consumption growth is low.[7] In term of the earlier decomposition, the key covariance term on the right-hand side now becomes

$$cov_t\left(\beta^j \frac{C_t}{C_{t+j}}, \frac{1}{\pi_{t,t+j}}\right)$$

What we can take away from this analysis is that fiscal or monetary policy induced changes in nominal yields must make investors either update their views on average real returns (which are directly linked to real consumption growth with this specific stochastic discount factor) and average inflation or the comovement between inflation and real consumption growth growth.[8] In particular, changes in fiscal and monetary policies could change investors' views of the government and thus lead them to update their perceptions of future economic growth and/or future inflation.

3 A Hitchhiker's Guide to Functional Time Series Methods

In this section, we give a high-level overview of the functional time series methodology that we use throughout our paper.[9] When large amounts of data are available on economic variables that are theoretically linked via a functional relationship (such as various nominal yields linked via the yield curve), such a functional approach can efficiently exploit this functional relationship.

We assume that observations of the nominal yield curve in a period t can be described by a function $y_t(\tau)$ defined over an interval I of possible maturities (between three months and 30 years in our case) taking real values (*that is,* $y_t : I \to \mathbb{R}$). The yield at time t for a security that matures in $t+\tau$ is thus given by $y_t(\tau)$ where τ is a value taken from the set I. We treat the function $y_t(\tau)$ as a random variable in a functional space, as it varies non-deterministically from one period to the next. To be concise, we will drop the argument τ from the y_t function unless needed.

The functional form of the yield curve we use here (Gürkaynak, Sack and Wright, 2007)

[7]The risk free real return on a j-period security with log utility is given by $\left[E_t\left(\beta^j \frac{C_t}{C_{t+j}}\right)\right]^{-1}$.

[8]Even with richer stochastic discount factors such as those derived using Epstein-Zin utility, consumption growth is still a key determinant - see for example Campbell (2017).

[9]More details are provided in the appendix or in Chang, Hu and Park (2022) and Chang, Park and Pyun (2023b).

allows us to obtain a yield for all values of τ between the aforementioned bounds of three months to thirty years. We describe in Appendix D how we, in practice, use a grid of values to represent the interval I and their corresponding images (yields) for each quarter t.

So far we have not restricted the yield curve in any way - the function $y_t(\tau)$ can take on arbitrary value for each maturity τ at any point in time t. We next describe the mild restriction we impose on the function $y_t(\tau)$ before turning to a description of a finite-dimensional approximation of this function, which we can then exploit in our empirical analysis.

3.1 Restrictions on the Yield Curve

In order to econometrically exploit the fact that all yields are linked via the yield curve, we will put one mild restriction on the yield curve. We only study yield curves that are in the space $H = \mathcal{L}^2(I)$, the space of square integrable functions.[10] While this space of functions is very general (it includes functions that are not continuous, for example), it still imposes a surprising amount of regularity. In particular, we can now define a scalar product and a norm in H: For f and g in the space H we obtain

$$\langle f,g\rangle = \int_I f(x)g(x)dx \quad \text{and} \quad \|f\| = \sqrt{\langle f,f\rangle}. \tag{5}$$

In addition to the inner product and the norm, we also can define a tensor[11].

$$(f \otimes g)v = \langle v,g\rangle f \tag{6}$$

for all v in H. In Appendix E we show how to use these constructs (scalar and tensor products) to define the expectation function and the covariance operator of random functions in H.

Using results from functional analysis[12] we find that the space H is a *separable Hilbert space*. These are spaces that admit a scalar product, such as the one defined above, and have *countable* bases. This means that every yield curve in H can be expressed as the linear

[10]The space of all (real) functions f_t defined over I such that $\int_I |f(x)|^2 dx < \infty$.

[11]If $H \equiv \mathbb{R}^n$, we have $f \otimes g = fg'$, *i.e.*, $f \otimes g$ reduces to the outer product, in contrast to the inner product $\langle f,g\rangle = f'g$, where f' and g' are the transposes of f and g. Note that $(f \otimes g)v = (fg')v = (v'g)f$ for all $v \in \mathbb{R}^n$ in this case

[12]See for example Folland (1999).

combination of countable many functions $\{v_i\}_{i=1,2,3,...}$:[13]

$$y_t = \sum_{i=1}^{\infty} \alpha_{it} v_i. \tag{7}$$

Since the functions $\{v_i\}$ are independent of t, once they are determined, the yield curve y_t is fully characterized by the sequence of real numbers $(\alpha_{1t}, \alpha_{2t}, \ldots)$. In other words, the yield curve can be analyzed through a sequence of real numbers, and every sequence of real numbers can be traced back to a yield curve by combining the basis functions $\{v_1, v_2, \ldots\}$ with the sequence $(\alpha_{1t}, \alpha_{2t}, \ldots)$.

This approach is different from models of the yield curve that start with focusing on the level, slope, and curvature of the yield curve (Diebold and Rudebusch, 2012): We are not imposing a particular set of functions to describe the yield curve - instead, we choose basis functions that jointly describe most of the fluctuations in the yield curve.

3.1.1 A Finite Dimensional Representation of the Yield Curve

The dimension of a space is given by the number of elements in its basis. By this logic, the space H is infinite dimensional as the basis $\{v_i\}_{i=1,2,3,...}$ that we used in (7) has infinitely many elements.

The next step in our approach is to define a finite-dimensional subspace of H. We do this by considering only functions resulting from a linear combination of the first m elements of the basis $\{v_1, v_2, \ldots, v_m\}$, these functions define the finite-dimensional space H_m (a subspace of H).

The function y_t is not an element of H_m given that we need more than just the first m elements of the basis to represent it as we can see in (7). However, we can consider the projection of y_t on H_m given by

$$\tilde{y}_t = \sum_{i=1}^{m} \alpha_{it} v_i. \tag{8}$$

This gives us an equation akin to an observation equation in a state space model

$$y_t = \sum_{i=1}^{m} \alpha_{it} v_i + w_t, \tag{9}$$

[13]Note that we have omitted the argument τ of the function, but v_i (and y_t) still refers to a *function*.

where $w_t = y_t - \tilde{y}_t$ is the approximation error we make by restricting ourselves to H_m. Under suitable conditions, this approximation error becomes asymptotically negligible. In what follows, we assume that $\{v_i\}$ is an orthonormal basis, i.e. $\|v_i\| = 1$ for all i and $\langle v_i, v_j \rangle = 0$ for all $i \neq j$. Under this assumption, we have

$$\alpha_{it} = \langle v_i, y_t \rangle$$

for all i and t.

Let us now introduce a mapping from H_m to $\mathbb{R}^m$

$$H_m \ni \tilde{y}_t \mapsto \alpha_t = \begin{pmatrix} \alpha_{1t} \\ \alpha_{2t} \\ \vdots \\ \alpha_{mt} \end{pmatrix} \in \mathbb{R}^m.$$

This mapping is one-to-one correspondence between H_m and $\mathbb{R}^m$.Therefore, with the basis $\{v_1, v_2, \dots, v_m\}$ and $\alpha_t = (\alpha_{1t}, \alpha_{2t}, \dots, \alpha_{mt})'$ we can recover $\tilde{y}_t$ through (8). The mapping is an isometry between H_m and $\mathbb{R}^m$, which preserves the norm.[14] As a result, we can study a vector autoregression (VAR) for α_t by least squares, rather than having to work directly in a functional space.

Using functional principal components, whose properties we discuss in Appendix C, we determine a basis of functions $\{v_i\}_{i=1,2,3,\dots}$ such that its first m elements generate $y_t^{(m)} = \alpha_t \in \mathbb{R}^m$ a "best" approximation of y_t. Note that we can thus effectively choose a very efficient set of basis functions for our purposes rather than restrict ourselves to an a priori chosen basis function such as monomials $\{1, \tau, \tau^2, \dots\}$.

Since functional principal components algorithm depend on the data, and since the sample of the yield curve we use vary with the external shock being analyzed, so does the portion of the variability explained by these approximations. Our choice of $m = 3$ explains more than 90% of the variability of y_t in every case. This principal components analysis (detailed in the appendix) also delivers a time series for the vector of weights $\alpha_t = \begin{pmatrix} \alpha_{1t} & \alpha_{2t} & \dots & \alpha_{mt} \end{pmatrix}'$.

[14] We can show that

$$\|\tilde{y}_t\|^2 = \sum_{i=1}^{m} \langle v_i, y_t \rangle^2 = \sum_{i=1}^{m} \alpha_{it}^2 = \|\alpha_t\|^2,$$

where we use the same notation $\|\cdot\|$ to denote the norm of a function in H_m and the norm of a vector in $\mathbb{R}^m$.

Figure 2: Description of the Functional Principal Components

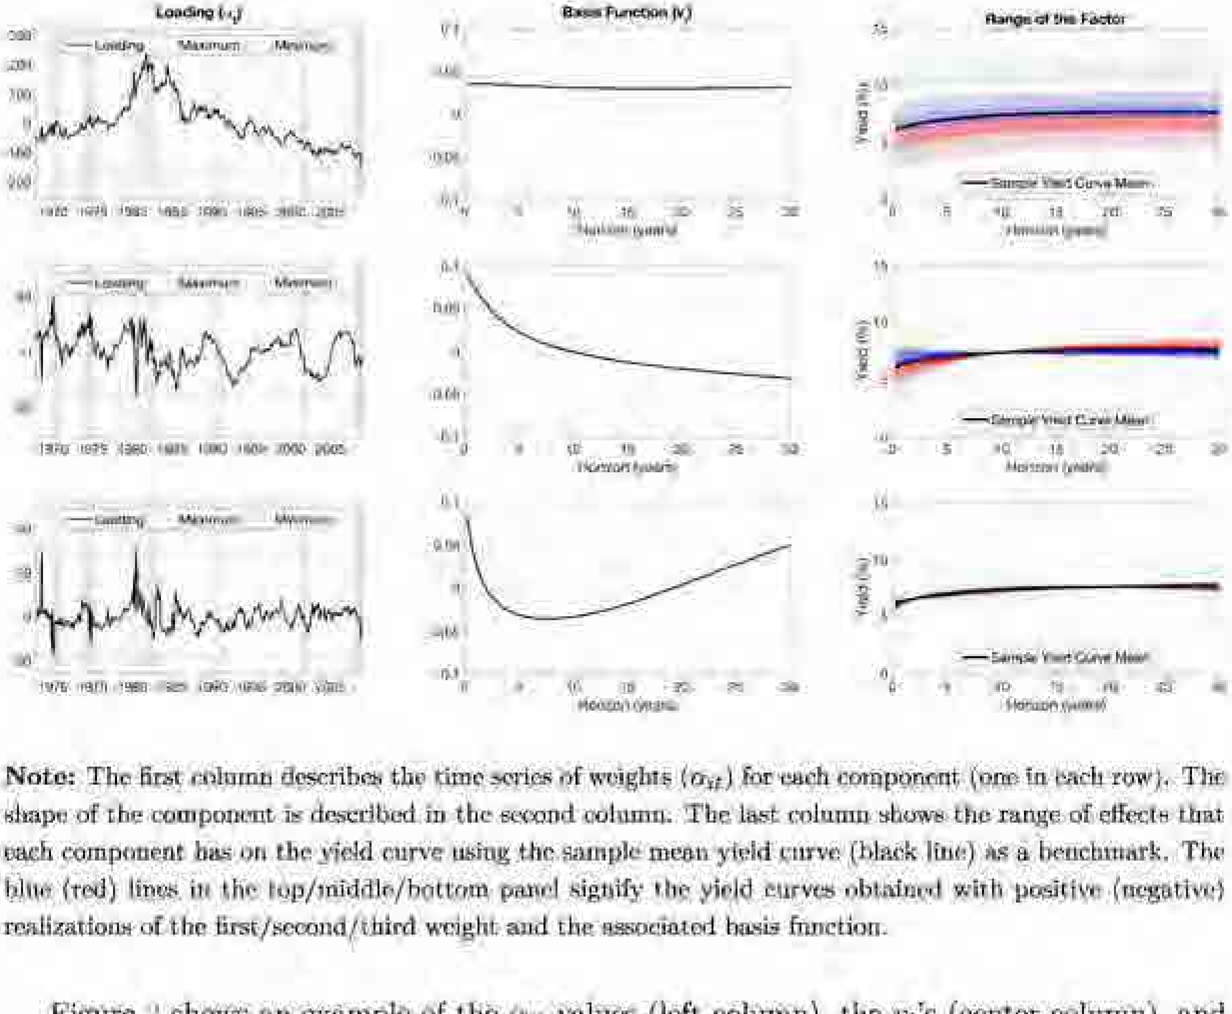

Note: The first column describes the time series of weights (α_{it}) for each component (one in each row). The shape of the component is described in the second column. The last column shows the range of effects that each component has on the yield curve using the sample mean yield curve (black line) as a benchmark. The blue (red) lines in the top/middle/bottom panel signify the yield curves obtained with positive (negative) realizations of the first/second/third weight and the associated basis function.

Figure 2 shows an example of the α_{it} values (left column), the v_i's (center column), and the range of yield curves generated by time series fluctuations in the α vector, using our yield curve data as described in Section [illegible], and in particular the sample mean as the benchmark value that is perturbed by movements in α.

This vector α_t cannot be directly interpreted as yields, as the measurement equation highlights that only together with the basis functions $\{v_1, v_2, \ldots, v_m\}$ can we recover the yield curve. It does, however, serve as the state in our state-space model for the yield curve.[15] An important feature of this approach is that for a fixed value of τ, the yield $y_t(\tau)$ is a *linear* combination of the elements of α_t, which makes construction of impulse responses straightforward since we assume a linear law of motion for that vector, as we discuss next.

3.2 The Dynamics of α_t and the Identification of Impulse Responses

We posit a VAR law of motion for α_t and the instrument for the policy shock of interest m_t.[16] In particular, we focus on a VAR(1) for the sake of parsimony:

$$\gamma_t = A\gamma_{t-1} + u_t, \tag{10}$$

where $\gamma_t \equiv [m_t\ \alpha_t']'$. From an applied perspective, our approach can be thought of as modeling observations on the yield curve (and the instrument) at each point in time t through a state-space framework with a set of observation equations (equation 9 and the identity $m_t = m_t$) that link the yield of an asset with a specific maturity to a set of basis functions that depend on the maturity and weights on each basis function, which vary over time, but do not depend on maturity. These weights represent (a subset of) the states in our state-space model, which we model as a Vector Autoregression (VAR) as in equation 10.

We identify the shock of interest by assuming a linear relationship between the forecast error u_t and the vector of structural shocks of interest e_t as

$$u_t = \Omega e_t, \tag{11}$$

and assume that Ω is computed via the lower triangular Cholesky decomposition of the covariance matrix of u_t so that $E(e_t e_t') = I$,[17] as proposed by Plagborg-Møller and Wolf (2021). The policy shock of interest is related to the first element of e_t, as we discuss below. This approach has a number of advantages, even beyond its simplicity. First, it automatically cor-

[15]The analogy to state space models might be slightly misleading because we first compute the states via principal components and then go on to model the law of motion for the estimated states, whereas standard applications of state space models often employ a filtering algorithm (think about the Kalman Filter, for example) that exploits a posited law of motion for the states when estimating the states. Our approach is instead very much reminiscent of the standard two-step approach to linear factor models in standard time series analyses (see, for example, Stock and Watson (2016)). The resulting model of the yield curve is still in state-space form.

[16]We estimate a separate VAR for each instrument because the sample sizes for the different instruments are not the same.

[17]I denotes the identity matrix.

rects for possible autocorrelation of the instrument and dependence of the instrument on past yield curve movements (which are generally thought to encode macroeconomic outcomes). To see this, it is useful to write the first equation of the set of Equations (10), using Equation (11):

$$m_t = A_{1,1}m_{t-1} + \sum_{j=1}^{m} A_{1,j+1}\alpha_{jt} + \Omega_{1,1}e_t^1, \tag{12}$$

where $A_{i,j}$ is the element of the matrix A in row i and column j. Following Plagborg-Møller and Wolf (2021), it is worthwhile to point out that this identification approach will correctly identify normalized impulse responses even if the yield curve itself does not contain enough information to identify the shock of interest ε_t (i.e., non-invertibility) and if there is measurement error w_t present in e_t^1, so that $e_t^1 = \theta\varepsilon_t + w_t$, where $\theta \neq 0$ is a parameter that influences the strength of identification and w_t is an i.i.d. measurement error. This comes at a cost, as we can only identify normalized impulse responses if there is non-invertibility. Throughout this paper, we plot impulse responses that increase the first element of e_t by one unit (which is equal to a one standard deviation change in the first element of the one-step ahead forecast error u_t). This has the advantage of giving us some sense of magnitude of the effects of the shock is indeed invertible.

In terms of inference, we use a bootstrap procedure that is detailed in Appendix F. [18] α_t and its associated basis functions have a clear interpretation in our application, as we highlight in Figure (2), which plots the basis functions (center column) associated with the first three elements of α_t - the basis functions resemble the level, slope, and curvature of the yield curve (Diebold and Rudebusch, 2012). Note, however, that we did not impose these shapes ex ante.

4 Yield Curve Data and Instruments

For the nominal yield curve, we use the data constructed by Gürkaynak, Sack and Wright (2007) that can be downloaded from the Board of Governors' website[19]. Our sample for the yield curve begins in June 1961. We use quarterly data; in particular, we use the curve observed the last day of each quarter as our quarterly yield curve. This ensures that shocks occurring at any point in a quarter can influence the yield curve in that same quarter. The exact sample to determine the response of the yield curve to each policy shock is the largest

[18]This bootstrap procedure is valid, as shown by Chang, Park and Pyun (2023b).
[19]https://www.federalreserve.gov/data/nominal-yield-curve.htm