realized consumption growth rates.

real consumption growth rates minus observed growth rates. Graph A of Figure 2 considers three-, six-, and 12-month forecasting horizons, while Graph B considers two-, three-, four-, and five-year horizons. While the forecasts appear to track real consumption growth rates fairly well over the entire sample period, the model appears to systematically under predict real consumption growth during the late 1990s, especially at short horizons (Graph A). One interpretation of this result is that actual consumption growth rates are unexpectedly high here, at least relative to what the term structure of interest rates predicts, because of the significant stock market appreciation surrounding the Internet bubble and the consequent effects of this increase in stock market wealth on consumer spending.[16]

1. Sources of Improved Forecast Accuracy

While the results of Table 3 are consistent with the two-factor model giving more accurate forecasts at most horizons, the results provide no information as to whether this improvement is due to lower bias, lower variance, or both. We can, however, test for the significance of these sources of improvement by following Ashley, Granger, and Schmalensee (1980) whose test procedures are applicable

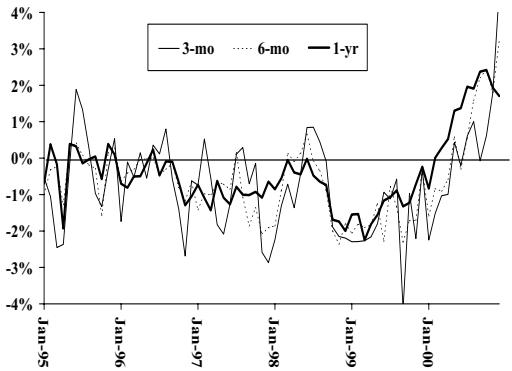
[16] See, for example, Poterba (2000) for an investigation of the effects of stock market wealth on consumption.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

FIGURE 2

The Two-Factor Model's Consumption Forecasting Errors

Graph A. Three-, Six-, and Twelve-Month Forecasting Horizons



Graph B. Two-, Three-, Four-, and Five-Year Forecasting Horizons

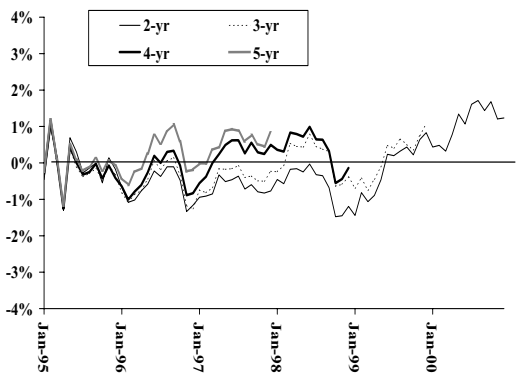


Figure 2 shows the time-series properties of the two-factor model's forecasting errors, defined as predicted minus observed real consumption growth rates. Model forecasts are obtained by using the preceeding 10 years of yield data to estimate the parameters needed to forecast real consumption growth rates. Proceeding recursively, we compute consumption growth forecasts from January 1995 through December 2000 and compare them to realized consumption growth rates.

even if the forecast errors are cross-correlated, autocorrelated, and have non-zero means.

In particular, letting $e_{1,t}$ and $e_{2,t}$ denote the forecast errors of the spread model and the two-factor model, respectively, if Δ_t denotes the difference in these forecast errors, $\Delta_t \equiv e_{1t} - e_{2t}$, and Σ_t denotes their sum, $\Sigma_t \equiv e_{1t} + e_{2t}$, with the corresponding sample mean, $\bar{\Sigma}$, then the regression equation,

$$\Delta_t = \beta_0 + \beta_1 \left(\Sigma_t - \bar{\Sigma}\right) + u_t, \tag{10}$$

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

allows us to test the null hypotheses that there is no difference in the biases of the competing forecasts, $H_0 : \beta_0 = 0$, and that there is no difference in the variances of the competing forecasts, $H_0 : \beta_1 = 0$.[17]

Table 4 tabulates the results. There for each horizon we decompose the difference in mean squared forecasting errors of the spread model vs. the two-factor model into the corresponding differences in squared mean errors and differences in the variances of their forecasting errors together with the statistical significance of these differences.

TABLE 4

Sources of Improved Forecast Accuracy

Time Horizon	Differences in Mean Errors Squared	Differences in Sample Variances	*p*-Value of Statistic Testing No Difference in Biases	*p*-Value of Statistic Testing No Difference in Variances
3 months	−0.15	0.08	0.30	0.74
6 months	0.04	0.15	0.75	0.44
1 year	0.29	0.30	0.02	0.11
2 years	0.59	0.62	< 0.01	< 0.01
3 years	1.11	0.11	< 0.01	0.32
4 years	1.23	−0.26	< 0.01	< 0.01
5 years	0.89	−0.25	< 0.01	< 0.01

Table 4 decomposes the differences in the sample mean squared forecasting errors (MSE) of the spread model with errors $e_{1,t}$ and the two-factor model with errors $e_{2,t}$ into the differences in their sample mean errors squared, m^2, plus the differences in their sample variances, s^2, $\text{MSE}(e_1) - \text{MSE}(e_2) = [m(e_1)^2 - m(e_2)^2] + [s^2(e_1) - s^2(e_2)]$. We also tabulate the *p*-values of Ashley, Granger, and Schmalensee's (1980) tests of the null hypotheses that there is no difference in the biases of the competing forecasts and that there is no difference in the variances of the competing forecasts. The Newey and West (1987) correction for heteroskedasticity and serial correlation is applied to take into account the overlapping nature of these forecasts.

At horizons of one year and longer, the two-factor model's forecasts are significantly less biased than those of the spread model. Furthermore, the two-factor model's forecasts are significantly less variable at the two-year horizon but more variable at horizons of four and five years. In general, the results of Table 4 point to the improvement in the two-factor model's mean squared forecasting errors stemming from their consistently smaller mean forecasting errors.

2. Using the Entire Term Structure of Spreads

Until now, the empirical analysis of the spread model has relied on just two points on the term structure, the five-year and three-month yields, to forecast real consumption growth rates at *all* horizons. By comparison, the two-factor model uses information from all of the sampled yields. Therefore, the improvement in the two-factor model's forecasting accuracy may simply reflect the fact that it makes use of the entire term structure, not just two points on it.

To investigate whether this is the case, we compare the two-factor model to a *system* of spread regressions in which each regression relies on the spread

[17] Since the sample mean squared error, MSE, can be decomposed into the sample variance, s^2, plus the sample mean error squared, m^2, we can write $\text{MSE}(e_1) - \text{MSE}(e_2) = [s^2(e_1) - s^2(e_2)] + [m(e_1)^2 - m(e_2)^2]$, which with further manipulation can be expressed as $\text{MSE}(e_1) - \text{MSE}(e_2) = \widehat{\text{cov}}(\Delta, \Sigma) + [m(e_1)^2 - m(e_2)^2]$ where $\widehat{\text{cov}}$ denotes sample covariance. Defining u_t to be a mean zero error term assumed to be independent of Σ, and $\bar{\Sigma}$ to be the sample mean of Σ, then the regression equation (10) and the corresponding tests immediately follow. Without loss of any generality, the discussion here and in the text assumes that both error means are positive. The methodology can be easily generalized to the other cases. See Ashley, Granger, and Schmalensee (1980) for details.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

calculated over that particular portion of the term structure corresponding to the horizon over which real consumption growth is forecasted. That is, we use the spread between τ-period and one-month yields to forecast the τ-period real consumption growth rate, where τ = three months, six months, one year, two years, three years, four years, and five years.[18] This system of regressions allows the spread model to take advantage of all of the sampled yields not just two. The results are given in Table 5. While the forecast accuracy of the spread model is now improved, the two-factor model still is more accurate at longer horizons.

TABLE 5

Comparing Predictive Accuracy in Forecasting Real Consumption Growth Rates: The Two-Factor Model vs. the Spread Model Using the Spread between τ-Period and One-Month Yields at the τ-Period Horizon

Time Horizon	Avg. of Differences in Squared Errors	*p*-Value of Diebold-Mariano Statistic
3 months	−0.41	0.15
6 months	−0.17	0.37
1 year	0.11	0.60
2 years	0.76	< 0.01
3 years	0.98	< 0.01
4 years	0.89	< 0.01
5 years	0.64	< 0.01

Table 5 compares the predictive accuracy of the two-factor model vs. the spread model in forecasting real consumption growth rates. The spread model now uses the spread between τ-period and one-month yields to forecast τ-period growth rates where τ = three months, six months, one year, two years, three years, four years, and five years. Predictive accuracy is measured by a model's corresponding squared error. We tabulate by forecast horizon the average across estimation dates of the differences between the spread model's and the two-factor model's squared errors. The asymptotic *p*-values of the Diebold-Mariano (1995) statistic testing the null hypothesis of no difference in the accuracy of these competing forecasts are also provided.

3. Comparison to the One-Factor Model

We also fit a simpler one-factor version of our term structure model and use it to forecast real consumption growth rates. Table 6 presents the results and shows that there are consequences to relying on the simpler model. In particular, at every horizon except three years, the two-factor model provides more accurate forecasts of real consumption growth. This improvement in forecast accuracy is statistically significant at horizons of less than one year as well as at the five-year horizon.[19]

4. Predictive Accuracy in Forecasting Real Consumption Growth Rates from January 1969 to December 1973

Thus far, we restricted our attention to forecasting consumption growth only in the post-Volcker experiment era. We now investigate how well the two-factor model performs in earlier data.

Monthly real consumption data is available from January 1959. To make these forecasting results comparable to our previous analysis, we implement the two-factor and spread models as before. In the case of the spread model, we use 10

[18] Like the three-month Treasury yield, the one-month Treasury yield is taken from CRSP's Fama file. The six-month yield is taken from the FRED database of the Federal Reserve Bank of St. Louis.

[19] In unreported results, we also explored a more computationally intensive three-factor model but it does not provide statistically more accurate forecasts than the two-factor model at any horizon.

TABLE 6

Comparing Predictive Accuracy in Forecasting Real Consumption Growth Rates: The Two-Factor vs. the One-Factor Model

Time Horizon	Avg. of Diff. in Squared Errors	*p*-Value of Diebold-Mariano Statistic
3 months	0.46	0.09
6 months	0.43	0.05
1 year	0.21	0.12
2 years	0.01	0.79
3 years	−0.05	0.28
4 years	0.18	0.15
5 years	0.46	< 0.01

Table 6 compares the predictive accuracy of the two-factor vs. the one-factor model in forecasting real consumption growth rates. Predictive accuracy is measured by a model's corresponding squared error. We tabulate by forecast horizon the average across estimation dates of the differences between the one-factor and the two-factor models' squared errors. The asymptotic *p*-values of the Diebold-Mariano (1995) statistic testing the null hypothesis of no difference in the accuracy of these competing forecasts are also provided.

years of yield and real consumption data, beginning in January 1959, to fit linear regressions of realized consumption growth rates against the spread between five-year and three-month yields to recursively forecast out-of-sample consumption growth rates from January 1969 to December 1973.[20] Alternatively, we use 10 years of yield data beginning in January 1959 to fit the two-factor model and then proceed recursively from January 1969 to December 1973 to imply out-of-sample real consumption growth forecasts.

From Table 7, we see that the accuracy in forecasting out-of-sample consumption growth rates, both absolutely, as measured by the magnitudes of the root mean squared forecast errors, and relatively, as compared to the term spread model, deteriorates over the earlier January 1969 to December 1973 forecast period. With respect to the two-factor model's absolute accuracy, its root mean squared errors tabulated in Table 7 are at least twice as large as the corresponding errors over the December 1995 to January 2002 period shown in Figure 1. From Table 7, we also see that the Diebold and Mariano test no longer rejects the null hypothesis of no difference in the accuracy of these competing forecasts at any horizon in the January 1969 to December 1973 forecast period.

The results suggest a deterioration in the ability of the term structure to forecast real consumption growth during the pre-Volcker experiment period. Reasons for this deterioration include the fact that consumption growth was more volatile and so more difficult to forecast accurately during this period. For example, an F-test rejects the null hypothesis that the variance of monthly consumption growth rates over the January 1969 to December 1973 period, 4.97%, equals the corresponding variance over the January 1995 to December 2000 period, 2.86%, with a p-value < 0.01 ($F = 2.95$). Additionally, as Cochrane (2001), p. 380 argues, "The data seem to suggest a change in regime between the 1970s and 1990s: in the 1970s most interest rate changes were due to inflation, while the opposite seems true now." To the extent that nominal interest rates indeed changed predominantly because of real rate news in our later forecast period as opposed to inflation news in the earlier forecast period is consistent with the term structure being better able

[20] By ending at December 1973, we ensure that our forecasts do not extend beyond December 1978, which Chapman (1997) delineates as the end of the pre-Volcker experiment era.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

TABLE 7

Predictive Accuracy of the Two-Factor Model in Forecasting Real Consumption Growth Rates (January 1969 to December 1973)

Time Horizon	Root Mean Squared Error (bps)	*p*-Value of Diebold-Mariano Statistic
3 months	310	0.13
6 months	271	0.12
1 year	234	0.17
2 years	203	0.68
3 years	150	0.85
4 years	130	0.85
5 years	124	0.38

Table 7 summarizes the predictive accuracy of the two-factor model in forecasting real consumption growth rates over the January 1969 to December 1973 forecast period and also compares its predictive accuracy to that of the spread model. The two-factor model's root mean squared errors, in basis points (bps), are tabulated by forecast horizon. The asymptotic *p*-values of the Diebold-Mariano (1995) statistic testing the null hypothesis of no difference in the accuracy of the two-factor and spread model forecasts are also provided.

to forecast real consumption growth over the January 1995 to December 2000 period.

5. Forecasting Yields vs. Forecasting Consumption

Recently Duarte (2004) and Duffee (2002) conclude that term structure models similar to ours produce relatively poor forecasts of future Treasury yields. For example, relying on monthly data from January 1952 until December 1994, Duffee demonstrates that a variety of term structure models produce out-of-sample forecasts no better than a simple random walk model of yields over the January 1995 through December 1998 forecasting period.

Duffee's conclusion that term structure models do a poor job at forecasting yields appears to be inconsistent with our evidence regarding the ability of the two-factor model to forecast consumption growth rates over the January 1995 to December 2000 period that includes Duffee's forecast period. In both cases, however, the conclusions pertain to how well the models perform relative to a benchmark. Unfortunately, a direct comparison here is difficult because the benchmarks differ in each case—a random walk model in forecasting yields vs. a term spread model in forecasting consumption growth rates.

We can, however, investigate the extent to which the absolute accuracy of term structure models in forecasting yields vs. consumption growth rates is dependent on the choice of forecasting period. Since neither Duarte nor Duffee provides out-of-sample forecasting results for earlier forecasting periods, we rely on the two-factor model to also forecast three-month, two-, and five-year yields at forecast horizons of three- and six-months over our January 1969 to December 1973 forecast period. For comparison purposes, we also produce these yield forecasts over the January 1995 to December 2000 forecast period. The results are tabulated in Table 8, Panel A, which gives the results for the later forecast period and Panel B, which gives results for the earlier period.

Even though the bond maturities differ slightly as do the forecast horizons, [21] the results in Panel A are consistent with those of Duffee. In particular, the root

[21] Duffee considers six-month, two-, and ten-year maturities and forecast horizons of three, six, and 12 months.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

TABLE 8

Predictive Accuracy of the Two-Factor Model in Forecasting Treasury Yields

Bond Maturity	Forecast Horizon	Root Mean Squared Error (bps)
Panel A. January 1995 to December 2000 Forecast Period		
3 months	3 months	40
2 years	3 months	55
5 years	3 months	59
3 months	6 months	72
2 years	6 months	83
5 years	6 months	86
Panel B. January 1969 to December 1973 Forecast Period		
3 months	3 months	102
2 years	3 months	81
5 years	3 months	65
3 months	6 months	149
2 years	6 months	113
5 years	6 months	90

Table 8 summarizes the predictive accuracy of the two-factor model in forecasting Treasury yields from January 1995 to December 2000 (Panel A) and January 1969 to December 1973 (Panel B). The two-factor model's root mean squared errors, in basis points (bps), are provided for bond maturities of three months, two and five years, and for forecast horizons of three and six months.

mean squared forecasting errors are similar in magnitude and increase with bond maturity and forecast horizon. By contrast in Panel B, similar to forecasting consumption growth rates, the two-factor term structure model's accuracy in forecasting yields deteriorates over the January 1969 to December 1973 period. For example, for each bond maturity and each forecast horizon, the magnitude of the corresponding root mean squared forecast error is larger. In addition, the model's forecast accuracy is now worse at the three-month bond maturity and improves with increasing bond maturity. The greater variability of yields over this forecasting period clearly adversely affects the term structure model's forecasting ability.

So while it may be difficult to compare the relative accuracy of term structure models in forecasting yields vs. consumption growth rates, their respective absolute accuracies behave similarly over the forecasting periods that we examine. In particular, the root mean squared out-of-sample forecasting errors decrease over the January 1995 to December 2000 period as compared to the January 1969 to December 1973 period.

IV. Conclusions

Investors set the yields of bonds maturing at different times in the future by taking into account the levels of consumption expected at those times. In this paper, we recover these investor expectations from a simple general equilibrium model of the term structure. By fitting this model to observed yields, we are able to imply ex ante forecasts of real consumption growth without relying whatsoever on consumption data. This is in contrast to the spread model whose forecasts are obtained by extrapolating the ex post relation between the term spread and subsequent realized consumption growth. By its very construction, the longer the forecast horizon, the more dated is this relation and so the less accurate are the spread model's forecasts. Relying on term structure data over the 1985–2000 sample period, our empirical results are indeed consistent with the increased predictive accuracy of the general equilibrium approach.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press

References

Ang, A.; M. Piazzesi; and M. Wei. "What Does the Yield Curve Tell Us about GDP Growth?" Unpubl. Working Paper, UCLA (2002).

Ashley, R.; C. W. J. Granger; and R. Schmalensee. "Advertising and Aggregate Consumption: An Analysis of Causality." *Econometrica*, 48 (1980), 1149–1167.

Brown, R. H., and S. M. Schaefer. "Interest Rate Volatility and the Shape of the Term Structure." *Philosophical Transactions of the Royal Society A*, 347 (1994), 563–576.

Chapman, D. A. "The Cyclical Properties of Consumption Growth and the Real Term Structure." *Journal of Monetary Economics*, 39 (1997), 145–172.

Cochrane, J. H. *Asset Pricing.* Princeton, NJ: Princeton Univ. Press (2001).

Cox, J. C.; J. E. Ingersoll; and S. A. Ross. "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica*, 53 (1985), 363–384.

Dai, Q., and K. J. Singleton. "Specification Analysis of Affine Term Structure Models." *Journal of Finance*, 55 (2000), 1943–1978.

Diebold, F. X. *Elements of Forecasting.* 2nd ed. Cincinnati, OH: South-Western Publishing (2001).

Diebold, F. X., and R. S. Mariano. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics*, 13 (1995), 253–263.

Duarte, J. "Evaluating an Alternative Risk Preference in Affine Term Structure Models." *Review of Financial Studies*, 17 (2004), 379–404.

Duffee, G. R. "Estimating the Price of Default Risk." *Review of Financial Studies*, 12 (1999), 197–226.

Duffee, G. R. "Term Premia and Interest Rate Forecasts in Affine Models." *Journal of Finance*, 57 (2002), 405–443.

Duffie, D. *Dynamic Asset Pricing Theory.* 3rd ed. Princeton, NJ: Princeton Univ. Press (2001).

Dunn, K. B., and K. J. Singleton."Modeling the Term Structure of Interest Rates under Non-Separable Utility and Durable Goods." *Journal of Financial Economics*, 17 (1986), 27–55.

Estrella, A., and G. A. Hardouvelis. "The Term Structure as a Predictor of Real Economic Activity." *Journal of Finance*, 46 (1991), 555–576.

Ferson, W. E., and G. M. Constantinides. "Habit Persistence and Durability in Aggregate Consumption: Empirical Tests." *Journal of Financial Economics*, 29 (1991), 199–240.

Hamilton, J. D., and D. H. Kim."A Re-Examination of the Predictability of Economic Activity Using the Yield Spread." *Journal of Money, Credit and Banking*, 34 (2002), 340–360.

Harvey, C. R. "The Real Term Structure and Consumption Growth." *Journal of Financial Economics*, 22 (1988), 305–334.

Harvey, C. R. "Forecasts of Economic Growth from the Bond and Stock Markets." *Financial Analysts Journal*, 45 (1989), 38–45.

Harvey, C. R. "The Term Structure and World Economic Growth." *Journal of Fixed Income*, 1 (1991), 4–17.

Harvey, C. R. "Term Structure Forecasts Economic Growth." *Financial Analysts Journal*, 49 (1993), 6–8.

Kamara, A. "The Relation between Default-Free Interest Rates and Expected Economic Growth is Stronger Than You Think." *Journal of Finance*, 52 (1997), 1681–1694.

Kessel, R. "The Cyclical Behavior of the Term Structure of Interest Rates." Occasional Paper No. 91, Cambridge, MA: National Bureau of Economic Research (1965).

Langetieg, T. C. "A Multivariate Model of the Term Structure." *Journal of Finance*, 35 (1980), 71–97.

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

Newey, W., and K. West. "A Simple Positive Semi-Definite, Heteroskedastic and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), 703–708.

Newey, W., and K. West. "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies*, 61 (1994), 631–653.

Plosser, C. I., and K. G. Rouwenhorst. "International Term Structures and Real Economic Growth." *Journal of Monetary Economics*, 33 (1994), 133–155.

Poterba, J. "Stock Market Wealth and Consumption." *Journal of Economic Perspectives*, 14 (2000), 99–118.

Roma, A., and W. Torous. "The Cyclical Behavior of Interest Rates." *Journal of Finance*, 52 (1997), 1519–1542.

https://doi.org/10.1017/S0022109000002295 Published online by Cambridge University Press