

w o r k i n g
p a p e r

03 08

**Pricing Kernels, Inflation, and
the Term Structure of Interest Rates**

by Ben R. Craig and Joseph G. Haubrich



FEDERAL RESERVE BANK OF CLEVELAND

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

Working papers are now available electronically through the Cleveland Fed's site on the World Wide Web: **www.clevelandfed.org**.

**Pricing Kernels, Inflation,
and the Term Structure of Interest Rates**

by Ben R. Craig and Joseph G. Haubrich

We estimate a discrete-time multivariate pricing kernel for the term structure of interest rates, using both yields and inflation rates. This gives a separate estimate of the real kernel and the nominal kernel, taking into account a relatively sophisticated dynamical structure and mutual interaction between the real and nominal side of the economy. Along with obtaining an estimate of the real term structure, we use the estimates to obtain a new perspective on how real and nominal influences interact to produce the observed term structure.

Keywords: Inflation, term structure, asset pricing

Ben R. Craig and Joseph G. Haubrich are at the Federal Reserve Bank of Cleveland and can be reached at ben.r.craig@clev.frb.org and joseph.g.haubrich@clev.frb.org. The authors thank George Pennacchi and Peter Ritchken for many interesting discussions on the term structure and Mike Dueker for help with the econometrics.

1. Introduction

Despite its importance for pricing, hedging, and the conduct of monetary policy, the relation between real and nominal interest rates remains somewhat elusive. The central difficulty is that until recently, real rates are not observed. Both the nominal term structure and inflation rates are observed, however, and these can act as inputs to a pricing kernel model that can help sort out the real and nominal elements. Our paper thus unites two somewhat separate lines of research: models based on pricing kernels and models that explicitly account for inflation.

The pricing kernel, or stochastic discount factor, is the stochastic process governing the prices of state contingent claims. Given a pricing kernel, we can compute the price of any financial asset. Given asset prices free from arbitrage opportunities, a pricing kernel exists.¹

The pricing kernel approach unifies many aspects of asset pricing theory. Some work only implicitly defines pricing kernels, deriving them from utility functions and returns technology (Cox, Ingersoll and Ross, 1985, Pennacchi, 1991) or from consumption patterns (Campbell, 1986). Starting directly with the pricing kernel, however, has advantages of flexibility, generality, and tractability, as Constantinides (1992) and Backus and Zin (1994) have shown to great effect. Given the problems with correctly specifying the return process in CIR type models (Constantinides 1992) and the failures of consumption based models (Hansen and Jagannathan, 1991), the pricing kernel approach offers a different avenue of progress.

Most work using the pricing kernel, however, has assumed a “nominal” pricing kernel,

¹ Sargent (1987) refers to it as the pricing kernel. Harrison and Kreps (1979) and Kreps (1981) use the concept to characterize price systems without arbitrage opportunities, where a variant appears as the martingale equivalent measure. See Huang and Litzenberger (1988), chapter 8, and Duffie (1992), chapters 2-4.

which confounds real interest rates and inflation. This approach is inadequate for questions of extracting real and nominal information, be it from the term structure, inflation indexed bonds, or CPI futures.

This paper explicitly adds inflation into the pricing kernel approach. It splits the nominal kernel into a real kernel and an inflation process, assuming neither independence nor a complete representative agent model. In this case, the kernel approach is particularly appealing. It allows interaction between the real and nominal sides without having to reach a consensus on the proper fully specified dynamic general equilibrium model with a role for money.²

This paper complements two other papers that specifically address the relation between real and nominal rates. Pennacchi (1991) directly models the real rate and the inflation rate in a Cox-Ingersoll-Ross framework and uses survey measures of expected inflation. We use a discrete-time kernel framework in conjunction with the actual inflation process, and estimate a richer dynamic structure on both the real and inflation components of interest rates. We also look at the term structure out to 120 months, whereas Pennacchi stops at 12. Evans (1998) uses UK index-linked gilts to look at real rates, expected inflation, and inflation risk premia. As mentioned before, we take as our two observables the actual inflation rate and nominal bond prices, in part because we have concerns about the comparability of US inflation-indexed securities, mainly because of liquidity and tax differences. We also impose more structure on the underlying kernel in the estimation procedure, a process useful for general hedging and monetary policy considerations.

This paper most closely follows Backus and Zin (1994) in adopting the log-linear

² For interesting work on the term structure in general equilibrium models with money, see Boudoukh (1993), Labadie (1994), den Haan (1995), Buraschi and Jiltsov (1999), T. Wu (2000), and S. Wu (2001).

discrete time approach to pricing bonds. This allows for complicated dynamics moving substantially beyond mean reversion. The added flexibility gives the model a better chance of matching the time series and cross-sectional properties of the interest rate data. On one level, this paper may be viewed as an elaboration of the Backus and Zin paper, formulating and estimating a multifactor specification of the pricing kernel. Without downplaying such a contribution, we prefer a different emphasis. Using the pricing kernel to combine inflation and interest rates provides a more detailed analysis of the interaction between real and nominal variables.

2. A Theoretical Framework for Pricing Kernels and Inflation

The basic relationship in asset pricing, relating the expected real return on an asset to the stochastic discount rate (or pricing kernel) takes the form

$$1 = E_t[M_{t+1}R_{t+1}^{n+1}] \quad (1)$$

where R_{t+1}^{n+1} is the one-period gross real return and M_{t+1} is the real pricing kernel or stochastic discount factor (See chapter 11 of Campbell, Lo, and MacKinlay 1997). Such a kernel exists under fairly weak assumptions, requiring little more than that bond traders prefer more consumption to less, and act on that preference to eliminate arbitrage. With more structure, we can split the kernel into real and nominal parts and estimate both from the above pricing relation.

For nominal claims, let $b_t^{\$n}$ be the dollar price at date t of an n -period discount bond: a claim to \$1 at date $t + n$. Then, with the nominal price index at date t denoted Q_t , the real price of a bond is

$$b_t^n = \frac{b_t^{\$n}}{Q_t}.$$

Applying the asset pricing condition to the real return on an n -period nominal bond results

in

$$1 = E_t[M_{t+1} \frac{b_{t+1}^{\$n-1}/Q_{t+1}}{b_t^{\$n}/Q_t}].$$

This implies

$$b_t^{\$n} = E_t[M_{t+1} b_{t+1}^{\$n-1} \frac{Q_t}{Q_{t+1}}].$$

For $n = 0$, $b_t^{\$0} = 1$, expressing the truism that \$1 today costs \$1. The price of \$1 tomorrow is

$$\begin{aligned} b_t^{\$1} &= E_t[M_{t+1} 1 \frac{Q_t}{Q_{t+1}}] \\ &= E_t[M_{t+1} \frac{1}{\Pi_{t+1}}] \end{aligned}$$

where Π_{t+1} is the *gross* inflation rate Q_{t+1}/Q_t . Note that this splits what might be termed the “nominal pricing kernel” into two parts: the real kernel and inflation. In contrast to some earlier work, we do not assume independence between inflation and the real side (Gibbons and Ramaswamy 1993), nor do we directly posit primitive assumptions about the nominal kernel M_{t+1}/Π_{t+1} (Brown and Dybvig, 1986, Pennacchi, Ritchken, and Sankarasubramiam, 1995).

The pricing kernel approach, or reverse engineering approach (to follow Backus and Zin) then proceeds as follows. It first specifies a process for M_t and Π_t and then uses that process to price the term structure, i.e., derive the yield on zero-coupon bonds of different maturities³. The time series and cross section properties implied by the theory are then matched with the data to derive the parameters of the underlying process—to tie down the pricing kernel and inflation. Once tied down, the two halves of the pricing kernel can function as a metric for assessing asset pricing theories and as an engine for pricing securities.

³ This contrasts with the consumption based view, in which the asset pricing equation takes the form $1 = E_t[\beta u'(c_{t+1})/u'(c_t)R_{t+1}^1]$ for time separable utility: the stochastic process for consumption and the form of the utility function determine the pricing kernel.

Kernel Structure

The general asset pricing condition (1) becomes a theory of bond pricing once we characterize the pricing kernel M_t and the inflation rate Π_t . A rich yet tractable choice is a log-linear bivariate process for the two parts of the kernel. Specifically, let the log of M_t follow an ARMA(1,1) in each shock, and let the log of Π_t follow an AR(2) in each shock. We choose the AR(2) process for inflation because classic Box-Jenkins analysis suggests it as a likely representation.

$$\begin{bmatrix} 1 - \phi L & 0 \\ 0 & 1 - \rho_1 L - \rho_2 L^2 \end{bmatrix} \begin{bmatrix} -\log M_t - \delta \\ -\log \Pi_t - \lambda \end{bmatrix} = \begin{bmatrix} 1 + \theta_1 L & 1 + \theta_2 L \\ v_\varepsilon & v_\eta \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \eta_t \end{bmatrix}. \quad (2)$$

In this equation, L denotes the lag operator and v_ε and v_η are constants that allow the shocks to have different influences on M_t and Π_t . The log formulation is useful because yields and interest rates are easier to work with than bond prices, and it allows us to exploit the well-known property of lognormal distributions, that is, if $X \sim N(\mu, \sigma^2)$ and if $\log Y = X$, then $\log E(Y) = \mu + \frac{\sigma^2}{2}$.

In this form, M_t and Π_t are not independent, but they do not depend directly on one another. One interpretation of (2) has ε_t as the real shock and η_t as the nominal shock, but this is not the only way to interpret the representation. Equally important, the process is not univariate. Although for either inflation or the real kernel, a univariate process could deliver the same autocorrelations, such a process would miss the interactions between the two. Even with a reduced form for the kernel, autocorrelations do not suffice to price the term structure. Different parameters delivering the same autocorrelations can result in different term structures (Campbell, Lo, and MacKinlay, 1997 chapter 11.2.1).

If only the autocorrelation properties are considered, the whole nominal kernel can be represented as an ARMA(4,4) process.⁴ As such, it allows richer dynamics than the

⁴ This uses the Granger and Morris (1976) formula, $\text{ARMA}(p, m) + \text{ARMA}(q, n) = \text{ARMA}[(p + q), \max(p + n, q + m)]$.

ARMA(2,3) kernel used by Backus and Zin.⁵ To reiterate, though, splitting out inflation involves more than merely using a more complicated ARMA process. It requires us to explicitly consider the interactions between real and nominal rates.

Equation (2) is rather difficult to work with directly. Of more use is the infinite-order moving average form. Equation (2) reduces to

$$\begin{aligned} -\log M_t &= \delta + \frac{1}{1 - \phi L} (1 + \theta_1 L) \varepsilon_t + \frac{1}{1 - \phi L} (1 + \theta_2 L) \eta_t \\ -\log \Pi_t &= \lambda + \frac{v_\varepsilon}{1 - \rho_1 L - \rho_2 L^2} \varepsilon_t + \frac{v_\eta}{1 - \rho_1 L - \rho_2 L^2} \eta_t. \end{aligned}$$

Putting this in MA(∞) form yields:

$$\begin{aligned} -\log M_t &= \delta + [1 + (\theta_1 + \phi)L + \phi(\theta_1 + \phi)L^2 \dots] \varepsilon_t + [1 + (\theta_2 + \phi)L + \phi(\theta_2 + \phi)L^2 \dots] \eta_t \\ -\log \Pi_t &= \lambda + \sum_{j=0}^{\infty} \psi_j v_\varepsilon \varepsilon_{t-j} + \sum_{j=0}^{\infty} \psi_j v_\eta \eta_{t-j}. \end{aligned}$$

Here ψ_j is defined recursively from the AR(2) process (as in the Yule-Walker equations).

Combining the two, as we must for asset pricing, recalling that $\log(M_t/\Pi_t) = \log M_t - \log \Pi_t$, yields

$$\begin{aligned} -\log M_t + \log \Pi_t &= (\delta - \lambda) + [(1 - v_\varepsilon) \varepsilon_t + \sum_{j=0}^{\infty} [(\theta_1 + \phi) \phi^j - \psi_{j+1} v_\varepsilon] \varepsilon_{t-j-1} \\ &\quad + [(1 - v_\eta) \eta_t + \sum_{j=0}^{\infty} [(\theta_2 + \phi) \phi^j - \psi_{j+1} v_\eta] \eta_{t-j-1}]. \end{aligned} \tag{3}$$

This, in conjunction with equation (1), prices assets. It describes how the pricing kernel evolves over time, or equivalently, how the rate at which we discount the future depends on both real and nominal shocks.

Pricing the Term Structure

⁵ The richer dynamics do not constitute a full generalization because our process allows only some subsets of the ARMA(4,4) process.

The framework for pricing is now in place. The next step in obtaining predictions about the term structure involves describing the assets to be priced. Recall that $b_t^{\$n}$ denotes the dollar price at date t of an n -period discount bond. The dollar *yield* on such a bond is

$$y_t^n \equiv -(1/n) \log b_t^{\$n}.$$

The short rate, or yield on a one-period bond, is

$$r_t \equiv y_t^1 \equiv -\log b_t^{\$1}. \quad (A)$$

Forward rates, also called implied future rates, are defined from the prices of longer term bonds:

$$f_t^n \equiv \log(b_t^{\$n}/b_t^{\$(n+1)}). \quad (B)$$

Definitions (A) and (B) imply that yields are averages of forward rates.

$$y_t^n = \frac{1}{n} \sum_{j=1}^n f_t^{j-1}.$$

This does not imply that forward rates are unbiased or accurate predictors of future short rates, as nothing in the model implies that forward rates are unbiased or accurate predictors of future short rates.

Bond prices and yields then follow using the techniques of Backus and Zin (1995, section 3 and appendix B). They show how to combine the log linearity of the stochastic process with the definition of the yield as a log to simplify the basic asset pricing equation into an expression involving only means and variances. It is then straightforward, if tedious, to compute the forward rates,

$$f_t^n = (\delta - \lambda) - (1/2)(A_{1,n}^2 \sigma_\varepsilon^2 + A_{2,n}^2 \sigma_\eta^2) + \sum_{j=0}^{\infty} \alpha_{1,n+1+j} \varepsilon_{t-j} + \sum_{j=0}^{\infty} \alpha_{2,n+1+j} \eta_{t-j} \quad (4)$$

the yields,

$$\begin{aligned}
y_t^n = & (\delta - \lambda) - \frac{\sigma_\varepsilon^2}{2n} \sum_{j=1}^n A_{1,j-1}^2 - \frac{\sigma_\eta^2}{2n} \sum_{j=1}^n A_{2,j-1}^2 \\
& + \frac{1}{n} \sum_{j=0}^{\infty} (A_{1,n+j} - A_{1,j}) \varepsilon_{t-j} + \frac{1}{n} \sum_{j=0}^{\infty} (A_{2,n+j} - A_{2,j}) \eta_{t-j}
\end{aligned} \tag{5}$$

and the short rate,

$$\begin{aligned}
r_t = & (\delta - \lambda) - (1 - v_\varepsilon) \frac{\sigma_\varepsilon^2}{2} - (1 - v_\eta) \frac{\sigma_\eta^2}{2} \\
& + \sum_{j=0}^{\infty} [(\theta_1 + \phi) \phi^j - \psi_{j+1} v_\varepsilon] \varepsilon_{t-j} + \sum_{j=0}^{\infty} [(\theta_2 + \phi) \phi^j - \psi_{j+1} v_\eta] \eta_{t-j}.
\end{aligned} \tag{6}$$

In these equations, $\alpha_{1,j}$ denotes the coefficient on ε_{t-j} in the nominal kernel (eq. 3) so that $\alpha_{1,0} = (1 - v_\varepsilon)$, $\alpha_{1,1} = (\theta_1 + \phi) - \rho v_\varepsilon$, and so forth. Similarly, $\alpha_{2,j}$ denotes the corresponding coefficient on η_{t-j} , so that $\alpha_{2,2} = (\theta_2 + \phi) \phi - \psi_2 v_\eta$. We let $A_{1,n} = \sum_{j=0}^n \alpha_{1,j}$. In moving from (5) to (6), since the short rate r_t is just y_t^1 , we use $A_{1,j+1} - A_{1,j} = \alpha_{1,j+1}$.

In equation (6), even though inflation and real rates can be additively decoupled, the Fisher equation does not hold. Because $E(\log \pi) \neq \log E(\pi)$, simply subtracting off expected inflation does not give the real interest rate. Mathematically, in the lognormal framework, that means variance terms, σ_ε^2 and σ_η^2 , show up in addition to expected values. Economically, it means that inflation risk matters. Equation (6) also shows that any shock moving inflation will almost surely move the real rates. It would be wrong to take a single piece of information, say that inflation is high this month, and assume that it indicates a decrease in real rates, as the shocks increasing inflation also move the real rate. Naive decompositions are misleading.

3. Estimating the Kernel

Splitting the kernel into real and inflation components means that information about inflation can help estimate the kernel. Following what Backus and Zin (1994) call the

“reverse-engineering” approach, this takes place in two steps. The first step uses equations (4-6) to obtain theoretical moment predictions for

- (i) Time-series (autocorrelation) properties of the short rate
- (ii) Mean values of the yield spreads (average slope of the term structure)
- (iii) Time-series properties of inflation.

The next step chooses parameters to give the best fit with the data. Specifically, that means using GMM (generalized method of moments) to choose θ_1 , ρ_1 , and so forth to match the theoretical moments as closely as possible to the sample moments in the data. Other methods, such as maximum likelihood, would work as well.

In principle, there is no reason to restrict assets and moments to discount bonds. The same techniques could use data on swaps, options, or structured notes, but fitting the term structure is challenge enough for one paper.

Moment Predictions

Starting from equations for yields (eq. 4-6) the model’s prediction for mean spreads and autocorrelations follows in a straightforward fashion. The moments naturally cluster into four areas: autocovariances of the short rate, mean values of yield spreads, autocovariances of inflation, and correlations between inflation, the short rate and yield spreads.

For the autocovariances of the short rate, collecting terms depending on k , the lag, results in

$$Cov(r_t r_{t-k}) = \sigma_\varepsilon^2 \left[\sum_{j=1}^{\infty} \alpha_{1,j} \alpha_{1,j+k} \right] + \sigma_\eta^2 \left[\sum_{j=1}^{\infty} \alpha_{2,j} \alpha_{2,j+k} \right] \quad (7)$$

For the mean yield spreads,

$$E(y_t^n - y_t^1) = \frac{\sigma_\varepsilon^2}{2} (1 - v_\varepsilon) - \frac{\sigma_\varepsilon^2}{2n} \sum_{j=1}^n A_{1,j-1}^2 + \frac{\sigma_\eta^2}{2} (1 - v_\eta) - \frac{\sigma_\eta^2}{2n} \sum_{j=1}^n A_{2,j-1}^2. \quad (8)$$

In evaluating (8) note that

$$\begin{aligned} \sum_{j=1}^n A_{1,j-1}^2 &= [1 - v_\varepsilon]^2 \\ &+ [1 - v_\varepsilon + (\theta_1 + \phi) - v_\varepsilon \psi_1]^2 \\ &+ [1 + (\theta_1 + \phi) + (\theta_1 + \phi)\phi - v_\varepsilon(1 + \psi_1 + \psi_2)]^2 \\ &+ \dots \end{aligned}$$

These are the squared coefficients of v_ε in the Wald representation of the kernel (3). For inflation, note that

$$E(-\log \Pi_t) = \lambda.$$

$$Var(-\log \Pi) = \frac{\rho_1}{1 - \rho_2} (v_\varepsilon^2 \sigma_\varepsilon^2 + v_\eta^2 \sigma_\eta^2). \quad (9)$$

$$E(\log \Pi_t \log \Pi_{t-1}) / Var(-\log \Pi_t) = \frac{1 - \rho_2}{(1 + \rho_2)[(1 - \rho_2)^2 - \rho_1^2]}.$$

Data and Estimation Procedures

The estimation procedure chooses the parameters of the bivariate ARMA process for the real kernel and inflation to make the theory “close” to the data in a precise sense. First, we need to impose one identifying restriction so that GMM can fit the moment conditions provided by the cross-section and time-series data for interest rates and inflation. Implementing GMM requires choosing the exact moments to match, that is, a set of autocovariances and yields spreads. It also requires choosing a data set.

The estimation applies GMM to the US term structure data set of McCulloch and Kwon (1993).⁶ Discarding data prior to the Treasury-Fed Accord, the estimation uses monthly data from January 1952 though February 1991. The parameters are estimated separately on two subsamples, January 1952 through October 1982 and November 1982

⁶ The program was written in OX, and run on a Dell Pentium 4.

through February 1991. Inflation is the monthly percent change in the CPI over the same period.

We estimate the nine parameters of our normalized system, $(\phi, \rho_1, \rho_2, \theta_1, \theta_2, \sigma_\varepsilon, \sigma_\eta, \nu_\varepsilon, \text{ and } \nu_\eta)$ using a total of 19 moment conditions. These moments are the variance and autocovariances of the short rate at 1, 3, 12, and 24 months, the spread between the short rate and yields at 3, 12, 60, and 120 months, the variance and the autocovariance of inflation at 1, 3, 12, and 24 months, and the cross-correlations of the yields with the inflation rate (short rate and 3, 13, 60 and 120 months). The widely separated autocovariances and yield spreads maximize the amount of information derived from the estimation procedure because interest rates show high correlation both among different maturities and with past rates. The five conditions each for the short rate autocovariances and the mean yield spreads (the spread between the long rate and the one-month rate) follow Backus and Zin (1994).

The GMM estimation uses a constant weighting matrix, Ω . The yield-curve moment conditions are represented as simple deviations whereas the autocovariance conditions and the cross-covariances are represented as deviations of squared interest rates. Because of this, choosing Ω to be the identity matrix heavily weights the yield-curve moments. Indeed, the magnitudes prove so dissimilar that estimates produced in this way clearly fit the yield curve well at the expense of the other moments. Yet Monte-Carlo studies (Hayashi 2000, chapter 3) suggest that using a constant Ω (or a single iteration to produce an $\hat{\Omega}$ based on the inverted sum of the squared deviations from the moment conditions), often produces estimates with better small sample properties than those produced by the multiple-iteration technique proposed by Hansen, Heaton, and Ogaki (1989).

Our solution was to choose Ω initially to be a diagonal matrix with the inverse of

average squared sample moments. This placed each moment condition at a roughly equal magnitude.

In addition, as with most highly parameterized term structure models, we need an additional identifying restriction. Surprisingly, because of the model's structure and the many moments available for use, we need only one. We impose $\sigma_\eta = \sigma_\varepsilon = \sigma$.

Estimates

Table 1 reports the output from the GMM estimation. The parameter estimates imply a theoretical term structure and theoretical autocovariances for inflation and the short rate. To continue the reverse engineering metaphor, the question is now comparability—how closely does the reverse engineered term structure look like the original given by the data?

An informal but informative answer comes from visually comparing the theoretical prediction with the values given by the data. Figure 1 plots the actual and estimated term structure (i.e. that implied by the model parameterization), expressed as spreads over the one-month rate. Qualitatively, the model reproduces the term structure's concave, upward slope, though both the fit and the estimates differ across subsamples. For the pre-82 period, the model matches 6 and 12 months well but seriously overpredicts the 5 year rate, coming in a bit low on the 10 year rate. This produces a yield curve with less of a kink at one year than actually seen in the data. Furthermore, the predicted curve actually slopes downward (slightly) between 5 and 10 years.

The model does better in the post-1982 sample, though it again comes in too high for the 5-year rate and too low for the 10-year rate. Perhaps unsurprisingly, the model has a bit of trouble matching the rather sharp kink in the actual yield curve at one year. As a result, it predicts too high a value at 3 months and too low at 5 years.

Figure 2 makes a similar comparison for the autocovariances of the short rate. Particularly for the post-1982 sample, the model has a difficult time matching the slow decline of short rate autocovariances (as did Backus and Zin). It was this pattern that led Backus and Zin, in an earlier paper, to consider fractionally differenced stochastic processes, which can mimic such hyperbolic decay. As Backus and Zin point out, such a specification produces a mean yield curve that eventually slopes downward, something not observed at the maturities used in their (and our) paper. The term structure often does slope downward at maturities above 10 years, however, so such nonstationarity is not ruled out *a priori* by the data. In our case, the numerical methods we used become unreliable with extreme persistence, so we ruled out nonstationarity on computational grounds. The fit for the later sample is noticeably better, in part because the autocovariances themselves show a more geometric decline.

Along with both the term spreads and autocovariances, the model must also fit two additional sets of restrictions involving inflation. Figure 3 compares the predicted and actual autocovariances of inflation. (Cross-correlations are not plotted. Backus and Zin, with a purely nominal model, do not use the inflation process.) The second-order autoregressive process we use has mixed results. In the earlier period, though the process mimics the slow decline in autocovariances seen at lags 1 through 24, it misses on the variance (lag 0). In the later period, the model matches well at lags 0, 12, and 24, but does less well at lags 1 and 3. The large value on the second AR coefficient leads to a spiked pattern reminiscent of negative autocorrelation (though here the correlations are all positive). While somewhat strange looking, similar patterns appear if an AR(2) is fit directly to the inflation data, so this seems a robust characterization of the process. The earlier period shows a pattern of autocovariances that are significantly lower and smoother: as expected,

the estimated inflation process differs greatly across samples.

A more formal test of the model's fit is given by the J-statistic, the sample size times the value of the objective function at the estimated values, distributed as χ^2 , with degrees of freedom equal to the number of overidentifying restrictions (see Hamilton 1994, section 14.1). The model fails the J-test. Unlike the simpler Backus and Zin model, which also fails, this model fails largely because of the additional cross-covariance moments of the yield curve with inflation. (A J-test without these moments does better). Putting a small weight on these moments allowed us to match the autocovariances of the short rate and the yield curve more closely. There is also some evidence that fitting the short rate process made it harder for the AR(2) process to fit the inflation process. In one sense, considering an additional variable—inflation—provides a more rigorous test than any model that looks only at interest rates, no matter how many latent factors considered. Not surprisingly, getting both interest rates and inflation right is harder than getting interest rates alone right.

Robustness Issues

The fit of our model must also be judged against alternatives. To this end we looked at whether all parameters were necessary (in the extreme case reducing to a single-factor model) and the effect of alternative identifying restrictions. We also consider the possibilities of generalizing the model.

Table 2 reports the tests on removing parameters. The tests reject the one-factor model. Clearly $\theta_1 \neq \theta_2$ and $v_\eta \neq v_\varepsilon$. Further tests reject a simpler characterization of the time-series properties of the system. The inflation process is certainly more complex than a simple AR(1) ($\rho_2 \neq 0$), and the same can be said for the real kernel process as well ($\theta_1 \neq 0$ and $\theta_2 \neq 0$). Simpler versions are all resoundingly rejected.

Other identification schemes provided generally similar results. Thus we don't report the results for $\sigma_\eta = 1$ and $v_\eta = 1$. Some other restrictions, such as $v_\eta = v_\varepsilon$ or $\theta_1 = \theta_2$ did not identify the system.

At this stage we do not nest our model in a more general model. It is in principle possible (with additional identifying restrictions) to add a richer time series component to inflation or the real kernel. Alternatively, a third factor might be added. Given the additional assumptions needed to identify the richer models, and the complexity of the resulting estimation system, we feel it more useful to gain a basic understanding of the strengths and limits of the current model.

4. Implications and Applications

It might be tempting to interpret our results so far as estimating a more complicated (two-factor) stochastic process. While such an exercise has some use, we think the model's implications for real rates and inflation are more interesting. Our ARMA(4,4) process is heavily constrained by the imposed structure of our inflation model. If merely fitting the term structure were our goal, a purely nominal model unconstrained by the inflation data would do a better job. Simple curve fitting based on the unconstrained nominal model is likely to result in a better fit of the nominal term structure with a more parsimonious ARMA structure. In contrast to such curve fitting, our structural model has a natural interpretation, and an added benefit of providing most of the identifying restrictions needed.

Our model allows a closer look at interest rate dynamics and real interest rates. The dynamics implied by the estimated model show how real and nominal shocks interact to produce the observed nominal term structure. The model's explicit distinction between real and nominal produces estimates of the mean *real* term structure and historical real

rates.

Interest rate dynamics

One advantage of the reverse-engineering approach is that it gives an estimate of the dynamics of the pricing kernel, and thus of the short-term interest rate. Figures 4 (pre-1982) and 5 (post-1982) provide one view of these dynamics via a series of impulse response functions for the two components of the short rate, the real rate, and inflation. These trace out the estimated responses to a one-standard-deviation shock over 36 months.

Just as the inflation process differs across time periods, the response of inflation to shocks differs as well. For the pre-1982 sample, the η shock has an immediate negative impact on the inflation rate of 14 basis points; for ε , the impact is a positive 9 bp. The difference in sign suggests there is some wisdom in calling η a “real” shock and ε an “inflation” shock. The impulse responses for both ε and η decay quite slowly, still being nearly one-half of their original value after 24 months.

For the post-1982 period, the one-standard-deviation shocks again have opposite signs, but with a rather different impact on inflation : 55 basis points for η and 8 bp for ε . In addition, the impulse response functions decay faster than they did in the earlier period, though not, as it happens, monotonically.

The real rate, with an underlying ARMA structure, displays less interesting dynamics, also illustrated in figures 4 and 5. A large contemporaneous shock emerges as the major feature in either subsample. The 1,000 basis point effect (a full 10%) in the pre-1982 period looks suspiciously large, though the 500 basis point effect post-1982 is a bit more reasonable. We simply don’t see frequent changes of that magnitude in the data. Possibly, negative correlation between ε_t and η_t , not picked up by our model, may lead to shocks that tend to offset one another. In both periods, the ε effect turns slightly negative at one

month and then approaches zero. The impact of η drops effectively to zero immediately after the instantaneous response.

The response of the nominal rate is of course a combination of the two effects, with the effect on the real interest rate dominating, particularly in the immediate response. Perhaps interestingly, the pattern, that of a large initial shock followed by overshooting, matches that predicted by many limited participation models (such as Fuerst, 1992), where a restrictive monetary policy shock increases interest rates today because people cannot adjust their portfolios. Next period, when the adjustment takes place, the lower money supply means prices fall, and the Fisher effect reduces real rates by the expected deflation. Of course, we have not explicitly modeled these effects nor identified any monetary policy shock, so this pattern can only be construed as suggestive. More broadly, however, the results indicate that movements in short-term nominal rates are dominated by changes in the underlying real rate.

One very clear result is that η and ϵ each contribute to both the inflation and the real components of interest rates, so these cannot be considered independent. A shock to inflation will also be a shock to real rates.⁷ In macroeconomic models, of course, it is not surprising that underlying shocks have both real and nominal effects, but our results underscore the danger of identifying “inflation” or “the real side” as independent factors behind the term structure.

Real Rates

One straightforward application of the estimated model is an estimate of the real term structure. This follows from using the techniques of section 2 on m_t , the real kernel, and effectively ignoring the ρ terms in equation (3). That is, we use the process for the real

⁷ Algebraically, of course, one could find shock values that produce a real effect with no impact on inflation, but that calculation is unlikely to be informative.

kernel,

$$-\log m_t = \delta + \frac{1}{1 - \phi L}(1 + \theta_1 L)\varepsilon_t + \frac{1}{1 - \phi L}(1 + \theta_2 L)\eta_t$$

to price real bonds, using an appropriately reduced version of equation (8). For the mean real yield spreads,

$$E(y_t^n - y_t^1) = \frac{\sigma_\varepsilon^2}{2} - \frac{\sigma_\varepsilon^2}{2n} \sum_{j=1}^n B_{1,j-1}^2 + \frac{\sigma_\eta^2}{2} - \frac{\sigma_\eta^2}{2n} \sum_{j=1}^n B_{2,j-1}^2.$$

Here we also do the recursion only on the real terms of A, which we denote as B.

$$\begin{aligned} \sum_{j=1}^n B_{1,j-1}^2 &= [1]^2 \\ &+ [1 + (\theta_1 + \phi)]^2 \\ &+ [1 + (\theta_1 + \phi) + (\theta_1 + \phi)\phi]^2 \\ &+ \dots \end{aligned} \tag{10}$$

Determining the mean real term structure then involves simply substituting estimated parameter values into equation (10). Figure 6 plots the mean real term structure estimated from the model and provides as a comparison the mean predicted nominal rates as well.

First note, as is expected, the real term structure is lower than the nominal, reflecting the positive average inflation in both samples, 1951-1982 and 1982-1991. The difference tends to widen with maturity, supporting the not surprising view that long bonds are more vulnerable to inflation risk. During the early period, however, the difference does decrease at the long end; we conjecture that the markets felt some sort of nominal anchor existed—high inflation rates would not be permitted to continue, and thus longer bonds had less of an inflation premium. For the early sample, however, the inflation premium is reasonably large: 21 basis points at 12 months, 37 bp at 10 years.

The post-1982 sample shows a steeper real yield curve and much smaller inflation premium. The 10-year 3-month real spread stands at 186 bp, compared with 62 in the

earlier period. The inflation premium is quite small: at one year, the premium is 8 basis points: at 10 years it has only increased to 10 bp.

Finally, the real term structure does indeed slope up: the slope of the nominal thus reflects more than inflation fears. This is most noticeable after 1982 but exists even in the earlier period. The most likely explanation for this real term premium is a risk premium. Holding a 10-year bond, even a real one, for a year has risk because the next period's price is uncertain. This analysis of the real rate illustrates both the advantages and drawbacks of the kernel approach: It has enough structure to obtain estimates of the real term premia, but unlike an equilibrium model, it cannot give a fully satisfactory answer in terms of consumption risk or liquidity differences.

5. Conclusions

Previous work on the term structure of interest rates has often made extreme assumptions. This paper shows how the pricing kernel approach can be used to relax some of those assumptions. By allowing inflation and the real economy to affect each other, it generates results both for the term structure and the time path of interest rates. It thus represents an advance on work that divorces inflation from the real economy, that ignores inflation by working directly with nominal rates, or that specifies precise functional forms for utility and money demand.

What does such a new approach buy us in terms of understanding the term structure?

For the average term structure, the estimated model confirms what should be an intuitive result: both real and nominal factors contribute to the slope and curvature. Inflationary expectations increase the slope and may increase or decrease the curvature, but the real term structure slopes upward. Fitting the term structure, though, expresses only the cross-sectional aspects of the model, which has time-series implications as well.

An η shock initially pushes inflation down, and an ε shock pushes inflation up; both push real rates higher, and the real effect has the dominant effect on short-term rates, at least initially. Both underlying shocks to the economy affect nominal rates through “real” and “nominal” channels. Economically, of course, this makes perfect sense, as an increase in output will tend to decrease prices if money and velocity stay constant, and with sticky wages a nominal shock can have real effects. Furthermore, the response to these shocks is not always monotonic. The correlation between inflation and real rates depends heavily on the lags considered.

In the end, there is no magical, simple relationship that explains the term structure. We hope that our approach respects that complexity but at least partially untangles the varied shocks and dynamics that determine interest rates.

References

- Backus, David K., and Stanley E. Zin, Reverse Engineering the Yield Curve, *NBER Working Paper* No. 4676, (1994).
- Boudoukh, Jacob, An Equilibrium Model of Nominal Bond Prices with Inflation–Output Correlation and Stochastic Volatility, *Journal of Money, Credit, and Banking* 25, no. 3 part 2, 636–665, (1993).
- Brown, Stephen J., and Phillip H. Dybvig, The Empirical Implications of the Cox, Ingersoll, Ross Theory of the Term Structure of Interest Rates, *Journal of Finance* 41, 617–632, (1986).
- Buraschi, Andrea, and Alexie Jiltsov, How Large is the Inflation Risk Premium in the U.S. Nominal Term Structure? Working Paper, December 1999.
- Campbell, John Y., Bond and Stock Returns in a Simple Exchange Model, *The Quarterly Journal of Economics* 101, 785–803 (1986).
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ (1997).
- Constantinides, George M., A Theory of the Nominal Term Structure of Interest Rates, *The Review of Financial Studies* 5, 531–552 (1992).
- Cox John, Jonathan Ingersoll, and Stephen Ross, A Theory of the Term Structure of Interest Rates, *Econometrica* 53, 385–407 (1985).
- den Haan, Wouter J., The Nominal Term Structure of Interest Rates in Real and Monetary Economies, *Journal of Economic Dynamics and Control* 19, 909–940 (1995).
- Duffie, D. *Dynamic Asset Pricing Theory*, Princeton University Press, Princeton, NJ, (1992).
- Evans, Martin D.D., The Term Structure of Real Interest Rates: New Estimates and Implications, *Journal of Finance* 53, 187–218 (1998).
- Fuerst, Timothy S., Liquidity, Loanable Funds, and Real Activity, *Journal of Monetary Economics* 29, 3–24 (1992).
- Gibbons, M., and Krishna Ramaswamy, A Test of the Cox, Ingersoll, and Ross Model of the Term Structure, *Review of Financial Studies* 6, 619–658 (1993).
- Granger, C. W. J., and M. J. Morris, Time Series Modelling and Interpretation, *Journal of the Royal Statistical Society A* 139, part 2, 246–257 (1976).
- Hansen, Lars P., John C. Heaton, and Masao Ogaki, GMMQ.EXP, Gauss Program, (1989).
- Hansen, Lars P., and Ravi Jagannathan, Implications of Security Market Data for Models of Dynamic Economies, *Journal of Political Economy* 99, 225–262 (1991).
- Harrison, M., and D. Kreps, Martingales and Arbitrage in Multiperiod Securities Markets, *Journal of Economic Theory* 20, 381–408 (1979).

- Hayashi, Fumio, *Econometrics* Princeton University Press, Princeton (2000).
- Huang, C., and R. Litzenberger, *Foundations for Financial Economics* North Holland, New York, (1988).
- Huber, Peter J., "The Behavior of Maximum Likelihood Estimates under Non-standard Conditions" *Proceeds of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, Berkeley, pp.221–233 (1967) .
- Kreps, David., Arbitrage and Equilibrium in Economies with Infinitely Many Commodities, *Journal of Mathematical Economics* 8, 215–240 (1981).
- Labadie, Pamela, The Term Structure of Interest Rates over the Business Cycle, *Journal of Economic Dynamics and Control* 18, 671–697 (1994).
- McCulloch, J. Huston., and Heon-Chul Kwon, U.S. Term Structure Data, Ohio State Working Paper WP93-6 (1993).
- Pennacchi, George, "Identifying the Dynamics of Real Interest Rates and Inflation: Evidence Using Survey Data" *The Review of Economic Studies* 4, pp. 53-86 (1991).
- Pennacchi, George, Peter Ritchken, and L. Sankarasubramanian, On Pricing Kernels and Finite State Variable Heath Jarrow Morton Models, Working Paper (1995).
- Sargent, Thomas. J., *Dynamic Macroeconomic Theory*, Harvard University Press, Cambridge, Massachusetts (1987).
- Wu, Shu, How Do Changes in Monetary Policy Affect the Term Structure of Interest Rates?, Working Paper, University of Kansas (2001).
- Wu, Tao, Macro Factors and the Affine Term Structure of Interest Rates, Working Paper, Yale University (2000).

Table 1
Subsample 2-factor ARMA Parameter Estimates

variable	estimate	
	1952-82	1982-91
ϕ	0.6544 (0.0016)	0.9225 (0.0027)
ρ_1	0.8163 (0.0025)	0.0300 (0.0128)
ρ_2	0.1510 (0.0024)	0.2292 (0.0430)
θ_1	-0.6802 (0.0022)	-0.9499 (0.0021)
θ_2	-0.6430 (0.0016)	-0.9318 (0.0033)
v_ε	-0.0096 (0.0002)	0.0176 (0.0053)
v_η	0.0129 (0.0004)	0.0959 (0.072)
σ_ε	0.1555 (0.0084)	0.0473 (0.0016)
σ_η	0.1555 (na)	0.0473 (na)
$\chi^2(11)J$	174.16	106.30
$Prob < J$	1.00	1.00
min value	-0.4838	-4.204

Notes: $N = 359$ for the 1952-1981 subsample, and $N = 111$ for 1982-1991 subsample. Standard errors of the estimates are in parentheses underneath the estimated coefficients. Convergence using numerical derivatives was achieved for all specifications.


Table 2

Tests of Parameter Reduction

Period 1 restriction	χ^2	d.o.f.	Pr < χ^2	Period 2 χ^2	d.o.f.	Pr < χ^2
$\rho_2 = 0$	3970.9	1	1.00	28.38	1.00	1.00
$\theta_1 = 0$	97599.	1	1.00	2.1×10^5	1.00	1.00
$\theta_2 = 0$	1.6×10^5	1	1.00	81824	1.00	1.00
$\rho_2, \theta_1, \theta_2 = 0$	1.9×10^5	3	1.00	2.3×10^5	3	1.00
$\theta_1 = \theta_2$	1323.3	1	1.00	87.89	1.00	1.00
$v_\eta = v_\varepsilon$	1448.6	1	1.00	59.47	1.00	1.00
$v_\eta = v_\varepsilon, \theta_1 = \theta_2$	1457.0	2	1.00	136.52	2	1.00

Notes: $N = 359$ for the 1952-1981 subsample, and $N = 111$ for 1982-1991 subsample. Tests for both periods based on the two factor estimates.

Average Nominal Term Structure : Pre-Oct 1982 Period



Average Nominal Term Structure: Post-Oct 1982 Period





Figure 1

Short rate Autocovariances: Pre-Oct 1982 Period



Short rate Autocovariances: Post-Oct 1982 Period





Figure 2

Autocovariance of inflation: Pre-Oct 1982 Period



Autocovariance of inflation: Post-Oct 1982 Period




Figure 3









Figure 4










Figure 5

Predicted Nominal and Real Yields: Pre-Oct 1982 Period



Predicted Nominal and Real Yields: Post-Oct 1982 Period




Figure 6

Federal Reserve Bank
of Cleveland
Research Department
P.O. Box 6387
Cleveland, OH 44101

Address Correction Requested:

Please send corrected mailing label to the
Federal Reserve Bank of Cleveland
Research Department
P.O. Box 6387
Cleveland, OH 44101

PRST STD
U.S. Postage Paid
Cleveland, OH
Permit No. 385