The Influence of Fiscal and Monetary Policies on the Shape of the Yield Curve

Yoosoon Chang* Fabio Gómez-Rodríguez† Christian Matthes‡

September 2, 2025

Abstract

We study how fiscal and monetary policy shape the nominal yield curve and associated term premia. Government spending affects the long end of the curve, while tax changes and monetary policy influence the short end at impact. Within spending categories, only government consumption shifts the short end, but these effects dissipate within a year. While monetary policy and government consumption operate primarily through expected short rates, other fiscal interventions affect yields mainly by altering term premia.

JEL Classification: E50, E62, G10

Keywords: Yield Curve, Fiscal Policy, Monetary Policy, Functional Time Series

We would like to thank seminar and workshop participants at Indiana University, Lehigh University, the European Winter Meetings of the Econometric Society, 2024 SNDE Symposium, 2025 Symposium on Econometric Theory and Applications, and 2025 IAAE conference. The views expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Costa Rica. Early versions of this paper were presented under the title 'How Do the U.S. Government's Decisions Affect Its Borrowing Costs?'.

*Indiana University, yoosoon@indiana.edu.

†Lehigh University & Central Bank of Costa Rica, fag221@lehigh.edu

‡University of Notre Dame, cmatthes@nd.edu.

1 Introduction

How do policy actions influence the yield curve? We study this question using a unified and flexible statistical framework, combined with standard identification strategies for fiscal and monetary policy shocks. Understanding these effects matters for several reasons.

First, the yield curve plays a central role in the economy. Many interest rates that shape macroeconomic outcomes—such as mortgage rates—track it closely. Second, recent theoretical work identifies conditions under which fiscal and monetary policy can have similar effects on the broader economy (Correia *et al.*, 2008; Wolf, 2021). We complement this literature by showing that, historically, fiscal policy—especially government spending—has affected the yield curve in ways that differ from monetary policy. Third, we examine how fiscal and monetary decisions influence government borrowing costs. While governments finance a large share of spending through debt issuance, surprisingly little work has studied how fiscal actions shape the yield curve for nominal government liabilities. Existing research focuses primarily on the macroeconomic effects of tax and spending changes (e.g., Romer and Romer, 2010; Mertens and Ravn, 2013; Blanchard and Perotti, 2002; Auerbach and Gorodnichenko, 2012; Ramey, 2011; Ramey and Zubairy, 2018), not on how those decisions affect the government's own financing terms.[1]

Our paper fills this gap by documenting how fiscal policy and monetary policy move the nominal yield curve and, in turn, affect government borrowing costs. These effects are critical for setting fiscal policy. In equilibrium models of optimal fiscal policy, governments must consider how their decisions—and possibly those of the central bank—shift the yield curve (see, e.g., Lucas and Stokey, 1983; Barro, 1979; Buera and Nicolini, 2004; Angeletos, 2002).[2] We estimate the effects of both types of policy within a single empirical framework, avoiding the comparability issues that arise when monetary and fiscal effects are studied separately (e.g., using results from Piazzesi, 2005; Ireland, 2015).

Beyond these theoretical motivations, there are empirical reasons to expect fiscal policy to influence the yield curve. Ang and Piazzesi (2003) and Evans and Marshall (2007), for example, show that macroeconomic factors are key drivers of the nominal yield curve. Moreover, a growing literature finds that fiscal policies can exert powerful effects on economic

[1] We focus on nominal U.S. government debt. The market for inflation-indexed bonds (TIPS) is smaller, less liquid, and covers a shorter sample. Moreover, the nominal yield curve is already a central object of study in economics and finance.

[2] Even without assuming rational expectations, understanding how fiscal actions affect prices remains important for fiscal policy design; see Karantounias (2020).

activity. Whether and how those policies translate into changes in borrowing costs is the central question we address.

Figure 1: Correlation Between Policy Shocks and Yield Curve Changes

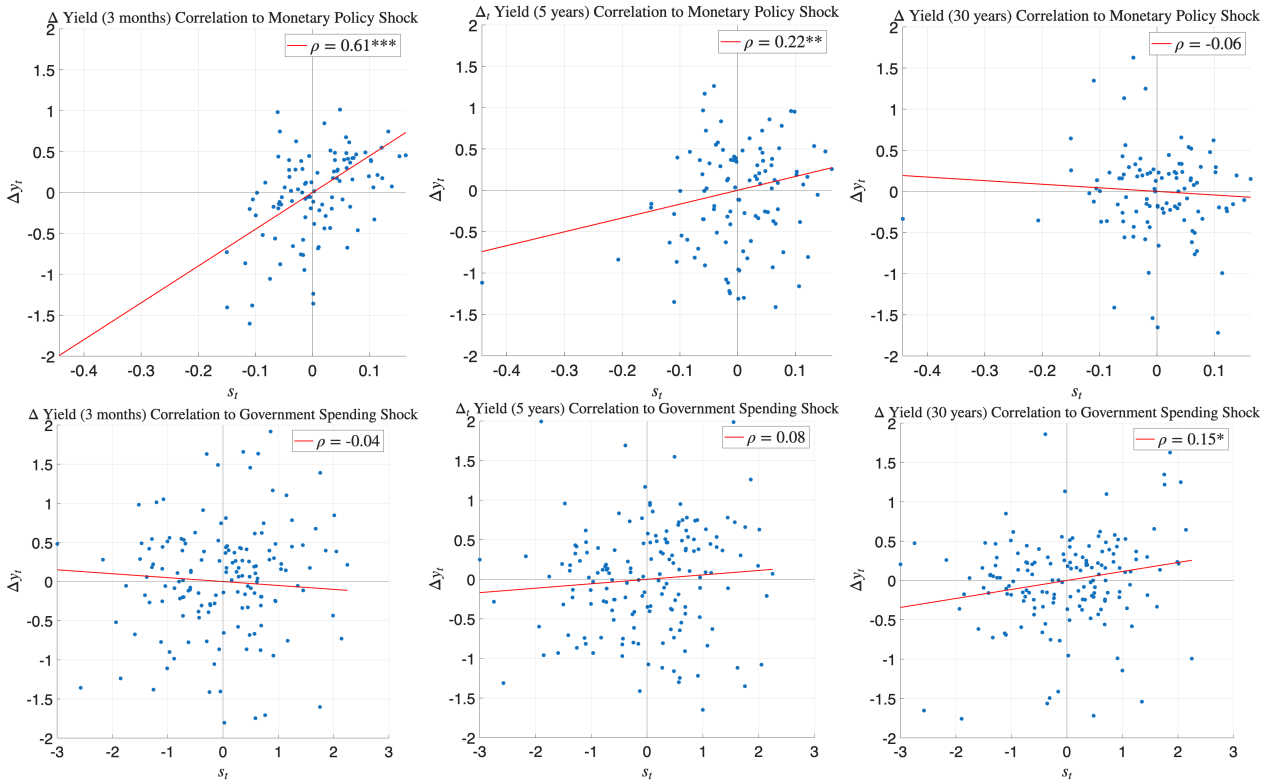

Note: Scatter-plots showing changes in yields Δy_t and instruments for policy shocks: monetary policy (Aruoba and Drechsel, 2022) and government spending shock (Auerbach and Gorodnichenko, 2012). We report the estimated correlation in the legend of each panel, where the significance levels are denoted by $^{*}: 10\%, ^{**}: 5\%, ^{***}: 1\%$. Information on data sources can be found in Appendix B.

Figure 1 provides a first look at how government decisions move the yield curve. It plots quarterly changes in yields at different maturities against instruments for monetary policy shocks (Aruoba and Drechsel, 2022) and government spending shocks (Auerbach and Gorodnichenko, 2012). Both sets of shocks are significantly correlated with changes in yields, but they affect different segments of the curve: monetary policy moves short maturities, while government spending shifts longer maturities.

These scatterplots are only suggestive. First, we do not directly observe structural shocks but instead rely on noisy instruments, and the figures do not correct for this measurement error. Second, they do not account for how the yield curve moves as a whole. Our empirical framework addresses both issues and extends the analysis to include tax changes and dis-

aggregated components of government spending. Still, the broad pattern—different policy shocks affecting different parts of the yield curve—remains one of our key findings.

To study the yield curve without imposing strong functional form assumptions, we adopt recent tools from the theory of functional time series (Chang *et al.*, 2016). We treat the yield curve at each point in time as a realization of a random function. With minimal assumptions, any such function can be expressed as a sum of basis functions weighted by time-varying coefficients. A key insight from the functional data literature is that a small number of basis functions often suffice to approximate the curve well. Our analysis focuses on these time-varying loadings, which summarize the dynamics of the yield curve. We show that this representation admits a natural state-space formulation.

This framework allows us to use the full yield curve—rather than just a few points or linear combinations of yields—to analyze how it evolves over time. In contrast to standard principal components approaches, which reduce dimensionality by analyzing fixed-weight combinations of selected yields, our method captures changes in the curve as a function, allowing us to efficiently use all available information. We study the forecasting performance of these different approaches for nominal yields in Appendix G. For further discussion and comparisons of the functional and traditional approaches, see Chang, Durlauf, Lee and Park (2023a) and Bjørnland, Chang and Cross (2023).

Beyond tracking the evolution of the yield curve, our goal is to identify the causal effects of government actions. To do so, we use standard instruments for fiscal and monetary shocks. Specifically, we employ exogenous variation in total government spending, defense spending, government consumption, and investment from Auerbach and Gorodnichenko (2012); tax shocks identified by Romer and Romer (2010); and monetary policy shocks constructed by Aruoba and Drechsel (2022). We estimate how these instruments relate to the dynamics of the yield curve—captured by the loadings on the basis functions—which allows us to compute impulse responses of the entire curve to each type of policy shock.

Our paper builds on, but is distinct from, several related contributions. Berndt *et al.* (2012), for example, examine how defense spending shocks influence government financing decisions, such as net surpluses and portfolio returns. Our focus, instead, is on how different types of fiscal policy shift nominal borrowing costs across the maturity spectrum. Plosser (1987) studies how fiscal forecast errors propagate to yields, while we estimate causal effects using external instruments. Dai and Philippon (2005) analyze the yield-curve impact of fiscal policy in a no-arbitrage framework with Blanchard and Perotti (2002)-style identification. We adopt a more flexible statistical approach that incorporates both fiscal and monetary policy

shocks and emphasizes causal identification via instruments, arguably the most common approach in empirical macroeconomics nowadays.

The closest applied paper using a functional approach is Inoue and Rossi (2021), who apply a Nelson-Siegel specification within a VAR to estimate how unconventional monetary policy affects the level, slope, and curvature of the yield curve. Chang *et al.* (2021) uses functional data methods to study how cross-sectional micro distributions respond to macroeconomic shocks and relate to aggregate dynamics.

In the next section, we draw on economic theory to further motivate our study and to provide potential mechanisms for how policy affects the yield curve. Section 3 outlines our empirical methodology with an emphasis on accessibility for applied researchers. We then present our main results.

2 Two Concepts from Economic Theory

We begin by illustrating one link between fiscal and monetary policy changes and the nominal yield curve via the government budget constraint, and then introduce a yield decomposition that defines a term premium.

Following Berndt *et al.* (2012), we analyze the government budget constraint but focus on the *nominal* version to match our interest in nominal yields. Specifically,

$$B_{t+1} = R^b_{t+1}\Big(B_t - S_t\Big), \tag{1}$$

where B_t is the nominal value of outstanding debt at the start of period t, S_t is the nominal primary surplus, and R^b_{t+1} is the nominal gross return on the government's portfolio between t and $t+1$.[3]

By log-linearizing, we can approximate this constraint (see also Berndt *et al.*, 2012) as

$$ns_t - b_t = E_t\Big[\sum_{j=1}^{\infty} \rho^j \left(r^b_{t+j} - \Delta ns_{t+j}\right)\Big], \tag{2}$$

where ns_t denotes the weighted log nominal primary surplus ratio (think of it as a measure of nominal surpluses), $b_t = \log B_t$, $r^b_t = \log R^b_t$, and $\rho \in (0,1)$. Intuitively, changes in $ns_t - b_t$ must show up in either future returns on the government portfolio (r^b_{t+j}) or future surpluses.

[3] Hall and Sargent (2011) propose methods for computing theory-consistent measures of R^b_{t+1}. See also Hilscher *et al.* (2014), Giannitsarou and Scott (2008), and Chung and Leeper (2007).

Berndt *et al.* (2012) focus on how defense spending affects this decomposition. Our emphasis is broader: we study how changes in different components of government spending, tax rates, and monetary policy alter the government's borrowing costs. These costs, captured by r_t^b, are encoded in the nominal yield curve (see Hall and Sargent, 2011). We therefore examine the effects of policy actions across the *entire* yield curve.

We next use the expectations hypothesis to define a time-t term premium $r_t(\tau)$ for a bond of maturity τ. Let $y_t(\tau)$ be the yield of that maturity, and let τ_0 denote a shorter reference maturity. The expectations hypothesis implies

$$y_t(\tau) = \underbrace{\frac{1}{n}\sum_{i=0}^{n-1} E_t\left[y_{t+i}(\tau_0)\right]}_{\text{expectation}} + \underbrace{r_t(\tau)}_{\text{term premium}}, \tag{3}$$

where n is the number of periods between τ_0 and τ. Our modeling framework will produce estimates of the expectation term, and we define the term premium $r_t(\tau)$ as the difference between the observed yield $y_t(\tau)$ and this expectation component.

3 A Hitchhiker's Guide to Functional Time Series Methods

We now provide a high-level explanation of the functional time series methodology used throughout this paper.[4] This approach is particularly effective when large amounts of data describe economic variables that are theoretically linked via a functional relationship, such as nominal yields along the yield curve. By treating these data as realizations of a random function, we can leverage their inherent structure more efficiently than traditional methods.

In Appendix G, we provide evidence that such a functional time series approach delivers competitive and often superior forecasts for nominal yields relative to standard VAR and principal components approaches.

We assume that observations of the nominal yield curve at time t can be represented by a function

$$y_t : I \to \mathbb{R},$$

where I is an interval of maturities (from three months to 30 years in our case). For a security maturing in $t+\tau$, the yield at time t is $y_t(\tau)$, with $\tau \in I$. We treat the entire function $y_t(\tau)$

[4] See the Appendix or Chang *et al.* (2022, 2023b) for further details.

as a random variable in a functional space, acknowledging that it varies stochastically across time. For brevity, we often write y_t instead of $y_t(\tau)$.

The data for the yield curve we use here (Gürkaynak, Sack and Wright, 2007) allows us to obtain a yield for all values of τ between the aforementioned bounds of three months to thirty years. In Appendix B we discuss their approach to computing yield curves and how it relates to our functional principal components approach. We describe in Appendix D how we, in practice, use a fine 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.[5] 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

$$\langle f, g\rangle = \int_I f(x)g(x)dx \quad \text{and} \quad \|f\| = \sqrt{\langle f, f\rangle}. \tag{4}$$

In addition to the inner product and the norm, we also can define a tensor product[6]

$$(f \otimes g)v = \langle v, g\rangle f \tag{5}$$

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.

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

[6] 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.

Using results from functional analysis[7] 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 combination of countably many functions $\{v_i\}_{i=1,2,3,\ldots}$:[8]

$$y_t = \sum_{i=1}^{\infty} \alpha_{it} v_i. \tag{6}$$

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 $\{v_i\}_{i=1,2,3,\ldots}$ that we used in (6) 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 (6). However, we can consider the projection of y_t on H_m given by

$$\tilde{y}_t = \sum_{i=1}^{m} \alpha_{it} v_i. \tag{7}$$

[7] See for example Folland (1999).

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

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

$$y_t = \sum_{i=1}^{m} \alpha_{it} v_i + w_t, \tag{8}$$

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 $\{v_1, v_2, \ldots, v_m\}$ and the vector $\alpha_t = (\alpha_{1t}, \alpha_{2t}, \ldots, \alpha_{mt})'$, we can recover $\tilde{y}_t$ through (7). The mapping is an isometry between H_m and $\mathbb{R}^m$, which preserves the norm.[9] 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 $\{v_i\}_{i=1,2,3,\ldots}$ such that its first m elements generate an 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 a priori chosen basis functions such as monomials $\{1, \tau, \tau^2, \ldots\}$. Comparing our approach to a standard principal components approach, it might be useful to note that the basis function here serve the same role as the weights

[9] 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 $\mathbb{R}^m$.

on the principal components in standard finite dimensional principal component analysis. Intuitively, the functional approach, because it uses infinite dimensional data, replaces the sums present in standard finite dimensional principal component analysis with integrals, but the underlying ideas remain the same.

This principal components analysis also delivers a time series for the vector of loadings $\alpha_t = (\alpha_{1t}, \alpha_{2t}, \ldots, \alpha_{mt})'$. Figure 2 displays 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_t vector (right column). This range is derived from the yield curve data described in Section 4. Each plot shows how variations in a single element of α_t, denoted α_{it}, influence the sample mean yield curve (shown in black) in the direction of the corresponding basis vector v_i for $i = 1, 2, 3, 4$. Note how the first three basis functions resemble the level, slope, and curvature factors common in the yield curve literature (Diebold and Rudebusch, 2012). However, we did not impose these shapes ex ante. The vector α_t is not directly interpretable as yields; only in conjunction with the basis functions $\{v_1, v_2, \ldots, v_m\}$ does it reconstruct the yield curve.

Figure 2: Description of the Functional Principal Components

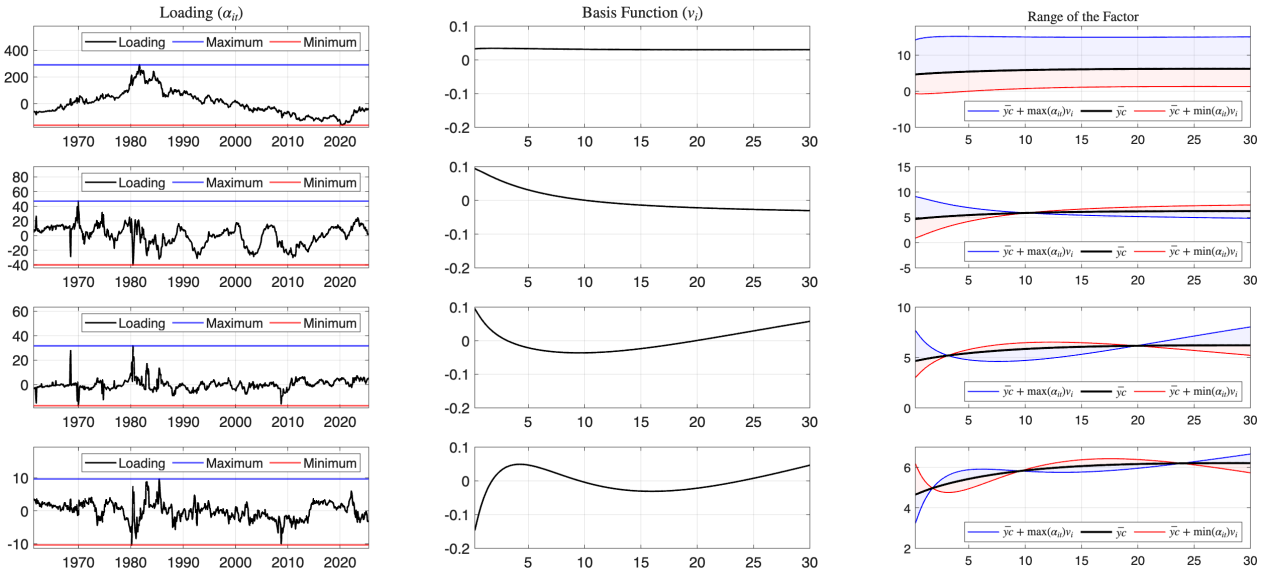

Note: The first column shows the time series of loadings (α_{it}) for each component (one in each row). The shape of the component is described in the second column. The last column shows the range of effects that each component has on the yield curve using the sample mean yield curve (black line) as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

Nonetheless, α_t serves as the *state* in our state-space representation of the yield curve.[10]

A crucial feature of this framework is that for each maturity τ, the yield $y_t(\tau)$ is a *linear* function of α_t. This linearity simplifies the construction of impulse responses, given the linear dynamics we impose on α_t (discussed in the next section).

3.2 The Dynamics of α_t and the Identification of Impulse Responses

We specify a VAR law of motion for α_t and the policy-shock instrument z_t.[11] In particular, we use a VAR(p) model and later discuss how we select p jointly with the number of factors m:

$$\gamma_t = c + \sum_{l=1}^{p} A_l \gamma_{t-l} + u_t, \tag{9}$$

where $\gamma_t \equiv [z_t \;\; \alpha_t']'$. This equation can be seen as the state equation for a state-space model for the yield curve at each time t.

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

$$u_t = \Omega e_t, \tag{10}$$

and assume that Ω is computed via the lower triangular Cholesky decomposition of the covariance matrix of u_t so that $E(e_t e_t') = I$,[12] as proposed by Plagborg-Møller and Wolf (2021). The policy shock of interest, denoted ε_t, is related to the first element e_t^1 of e_t, as we discuss below. This approach has a number of advantages, even beyond its simplicity. First, it automatically corrects 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 rewrite the equation for z_t implied by

[10] Here, the analogy to state-space models is slightly loose because we first estimate α_t via principal components and then specify its law of motion. In contrast, standard state-space applications often use a filtering algorithm (e.g., the Kalman Filter) that leverages the assumed law of motion to estimate the states. Our approach instead resembles the two-step estimation in linear factor models (e.g., Stock and Watson, 2016) by first extracting factors and then estimating the dynamics of these factors. However, the resulting yield-curve model is still in state-space form.

[11] We estimate a separate VAR for each instrument because the sample sizes differ across instruments.

[12] I denotes the identity matrix.

Equation (9) using Equation (10):

$$z_t = \sum_{l=1}^{p} A_l^{1,1} z_{t-l} + \sum_{l=1}^{p} \sum_{j=1}^{m} A_l^{1,1+j} \alpha_{j,t-l} + \Omega_{1,1} e_t^1, \tag{11}$$

where $A_l^{i,j}$ is the element of the matrix A_l 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 u_t present in e_t^1, so that $e_t^1 = \theta \varepsilon_t + u_t$, where $\theta \neq 0$ is a parameter that influences the strength of identification and u_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 if the shock is indeed invertible.

For statistical inference, we use a bootstrap procedure that is detailed in Appendix F.[13]

4 Yield Curve Data and Instruments

Our nominal yield-curve data come from Gürkaynak, Sack and Wright (2007), available on the Board of Governors' website.[14] The sample begins in June 1961. We use quarterly data, taking the yield curve observed on the last day of each quarter so that any shocks within the quarter can affect that quarter's yield curve. For each policy shock, we use the longest intersection between the yield-curve sample and the relevant shock series. We opt for quarterly data because many of our fiscal shocks are available only at that frequency.

As instruments for government spending shocks, we rely on the VAR-identified shocks of Auerbach and Gorodnichenko (2012).[15] Using VAR-identified shocks as instruments is common in applied work (e.g., Känzig, 2021). For tax shocks, we use the exogenous series documented by Romer and Romer (2010). We take our monetary policy shocks from Aruoba and Drechsel (2022), who employ machine learning and natural language processing to aug-

[13] This bootstrap procedure is valid, as shown by Chang, Park and Pyun (2023b).

[14] https://www.federalreserve.gov/data/nominal-yield-curve.htm

[15] We deviate slightly from Auerbach and Gorodnichenko (2012) by controlling for forecasts of overall government spending in *all* VAR specifications to ensure our identified shocks are truly unforecastable. The results are very similar if we follow Auerbach and Gorodnichenko (2012) exactly.

ment FOMC staff forecasts with information from internal FOMC documents. The deviation of actual FOMC decisions from these enriched forecasts serves as the monetary policy shock measure. A key advantage of this series is its longer sample coverage compared to popular high-frequency instruments based on interest-rate futures (Kuttner, 2001; Gertler and Karadi, 2015). In Appendix A.1, we show that our findings are robust to using the shock measure in Miranda-Agrippino and Ricco (2021) instead.

5 Response of the Yield Curve to Policy Shocks

We begin by plotting impulse responses for the entire yield curve, illustrating how yields at different maturities change h periods after a shock. Here, h is the *horizon* of the response, distinct from the *maturity* τ on the x-axis (in years). To estimate our model, we must choose the number of factors m and the lag length p. We select these jointly by minimizing average one-step-ahead forecast errors for yields with maturities between three months and 30 years. As a benchmark, we pick m and p from a VAR containing only α_t, allowing us to use one specification for all results. This procedure yields $m = 4$ and $p = 1$. These four factors explain 99.87% of the yield-curve variance. In Appendix A.3, we confirm that our findings are robust to choosing m and p separately for each shock, allowing for differences in observables and sample periods.

We first examine the impact response ($h = 0$) of various fiscal and monetary policy shocks. Our yield data represent quarter-end yields, so these impact responses capture changes within the same quarter as the shock. In all figures, the yield curve's response to monetary policy appears in light blue for ease of comparison.

Figure 3 (left panel) shows the response to a total government spending shock, identified using the Auerbach and Gorodnichenko (2012) instrument. In contrast to the short-end response triggered by monetary policy —significant out to maturities of about 10 years— government spending does not significantly affect yields under 12 years, after which the response becomes significant at longer maturities. The right panel shows the response to a tax shock, identified via the exogenous tax changes in Romer and Romer (2010). Interestingly, its impact response resembles the monetary policy response but in the opposite direction, and remains significant up to maturities of 15 years, suggesting a more persistent yield-curve impact than monetary policy.

We next investigate whether different components of government spending have the same effect on the yield curve. This question is motivated by Boehm (2020), who shows that the

Figure 3: Response at Impact of the Yield Curve to a Government Spending Shock (Auerbach and Gorodnichenko, 2012) and Exogenous Tax Change (Romer and Romer, 2010)

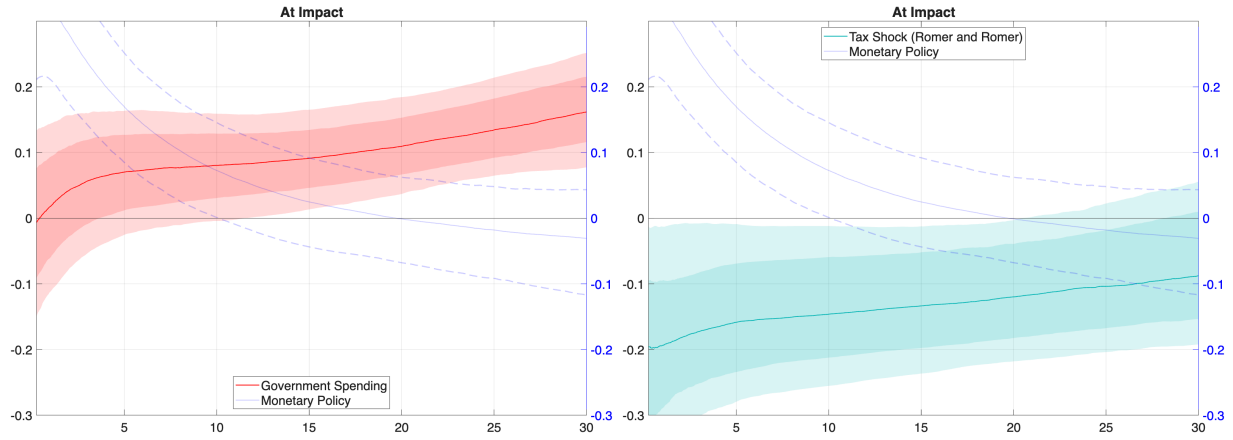

Note: For reference the response to monetary policy (Aruoba and Drechsel, 2022) and its 90% confidence band is represented in blue. The lighter (darker) shade signifies 90% (68%) confidence bands. 90% and 68% confidence bands estimated using bootstrap methods (see Appendix F).

fiscal multiplier can vary significantly between government investment and consumption. To measure shocks to defense spending, government consumption, and government investment, we use instruments constructed as in Auerbach and Gorodnichenko (2012).

Figure 4 (left panel) shows that defense spending shocks generate a response similar to overall government spending, significantly affecting only the long end of the yield curve.[16] The same pattern emerges for government investment, as shown in the right panel. Government consumption, however, moves only the very short end of the yield curve, and its effect is even more concentrated at short maturities than that of monetary policy.

We now consider responses for up to five years after the shock. To visualize these, we select certain maturities[17] and plot their impulse responses over time. Figure 5 reports these results for overall government spending, tax shocks, and monetary policy shocks, respectively.

For government spending shocks, the immediate effect grows larger at longer maturities, aligning with our previous analysis. After the shock, responses at all maturities quickly return to zero, except for the 30-year maturity, which displays slightly more persistence. Beyond that, the overall response shapes remain broadly similar across horizons. Tax shocks cause similar changes at all maturities, and these also revert to zero soon after the initial

[16] One might wonder why we do not use the defense news shock of Ramey (2011). Our sample starts after the Korean War, and as Ramey (2016) notes, that instrument has limited relevance in any sample beginning after the Korean War.

[17] For more maturities see Figures A-3–A-5 in the appendix.

Figure 4: Response at Impact of the Yield Curve to a Government Defense Spending Shock (left), Government Consumption Spending Shock (center), and, Government Investment Spending Shock (right) (Auerbach and Gorodnichenko, 2012)

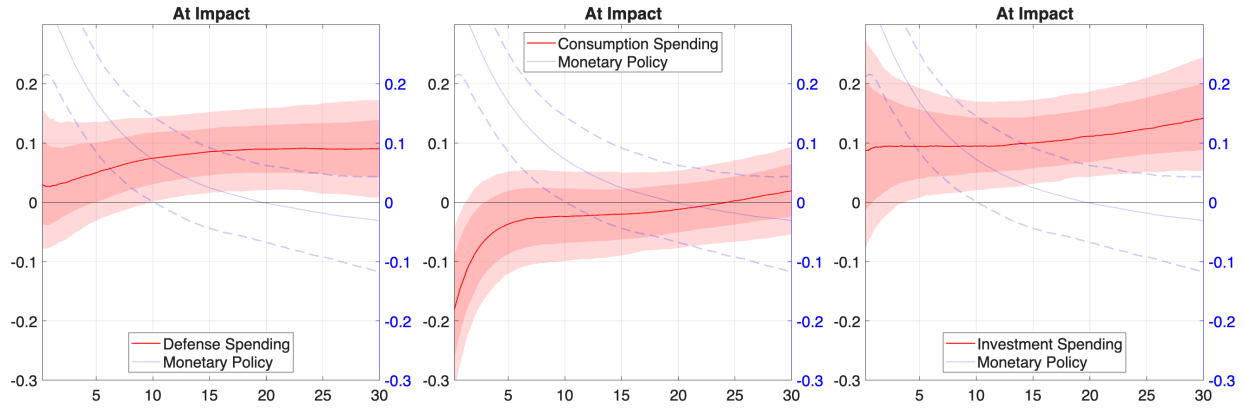

Note: For reference the response to monetary policy (Aruoba and Drechsel, 2022) and its 90% confidence band is represented in blue. The lighter (darker) shade signifies 90% (68%) confidence bands. 90% and 68% confidence bands estimated using bootstrap methods.

negative impact on yields. In contrast, monetary policy shocks reveal notable differences across maturities. At longer maturities, we observe a brief decline, possibly reflecting how higher current short-term rates can dampen inflation expectations, allowing the central bank to reduce short-term nominal rates in the future.

Finally, we decompose the impulse responses into average expected short rates up to each maturity and the term premium. Recall the decomposition in Equation (3), and note how the yield for a specific maturity τ depends on α_t:

$$y_t(\tau) = \sum_{i=1}^{m} \alpha_{it}\, v_i(\tau) + w_t(\tau),$$

where we now highlight τ. The expected future yield at horizon h is

$$E_t\Big[y_{t+h}(\tau)\Big] = \sum_{i=1}^{m} E_t[\alpha_{i,t+h}]\, v_i(\tau),$$

and we can compute $E_t[\alpha_{i,t+h}]$ using the VAR in Equation (9). This lets us obtain impulse responses for the expectations component of $y_t(\tau)$, and from those we derive the term premium's response as the difference between the total yield response and the response of the expectations component.

Figure 6 reports the decomposition at impact ($h = 0$) for all maturities from three months

Figure 5: Response of yields of specific maturities (given in the title of each subplot) to a government spending shock(top, Auerbach and Gorodnichenko (2012)), tax shock (middle, Romer and Romer (2010)), and monetary policy shock (bottom, Aruoba and Drechsel (2022)).

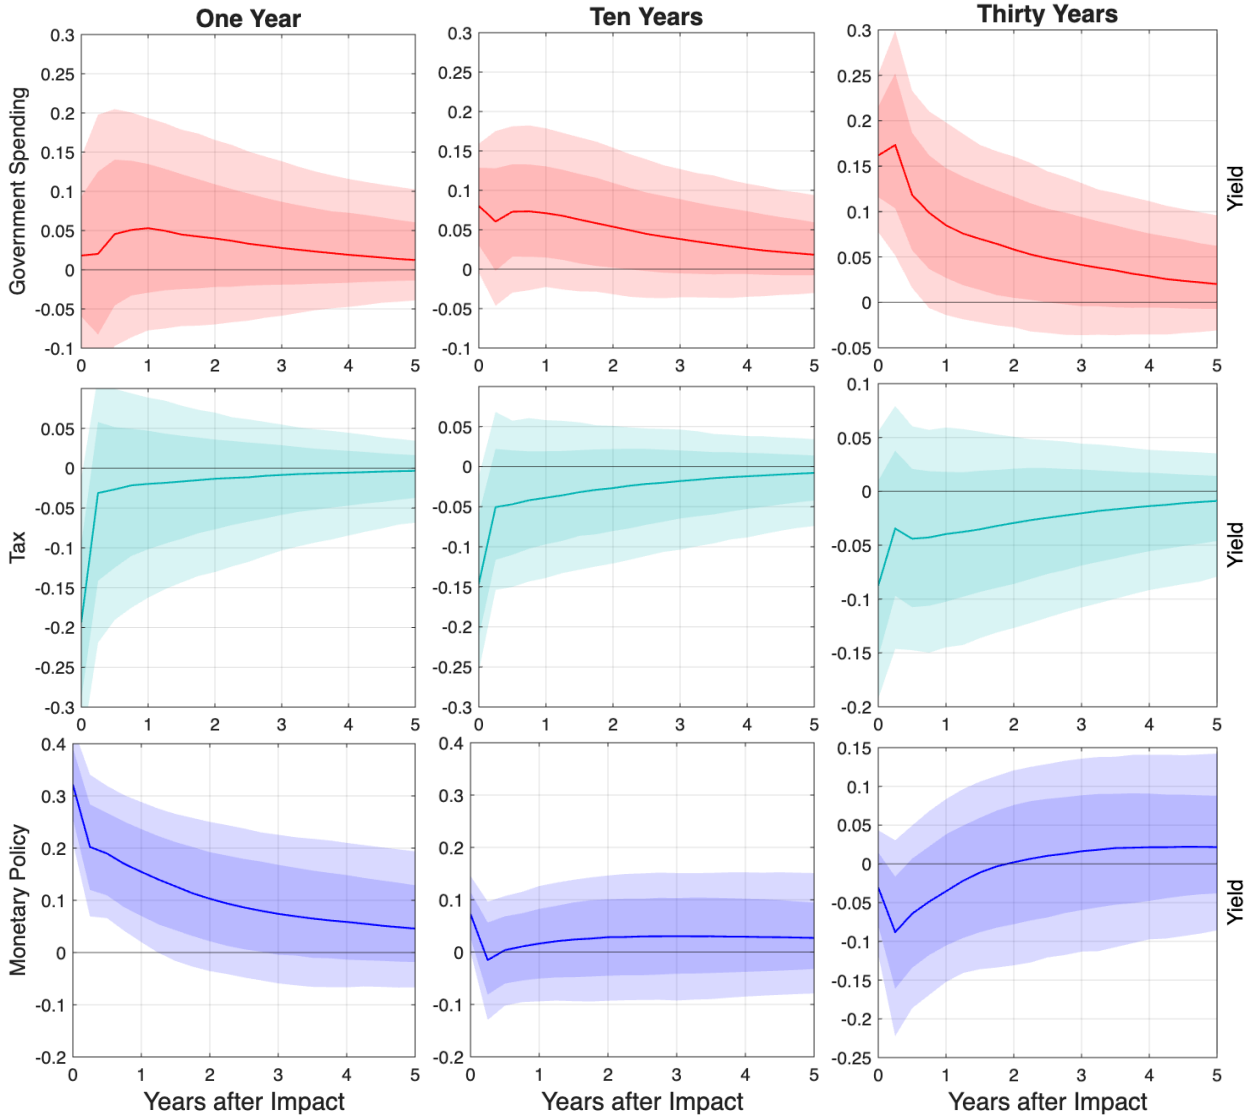

Note: The x-axis denotes periods after the shock occurred (h). The lighter (darker) shade signifies 90% (68%) confidence bands. 90% and 68% confidence bands estimated using bootstrap methods.

to 30 years, comparing government spending, tax, and monetary policy shocks. Government spending does not affect average expected short rates, but it moves the term premium significantly at the long end of the yield curve. Tax shocks instead influence the term premium for shorter maturities and the expectation component at the very short end of the yield curve. Monetary policy responses are driven mainly by shifts in expected current and future short rates, except at the longest maturities. Comparing the left and right panels of Figure 6 reveals that fiscal policy primarily affects yields through term premia, while monetary policy, unsurprisingly, operates via changes in expected short rates. This result suggests that monetary policymakers, insofar as they adjust short-term rates, do not substantially respond to these fiscal shocks over the time horizons we consider.

Figure 6: Decomposition of the responses at impact for the average short-term rate expectation (top) and the term-premium (bottom) for total government spending shock (left), the tax shock (center), and the monetary policy shock (right).

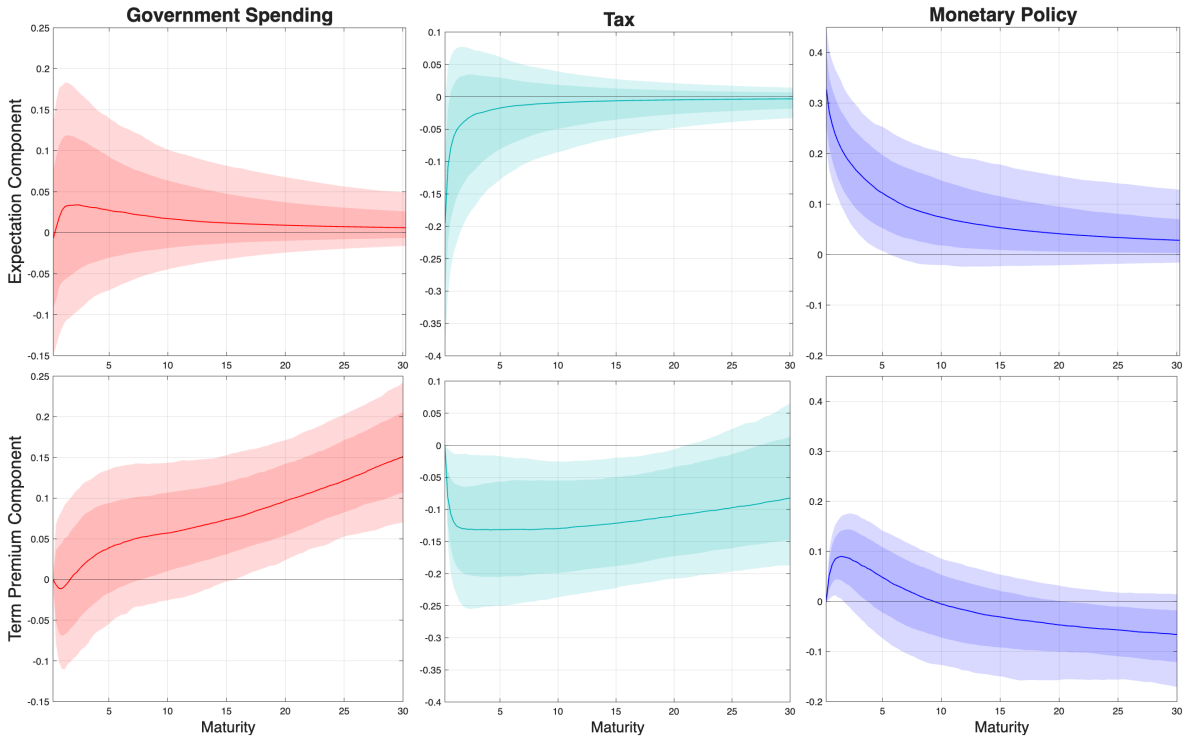

Finally, Figure 7 shows the same decomposition for various components of government spending. Interestingly, government consumption spending shocks again stand out as inducing a similar response qualitatively to monetary policy shocks - both move short rate expectations of short horizons, albeit in opposite directions. Investment and defense spending

shocks show patterns that are qualitatively similar to those of overall government spending shocks.

Figure 7: Decomposition of the responses at impact for the average short-term rate expectation (top) and the term-premium (bottom) for government defense spending (left),consumption (center) and investment (right).

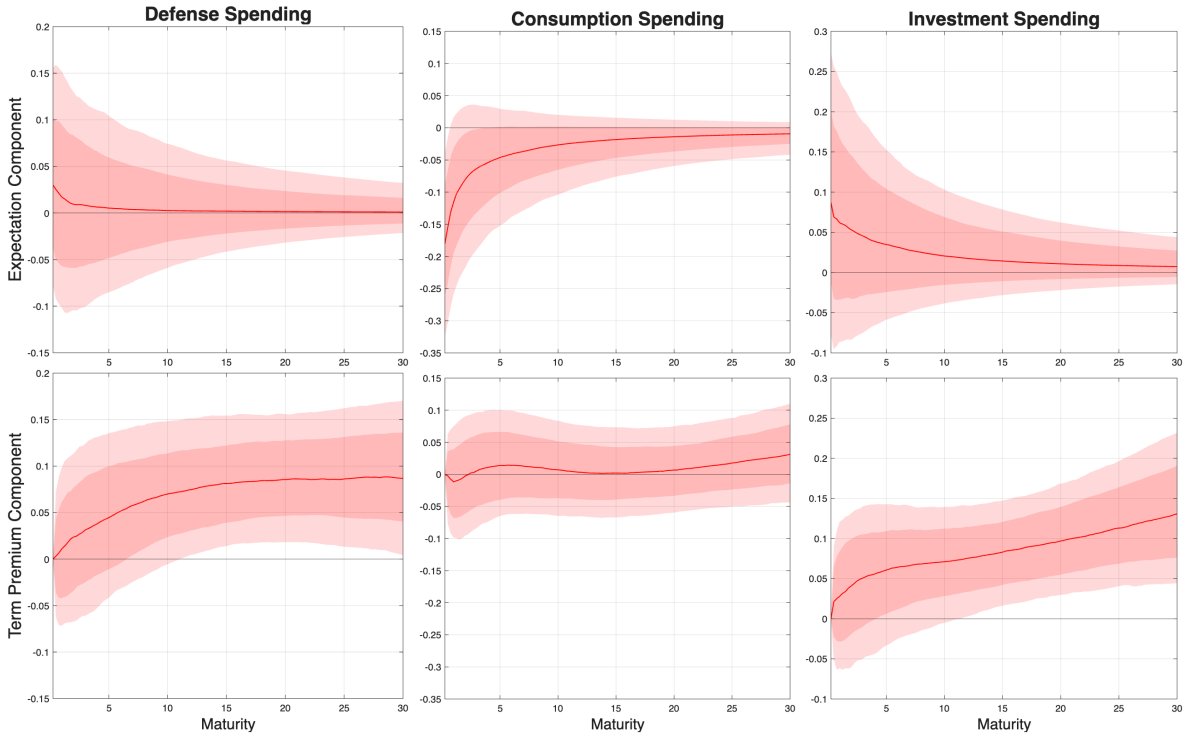

6 Conclusion

We study the effects of monetary and fiscal policies on the yield curve and find that they have qualitatively very different consequences for the yield curve - government spending affects the long end of the curve, while tax changes and monetary policy influence the short end at impact. Within spending categories, only government consumption shifts the short end, but these effects dissipate within a year. These findings are useful for both policymakers, who often view the yield curve as a major aspect of policy transmission, in addition to directly encoding a government's borrowing costs. Furthermore, our results can be useful as calibration targets for macroeconomists who want to develop quantitative equilibrium models that take the yield curve seriously.

References

Ang, A. and Piazzesi, M. (2003). A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, **50** (4), 745–787.

Angeletos, G.-M. (2002). Fiscal Policy with Noncontingent Debt and the Optimal Maturity Structure. *The Quarterly Journal of Economics*, **117** (3), 1105–1131.

Aruoba, B. and Drechsel, T. (2022). *Identifying Monetary Policy Shocks: A Natural Language Approach*. CEPR Discussion Papers 17133, C.E.P.R. Discussion Papers.

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy*, **4** (2), 1–27.

Barro, R. J. (1979). On the Determination of the Public Debt. *Journal of Political Economy*, **87** (5), 940–971.

Berndt, A., Lustig, H., and Yeltekin, S. (2012). How Does the US Government Finance Fiscal Shocks? *American Economic Journal: Macroeconomics*, **4** (1), 69–104.

Bjørnland, H. C., Chang, Y. and Cross, J. (2023). Oil and the stock market revisited: A mixed functional var approach. *Working Paper*.

Blanchard, O. and Perotti, R. (2002). An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output. *The Quarterly Journal of Economics*, **117** (4), 1329–1368.

Boehm, C. E. (2020). Government consumption and investment: Does the composition of purchases affect the multiplier? *Journal of Monetary Economics*, **115**, 80–93.

Buera, F. and Nicolini, J. P. (2004). Optimal maturity of government debt without state contingent bonds. *Journal of Monetary Economics*, **51** (3), 531–554.

Chang, M., Chen, X. and Schorfheide, F. (2021). *Heterogeneity and Aggregate Fluctuations*. NBER Working Papers 28853, National Bureau of Economic Research, Inc.

Chang, Y., Durlauf, S. N., Lee, S. and Park, J. Y. (2023a). A trajectories-based approach to measuring intergenerational mobility. *NBER Working Paper 31020*.

—, Hu, B. and Park, J. Y. (2022). Econometric analysis of functional dynamics in the presence of persistence. *Working Paper*.

—, Kim, C. S. and Park, J. Y. (2016). Nonstationarity in time series of state densities. *Journal of Econometrics*, **192** (1), 152 – 167.

—, Park, J. Y. and Pyun, D. (2023b). From functional autoregressions to vector autoregressions. *Working Paper*.

Chung, H. and Leeper, E. (2007). *What Has Financed Government Debt?* CAEPR Working Papers 2007-015, Center for Applied Economics and Policy Research, Department of Economics, Indiana University Bloomington.

Correia, I., Nicolini, J.-P. and Teles, P. (2008). Optimal fiscal and monetary policy: Equivalence results. *Journal of Political Economy*, **116** (1), 141–170.

Dai, Q. and Philippon, T. (2005). *Fiscal Policy and the Term Structure of Interest Rates*. NBER Working Papers 11574, National Bureau of Economic Research, Inc.

Diebold, F. X. and Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & economic statistics*, **20** (1), 134–144.

— and Rudebusch, G. D. (2012). *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*. No. 9895 in Economics Books, Princeton University Press.

Evans, C. L. and Marshall, D. A. (2007). Economic determinants of the nominal treasury yield curve. *Journal of Monetary Economics*, **54** (7), 1986–2003.

Folland, G. B. (1999). *Real analysis: modern techniques and their applications*, vol. 40. John Wiley & Sons.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, **7**, 44–76.

Giannitsarou, C. and Scott, A. (2008). Inflation Implications of Rising Government Debt. In *NBER International Seminar on Macroeconomics 2006*, NBER Chapters, National Bureau of Economic Research, Inc, pp. 393–442.

Gürkaynak, R. S., Sack, B. and Wright, J. H. (2007). The us treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, **54** (8), 2291–2304.

Hall, G. J. and Sargent, T. J. (2011). Interest Rate Risk and Other Determinants of Post-WWII US Government Debt/GDP Dynamics. *American Economic Journal: Macroeconomics*, **3** (3), 192–214.

Hilscher, J., Raviv, A. and Reis, R. (2014). *Inflating Away the Public Debt? An Empirical Assessment*. NBER Working Papers 20339, National Bureau of Economic Research, Inc.

Inoue, A. and Rossi, B. (2021). A new approach to measuring economic policy shocks, with an application to conventional and unconventional monetary policy. *Quantitative Economics*, **12** (4), 1085–1138.

Ireland, P. N. (2015). Monetary policy, bond risk premia, and the economy. *Journal of Monetary Economics*, **76**, 124–140.

Karantounias, A. (2020). *Doubts about the model and optimal policy*. Tech. rep., Federal Reserve Bank of Atlanta.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics*, **47**, 523–544.

Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review*, **111** (4), 1092–1125.

Lucas, R. J. and Stokey, N. L. (1983). Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics*, **12** (1), 55–93.

Mertens, K. and Ravn, M. O. (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, **103** (4), 1212–1247.

Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, **13** (3), 74–107.

Piazzesi, M. (2005). Bond Yields and the Federal Reserve. *Journal of Political Economy*, **113** (2), 311–344.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica*, **89** (2), 955–980.

Plosser, C. I. (1987). Fiscal policy and the term structure. *Journal of Monetary Economics*, **20** (2), 343–367.

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. In J. B. Taylor and H. Uhlig (eds.), *Handbook of Macroeconomics*, *Handbook of Macroeconomics*, vol. 2, *0*, Elsevier, pp. 71–162.

Ramey, V. A. (2011). Identifying Government Spending Shocks: It's all in the Timing. *The Quarterly Journal of Economics*, **126** (1), 1–50.

— and Zubairy, S. (2018). Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. *Journal of Political Economy*, **126** (2), 850–901.

Romer, C. D. and Romer, D. H. (2010). The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review*, **100** (3), 763–801.

Stock, J. and Watson, M. (2016). Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics. In J. B. Taylor and H. Uhlig (eds.), *Handbook of Macroeconomics*, *Handbook of Macroeconomics*, vol. 2, *0*, Elsevier, pp. 415–525.

Wolf, C. K. (2021). *Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result*. Working Paper 29193, National Bureau of Economic Research.

Online Appendix for "The Influence of Fiscal and Monetary Policies on the Shape of the Yield Curve"

A Additional Figures

A.1 Response of the Yield Curve to Monetary Policy Shocks Identified Using High Frequency Variation

When comparing the responses of the yield curve to monetary policy shocks, two different studies provide valuable insights. Aruoba and Drechsel (2022) and Miranda-Agrippino and Ricco (2021) utilize different shocks to analyze the transmission of monetary policy shocks. The shock identified by Aruoba and Drechsel (2022) focuses on a natural language analysis, while the shock from Miranda-Agrippino and Ricco (2021) represents a different perspective, based on high frequency estimation.

Figure A-1: Response at Impact of the Yield Curve to a Monetary Policy Shock (Miranda-Agrippino and Ricco, 2021)

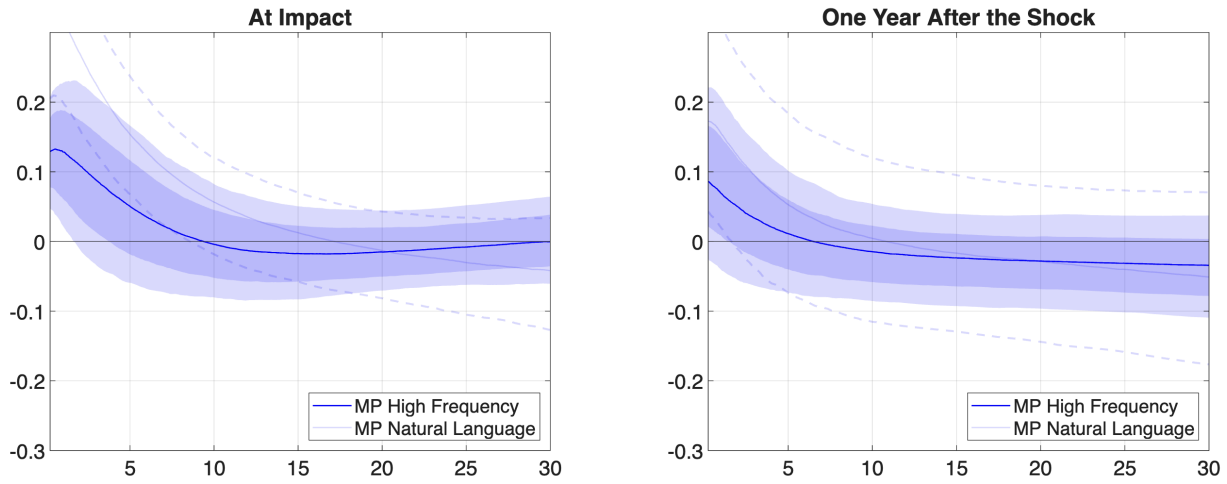

Note: For reference the response to monetary policy (Aruoba and Drechsel, 2022) and its 90% confidence band is represented in light blue. The lighter (darker) blue shade signifies the 90% (68%) confidence band for the response to high frequency monetary policy shock (Miranda-Agrippino and Ricco, 2021). The confidence bands are estimated using bootstrap methods.

A.2 Additional Results

The following figure shows the response of the yield curve to shocks to components of government spending one year after the shock. These plots confirm that the effects of fiscal policy become very weak or even negligible after just one year.

Figures A-3-A-5 present responses over horizons of up to five years for a broader set of maturities, thereby extending the analysis shown in Figure 6 of the main text.

Figure A-2: Response after one year of the yield curve to a Government Defense Spending Shock (left), Government Consumption Spending Shock (center), and, Government Investment Spending Shock (right) (Auerbach and Gorodnichenko, 2012)

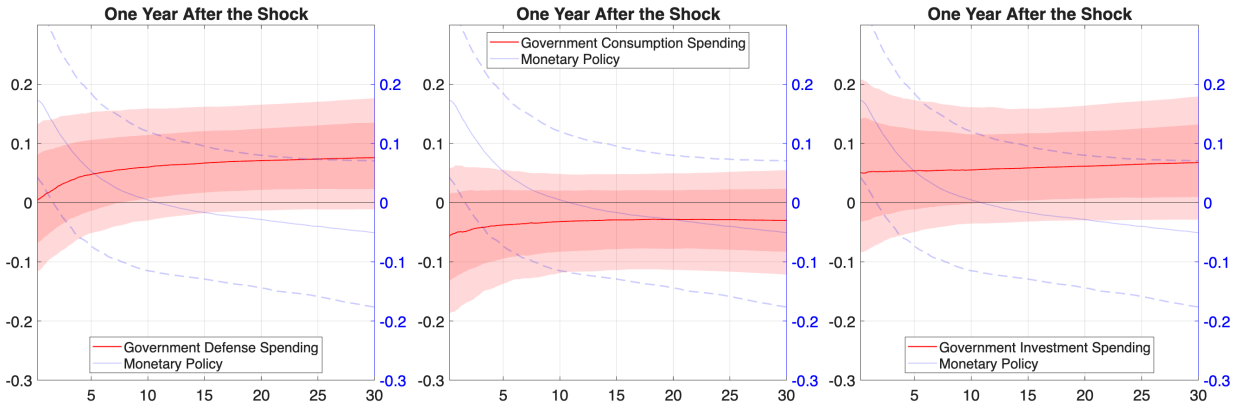

Note: For reference the response to monetary policy (Aruoba and Drechsel, 2022) and its 90% confidence band is represented in light blue. The lighter (darker) blue shade signifies the 90% (68%) confidence band for the response to different government spending shocks(Auerbach and Gorodnichenko, 2012). The confidence bands are estimated using bootstrap methods.

A.3 Robustness to Different Choices of m and p

In this section, we compare impulse responses (at impact) when choosing m and p optimally for each shock separately rather than using the benchmark values $m = 4$ and $p = 1$ that were obtained using the entire sample of the yield curve data without using a specific shock measure in the VAR.

B Data

In this section, we provide a comprehensive overview of the various data sources utilized in this paper.

B.1 Yield Curve Data

The yield curve data we use as a starting point are taken from the Federal Reserve Board[1], based on the model by Gürkaynak, Sack and Wright (2007). From this source, we obtain daily estimates of six parameters: $\beta_{0t}, \beta_{1t}, \beta_{2t}, \beta_{3t}, \tau_{1t}$, and τ_{2t}. These parameters are used to

[1] https://www.federalreserve.gov/data/nominal-yield-curve.htm

Figure A-3: Response of yields of specific maturities (given in the title of each subplot) to a government spending shock.

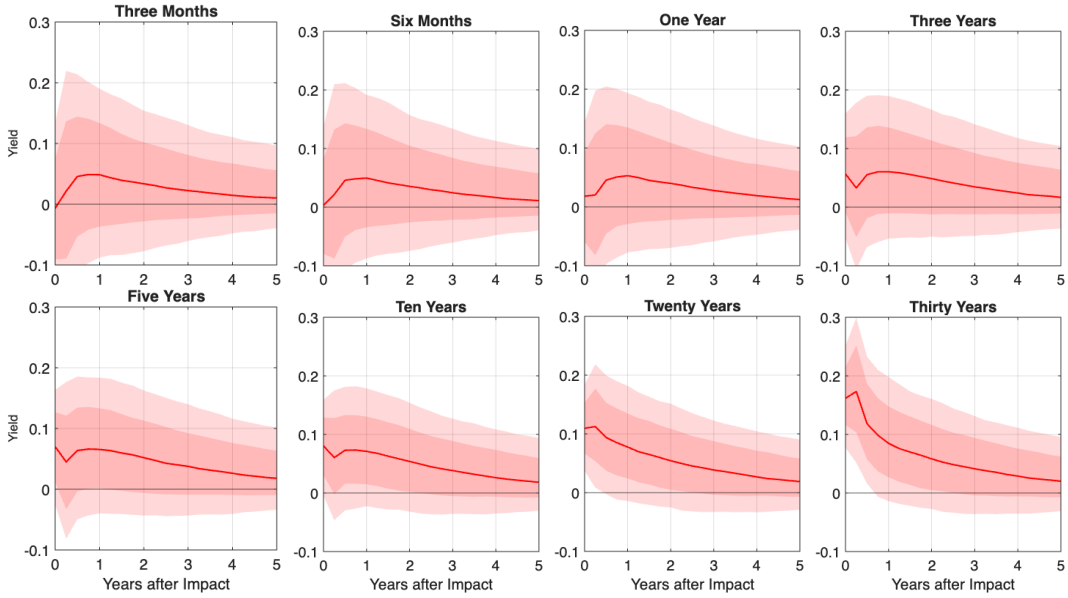

Note: The x-axis denotes periods after the shock occurred (h). The lighter (darker) shade signifies 90% (68%) confidence bands. 90% and 68% confidence bands estimated using bootstrap methods.

Figure A-4: Response of yields of specific maturities (given in the title of each subplot) to a tax shock.

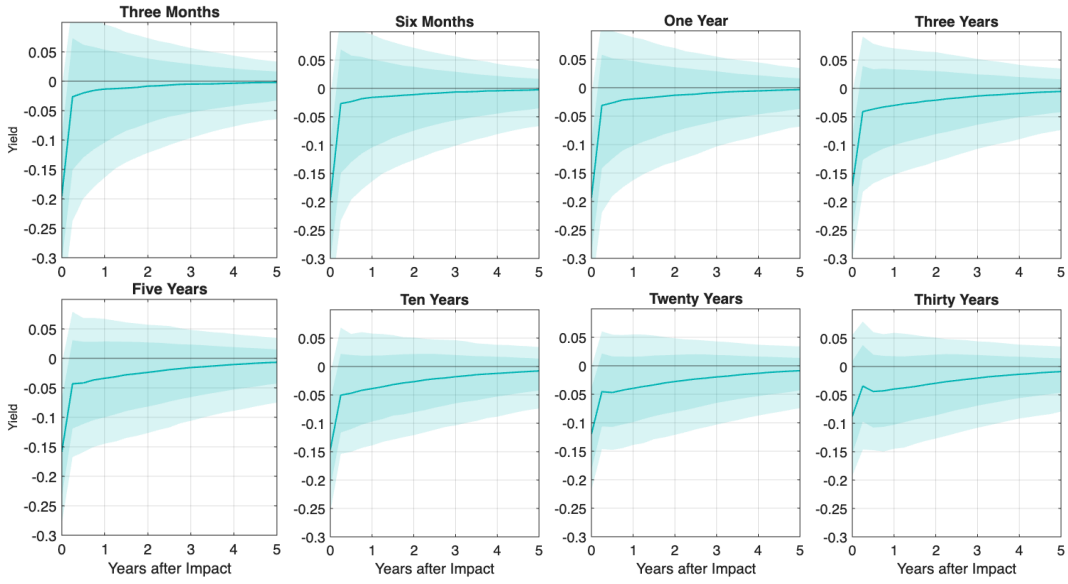

Note: The x-axis denotes periods after the shock occurred (h). The lighter (darker) shade signifies 90% (68%) confidence bands. 90% and 68% confidence bands estimated using bootstrap methods.

Figure A-5: Response of yields of specific maturities (given in the title of each subplot) to a monetary policy shock.

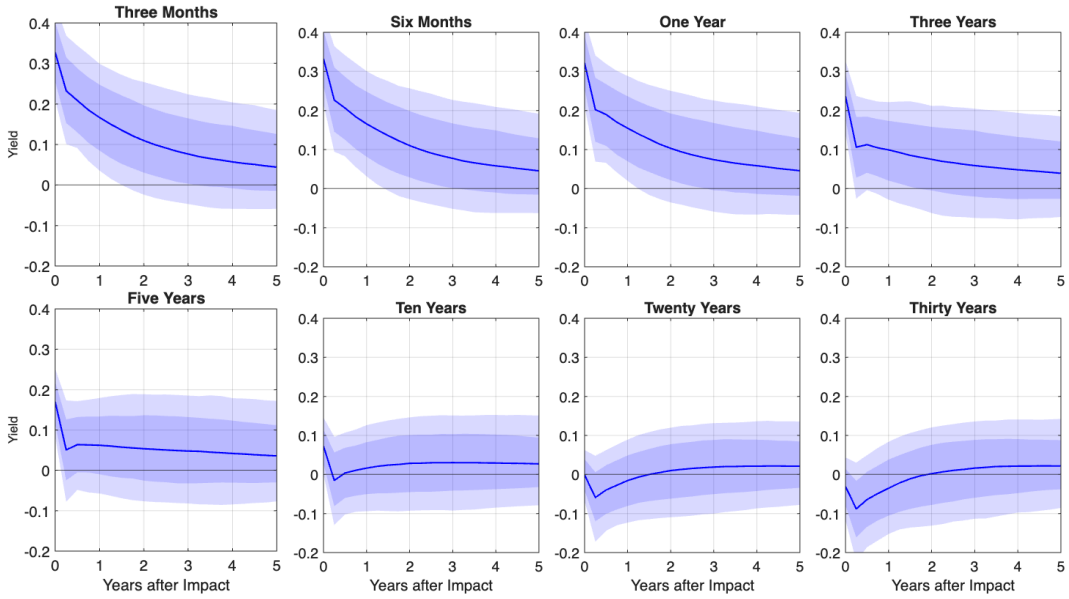

Note: The x-axis denotes periods after the shock occurred (h). The lighter (darker) shade signifies 90% (68%) confidence bands. 90% and 68% confidence bands estimated using bootstrap methods.

Figure A-6: Response (at Impact) of the Yield Curve to a Government Spending and Tax Shock.

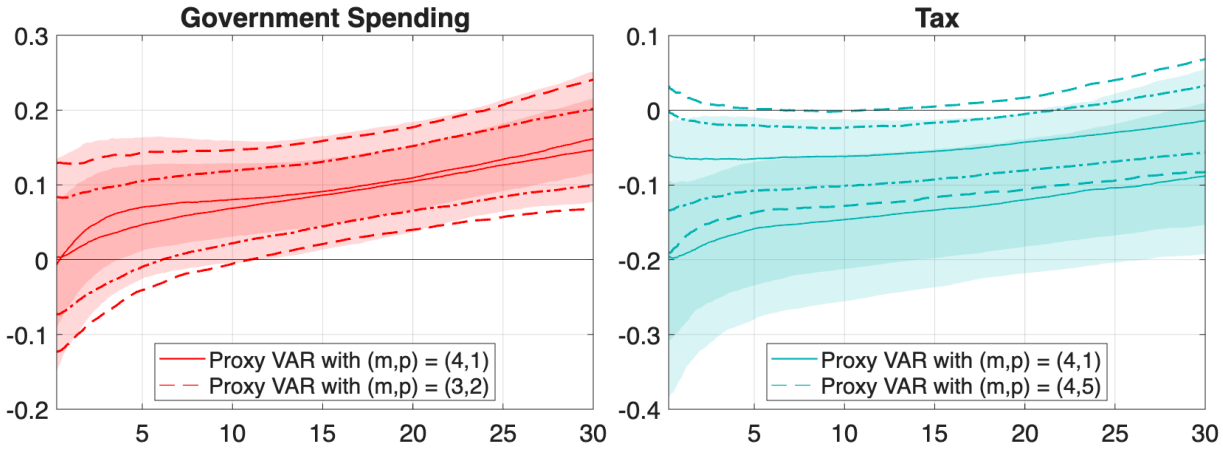

Note: This figure presents yield curve responses at impact to the government spending (Auerbach and Gorodnichenko, 2012) and tax (Romer and Romer, 2010) shocks. Results are shown for two different optimal (m,p) parameter choices: one based on the full yield curve sample, and one based on the restricted sample that matches the availability of the corresponding shock data.

construct the following yield curve function:

$$y_t(M) = \beta_{0t} + \beta_{1t}\left(\frac{1-e^{-\frac{M}{\tau_{1t}}}}{\frac{M}{\tau_{1t}}}\right) + \beta_{2t}\left(\frac{1-e^{-\frac{M}{\tau_{1t}}}}{\frac{M}{\tau_{1t}}} - e^{-\frac{M}{\tau_{1t}}}\right) + \beta_{3t}\left(\frac{1-e^{-\frac{M}{\tau_{2t}}}}{\frac{M}{\tau_{2t}}} - e^{-\frac{M}{\tau_{2t}}}\right)$$

Figure A-7: Response (at Impact) of the Yield Curve to Defense, Consumption and Investment Government Spending shocks. Comparison different general and specific values of m and p.

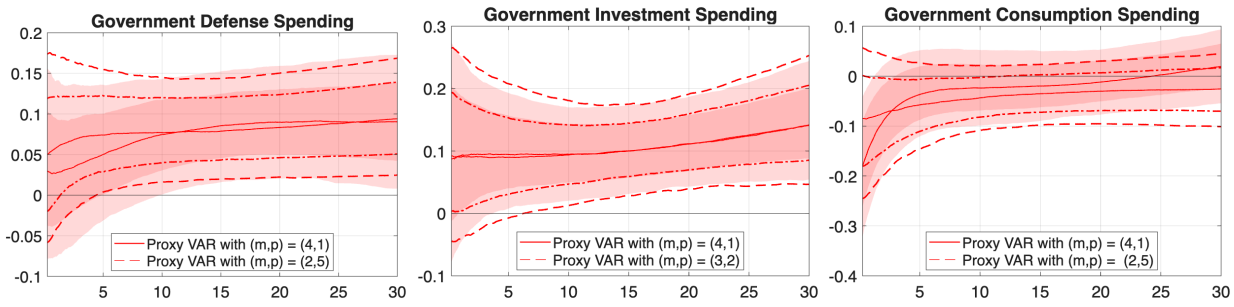

Note: This figure presents yield curve responses at impact to the government spending shocks variations as in Auerbach and Gorodnichenko (2012). Results are shown for two different optimal (m,p) parameter choices: one based on the full yield curve sample, and one based on the restricted sample that matches the availability of the shock data.

Conditional on τ_1 and τ_2, these specifications represent a linear combination of three basis functions. However, since these parameters vary over time—unlike our approach in the main text—the basis functions are not time-invariant. This means that our approach still involves a substantial dimensionality reduction relative to the original data. The daily sample of the yield curve starts on June 14, 1961, and is updated weekly.

B.2 External Shocks

Throughout the paper, we use several external shocks borrowed from the literature. We consider three main groups of shocks: monetary policy, government spending, and tax changes. In the following sections, we describe the data sources used to construct each of these external shocks.

B.2.1 Monetary Policy

The primary references for measuring changes in monetary policy are Aruoba and Drechsel (2022) and Miranda-Agrippino and Ricco (2021).

Aruoba and Drechsel (2022) This paper applies natural language processing techniques to analyze documents prepared by economists for FOMC meetings, aiming to capture the information available to the committee at the time of policy decisions. Using machine learning methods, the authors predict changes in the target interest rate based on this information and identify monetary policy shocks as the residuals. Their identified shock, cov-

ering the period from 1982Q3 to 2008Q4, is available at: http://econweb.umd.edu/ drechsel/files/Aruoba_Drechsel_Data.xlsx.

Miranda-Agrippino and Ricco (2021) This study constructs a high-frequency instrument for monetary policy shocks that accounts for informational rigidities. The series of shocks is available at: http://silviamirandaagrippino.com/s/Instruments_web-x8wr.xlsx. The sample period for this shock spans from 1991Q1 to 2009Q4.

B.2.2 Government Spending

Auerbach and Gorodnichenko (2012) We use the linear (non-regime-switching) version of the VAR model described in Auerbach and Gorodnichenko (2012) to obtain government spending shocks. The identification strategy (variable ordering), decomposition of the government spending variable, and treatment of predictable components of fiscal shocks remain consistent with the original study.

The variables included in the model are as follows: government spending, representing the log of real government (federal, state, and local) purchases (consumption and investment); government revenue, representing the log of real government receipts from direct and indirect taxes net of transfers to businesses and individuals; and output, representing the log of real gross domestic product (GDP) in chained 2000 dollars.

To account for anticipated shocks, we incorporate quarterly forecasts of fiscal and aggregate variables (government purchases, output, and taxes) from the University of Michigan's Research Seminar in Quantitative Economics (RSQE) macroeconometric model, the Survey of Professional Forecasters (SPF), and the forecasts prepared by the Federal Reserve Board (FRB) staff for FOMC meetings. These forecasts are included in the SVAR to eliminate the effects of "innovations" in fiscal variables that were predicted by professional forecasters.

The resulting shocks are obtained using the inverse of the Cholesky decomposition of the estimated covariance matrix and the estimated residuals. The sample period for these shocks spans from 1967Q3 to 2008Q4. Replication files are available at: http://doi.org/10.3886/E114783V1.

B.2.3 Tax Changes

We estimate the response of the yield curve to changes in tax legislation as described in Romer and Romer (2010).

Tax Changes The authors use narrative records, such as presidential speeches and congressional reports, to identify the size, timing, and principal motivation for all major postwar tax policy actions. Their analysis distinguishes between legislated changes motivated by prospective economic conditions and those driven by more exogenous factors. The sample period for this shock spans from 1961Q3 to 2007Q4. The shock series is available at: https://www.aeaweb.org/aer/data/june2010/20080421_app.zip.

C Functional Principal Components

The eigenvalues of a compact operator are defined as the non-zero scalars λ for which there exists a non-zero function v in the underlying vector space such that the operator S applied to v is a scalar multiple of v. More formally, for a compact operator S defined on a Banach space or a Hilbert space, the eigenvalue-eigenvector equation is given by:

$$Sv = \lambda v$$

Here, v is the eigenfunction associated with the eigenvalue λ, and λ is a scalar. Compact operators are typically defined on infinite-dimensional spaces, and their eigenvalues may include accumulation points or have a discrete or continuous spectrum depending on the properties of the operator and the underlying space. The eigenvalues of a compact operator provide important information about its spectral properties and behavior. In the case of a bounded compact operator on a Hilbert spaces, the eigenvalues are a sequence of numbers with 0 as only accumulation point.

For a time series of functions $y_1, y_2, y_3, \ldots, y_T$, we define the operator

$$S = \frac{1}{T}\sum_{i=1}^{T}(y_i - \bar{y}) \otimes (y_i - \bar{y})$$

where $\bar{y}$ is the sample mean. The operator S is the sample covariance operator of (y_t).

Let $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \cdots$ be the eigenvalues of S, and we call the associated eigenfunctions $v_1, v_2, v_3, \ldots$ the **functional principal components of** y_t.

There are multiple norms we can use for S. We use the trace norm $\|S\|$ to measure the variability of (y_t). The trace norm equals the summation of all eigenvalues of the compact

operator, viz.,

$$\|S\| = \lambda_1 + \lambda_2 + \lambda_3 + \cdots$$

A screeplot $s(m)$ measures the amount of the total variability explained by a subset of principal components:

$$s(m) = \frac{\lambda_1 + \lambda_2 + \cdots + \lambda_m}{\lambda_1 + \lambda_2 + \lambda_3 + \cdots}$$

for $m = 1, 2, \ldots$.

In Appendix D we show how to computationally obtain the functional principal components of a functional time series.

D How to Model the Yield Curve Computationally?

In the main text, we use the interval $I = [0.25, 30]$ representing maturities from three months to 30 years. We use a grid of 1024 equidistant points, from $x_1 = 0.25$ to $x_{1024} = 30$.

Computationally, the sample of the yield curve is a matrix Y of dimensions $T \times 1024$, where each row contains the values y_{ti} for $t = 1, 2, 3, \ldots, T$ and $i = 1, \ldots, 1024$, representing the yield in period t and at maturity x_i.

The main operation using the data are described as follows:

1. Sample mean yield curve: $\bar{y} = (\bar{y}_{\cdot 1}, \bar{y}_{\cdot 2}, \ldots, \bar{y}_{\cdot 1024})$. This represents the mean yield curve in the sample computed for each of the 1024 grid points.

2. The scalar product between two functions is the inner product of their vector representation.

3. The tensor product is the outer product of their vector representation.

4. The sample variance matrix S with dimensions 1024×1024, calculated as:

$$S = (Y - \bar{y})'(Y - \bar{y})$$

5. The estimated functional principal components are obtained from the eigenvalue decomposition of the matrix S. The functional principal components are given by the

eigenvectors of S, and the portion of the variance explained by each component is given by the eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \cdots$

E Random Functions

In this section, we describe some of the main concepts of random functions. Similarly to random variables and random vectors, we define a random function on a probability space $(\Omega, \mathcal{A}, P)$, which consists of a sample space Ω, an event space $\mathcal{A}$, and a probability measure P.

Expected function Given a random function f and an arbitrary element v of H, the scalar product $\langle f, v\rangle$ becomes a scalar random variable. This random variable has an expected value denoted as $\mathbb{E}\langle f, v\rangle$. The mapping

$$v \mapsto \mathbb{E}\langle f, v\rangle$$

is proven to be a linear functional from H to $\mathbb{R}$. By Riesz' representation theorem, there exists a non-random element in H referred to as "$\mathbb{E}f$", such that

$$v \mapsto \mathbb{E}\langle f, v\rangle = \langle \mathbb{E}f, v\rangle$$

In other words, we use the representation of a linear functional as the scalar product with a fixed element of H to characterize the expected function of the random function f.

Covariance operator If f and g are random functions taking values in H, then their covariance operator $\mathbb{E}(f \otimes g)$ is generally defined as a linear operator satisfying

$$\Big\langle u, [\mathbb{E}(f \otimes g)]v \Big\rangle = \mathbb{E}\langle u, f\rangle\langle v, g\rangle$$

for all u and v in H.

The combination of these two concepts allows us to define, for example, a functional white noise (ε_t) as follows: We set $\mathbb{E}\varepsilon_t = 0$ for all $t \geq 1$, and (ε_t) to be serially uncorrelated with $\mathbb{E}(\varepsilon_t \otimes \varepsilon_t) = \Sigma$ for all $t \geq 1$.

F Bootstrapping

We use the bootstrap to determine confidence intervals for the statistics we estimate regarding the yield curve and its reactions to external shocks.

In our investigation of the yield curve's response to an external shock, we consider a sample of the yield curve that corresponds to the sample of the external shock. The majority of the external shocks examined in this study occur at a quarterly frequency, so we utilize the most recent daily observation of the yield curve from the corresponding quarter. This ensures that we have two identical samples, in terms of size and frequency, for both the yield curve and the external shock.

For simplicity, the bootstrap method discussed here is illustrated using a VAR of order $p = 1$. However, this approach can be readily extended to cases with multiple lags, as suggested in the main text.

In the following, we outline the procedure for generating copies of functional time series and the external shock, each of size T:

1. Obtain the residuals from estimating model (9): $\hat{u}_t = c + \gamma_t - \hat{A}\gamma_{t-1}$.

2. Randomly select, with replacement, a sample of n residuals from the set $\{\hat{u}_1, \hat{u}_2, \ldots, \hat{u}_T\}$, and center (demean) it to obtain a new set of residuals (u_t^*).

3. Generate a new time series γ_t^* using the equation $\gamma_t^* = \hat{A}\gamma_{t-1}^* + u_t^*$ with the (u_t^*) generated in the previous step and the initial value $\gamma_0^* = \gamma_0$, the initial values of γ_t for the original sample. Note that $\hat{A}$ is the same as in the first step.

With the new copy of the time series $\{\gamma_t^*\}_{t=0,1,2,\ldots,T}$, estimate model (9) and obtain $\hat{u}_t^* = \gamma_t^* - \hat{A}^*\gamma_{t-1}^*$ for $t = 1, \ldots, T$. With the new estimation, obtain the usual $(1+m)$-dimensional impulse response functions γ_{t+h}^* of $\gamma_t^* = (z_t^*, \alpha_t^{*'})'$ as discussed in the main text. Then, use the basis functions $v_1, v_2, v_3, \ldots, v_m$ to recover the corresponding yield curve response y_{t+h}^* from the [2:m] components of the $(1+m)$-dimensional response vector γ_{t+h}^* associated with the α_{t+h}^* vector. Repeat the steps 2-3 a large number B of times (e.g., $B = 1000$). Calculate the desired confidence bands as the quantiles of the B saved estimates of y_{t+h}^*.

G Models of Yield Curve Dynamics

In this section, we evaluate how our functional time-series (FTS) specification compares with established approaches to modeling yield-curve dynamics. Since the FTS framework is used to trace responses to external monetary and fiscal shocks, it is important to verify that it also delivers competitive predictive accuracy. We assess out-of-sample performance using mean squared forecast errors (MSFEs) across multiple maturities and forecast horizons. We compare our FTS approach against two benchmark VAR models: one using yields at fixed maturities, and another using principal components extracted from conventional principal component analysis. We refer to these benchmark models as VAR and PCA, respectively. The comparison thus encompasses three specifications: the FTS model from the main text, the maturity-based VAR, and the PCA-based model.

For two models—one with the FTS specification and the other with an alternative (ALT), fix a maturity τ and forecast horizon h. For each forecast origin $t-h$, let $\hat{y}^{(i)}_{t|t-h}(\tau)$ denote the h-step-ahead forecast for period t from model $i \in \{\text{FTS}, \text{ALT}\}$. The corresponding forecast error is

$$FE^{(i)}_t(\tau, h) \;=\; y_t(\tau) - \hat{y}^{(i)}_{t|t-h}(\tau)\,.$$

Forecasts are produced with a rolling window of $W = 120$ months. To ensure comparability, we align the set of forecast origins across models (using the maximal lag order when relevant), so that both generate the same number $N(h)$ of h-step errors. For each (τ, h), the MSFE of model i is

$$\text{MSFE}^{(i)}(\tau, h) \;=\; \frac{1}{N(h)} \sum_t \Big[FE^{(i)}_t(\tau, h)\Big]^2,$$

and we report

$$\Delta\text{MSFE}(\tau, h) \equiv \text{MSFE}^{(\text{Alt})}(\tau, h) - \text{MSFE}^{(\text{FTS})}(\tau, h),$$

where positive values of ΔMSFE indicate that the FTS model yields lower forecast errors than the alternative model. To assess statistical significance, we use Diebold-Mariano tests (Diebold and Mariano, 2002) with HAC standard errors suitable for overlapping h-step forecasts.

We use the zero-coupon Treasury yield data of Gürkaynak *et al.* (2007), which provide a long monthly sample of over 700 observations. Applying a rolling 120-month window yields

several hundred forecast origins for each horizon.[2]

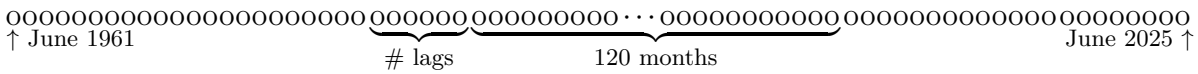

G.1 A Vector Autoregression with Multiple Maturities

A straightforward benchmark is to stack a set of maturities into a vector. Let k maturities $\tau_1 < \tau_2 < \cdots < \tau_k$, and define

$$Y_t = \Big(y_t(\tau_1),\, y_t(\tau_2),\, \ldots,\, y_t(\tau_k)\Big)'.$$

The main limitation of this approach is that it models only the selected maturities, whereas the FTS specification naturally produces forecasts for the entire curve. Nonetheless, we evaluate the out-of-sample accuracy of the FTS and VAR on equal footing, choosing the number of lags and components to minimize the MSFE separately for each specification.

We report results for multiple choices of both the number and the specific maturities included in Y_t. This collection of results provides reassurance that our comparison between the stacked-maturity VAR and the functional time-series approach is robust.

Figure A-8 reports the MSFE difference between the FTS and VAR across eight maturities (three months, one, two, five, ten, fifteen, twenty, and thirty years) and forecast horizons h (months ahead). A positive value implies an advantage for our FTS approach. The FTS gains are modest at short horizons ($h = 1$–3) but widen beyond $h = 5$, becoming both larger and statistically significant at several maturities. At long horizons ($h = 9$–12), the FTS framework consistently outperforms the VAR, underscoring its efficiency when forecasting further into the future.

Figure A-9 considers a smaller set of five maturities. Results mirror the earlier figure: little difference at short horizons (note that the y-axis range is narrower here than in the other panels), but consistent and statistically significant gains for the FTS model at longer horizons.

Figures A-10 and A-11 show results based on four selected maturities. The findings remain consistent: modest differences at short horizons, but clear FTS advantages at medium

[2]While the main text employs quarterly yield curve data due to availability constraints from the shock instruments, here we work with monthly frequency to increase the number of observed MSFE.

Figure A-8: Mean Squared Forecast Error Comparison

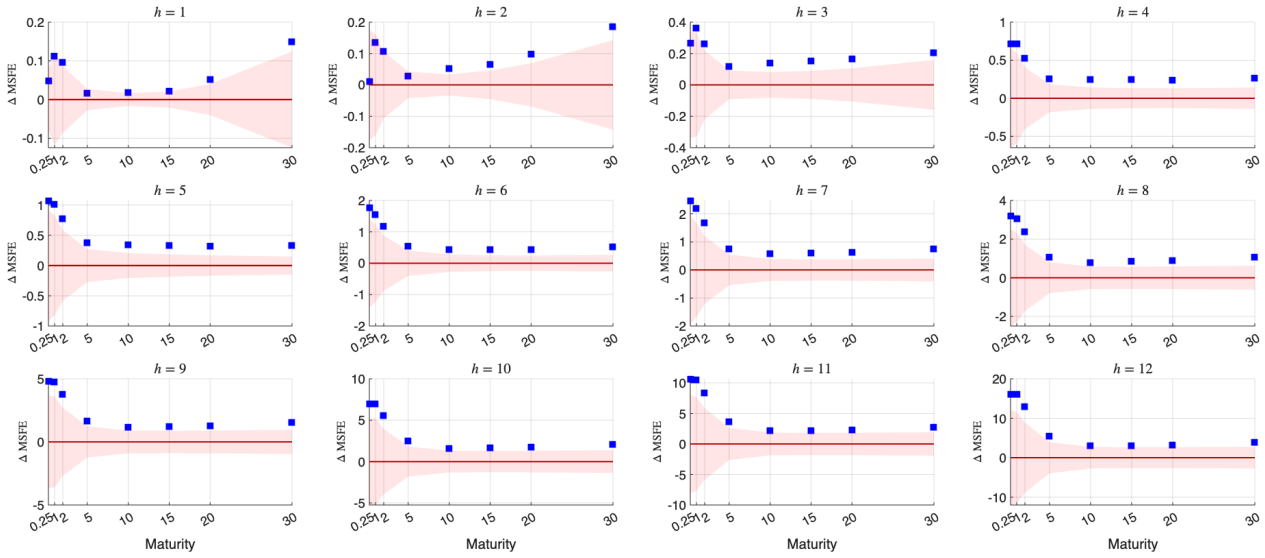

Note: Each panel reports the difference in MSFE between the FTS model and the VAR (blue squares), across eight maturities: three months, one, two, five, ten, fifteen, twenty, and thirty years. The shaded area denotes the one–standard deviation non-rejection region of the null hypothesis that the two approaches deliver equal predictive accuracy, as evaluated by the Diebold and Mariano (2002) statistic. Positive values of ΔMSFE imply that the VAR exhibits a larger MSFE than the FTS model.

Figure A-9: Mean Squared Forecast Error Difference: FTS *versus* VAR

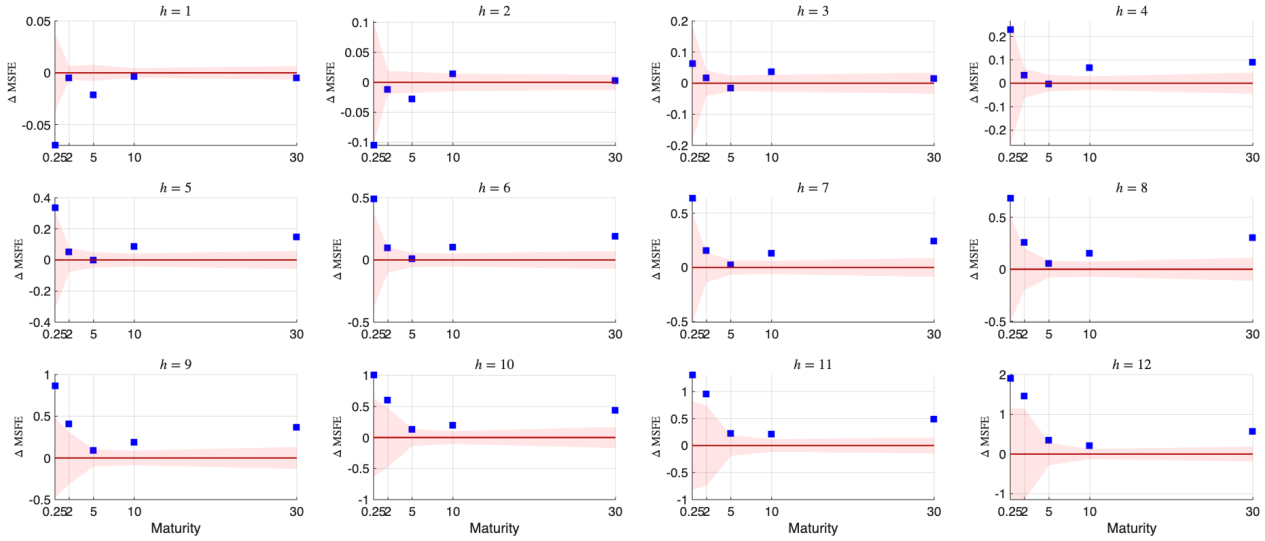

Note: Each panel reports the difference in MSFE between the FTS model and the VAR (blue squares), across five maturities: three months, two, five, ten, and thirty years. The shaded area denotes the one–standard deviation non-rejection region of the null hypothesis that the two approaches deliver equal predictive accuracy, as evaluated by the Diebold and Mariano (2002) statistic. Positive values of ΔMSFE imply that the VAR exhibits a larger MSFE than the FTS model.

Figure A-10: Mean Squared Forecast Error Difference: FTS *versus* VAR

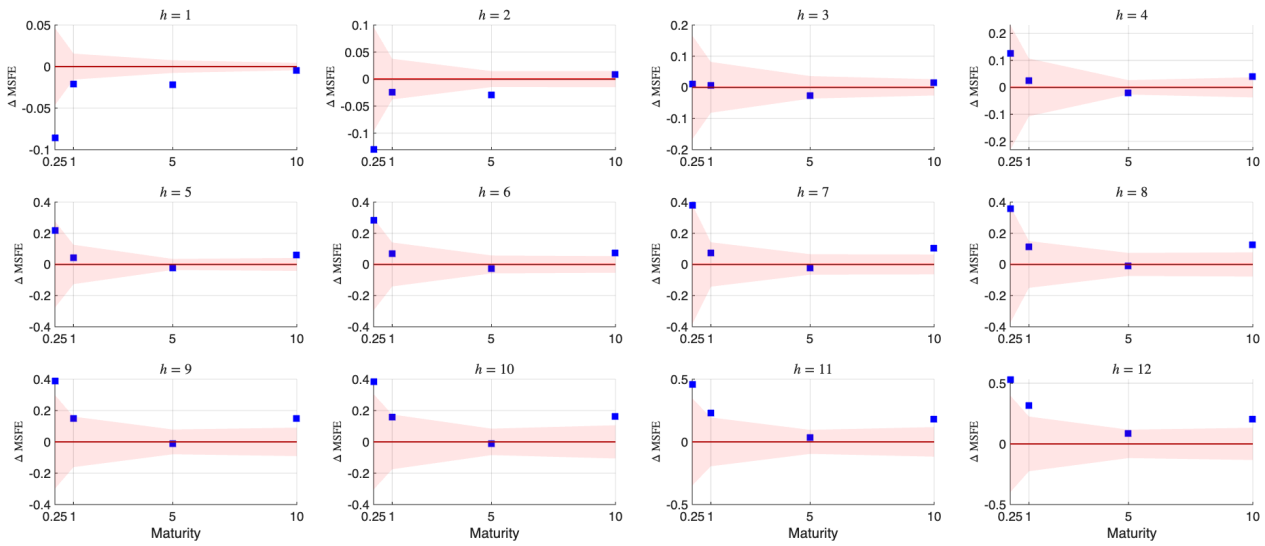

Note: Each panel reports the difference in MSFE between the FTS model and the VAR (blue squares), across four maturities: three months, one, five, and ten years. The shaded area denotes the one–standard deviation non-rejection region of the null hypothesis that the two approaches deliver equal predictive accuracy, as evaluated by the Diebold and Mariano (2002) statistic. Positive values of ΔMSFE imply that the VAR exhibits a larger MSFE than the FTS model.

and long horizons, with statistical significance in many cases. These two cases are particularly noteworthy because the VAR dimension equals the number of functional principal components used in the main text.

Figures A-12 extend the robustness analysis. The broader grid (three and six months, one, three, five, ten, fifteen, twenty, and thirty years) makes the advantage of the functional approach especially visible, as it scales more effectively with the dimensionality of the maturity space.

Summary. Across Figures A-8 to A-12, the FTS model systematically outperforms the VAR approach, except for very short horizons, where the difference in MSFEs is, however, not large. The superiority is not driven by the particular choice of maturities but reflects the inherent advantage of modeling the entire yield curve as a functional object. Gains are modest at short horizons but become substantial and statistically significant at medium and long horizons, confirming the efficiency of the functional specification.

G.2 Principal Components Vector

A second benchmark is principal component analysis (PCA), which summarizes a set of yields using a small number of common factors. If implemented on a dense grid of maturities

Figure A-11: Mean Squared Forecast Error Difference: FTS *versus* VAR

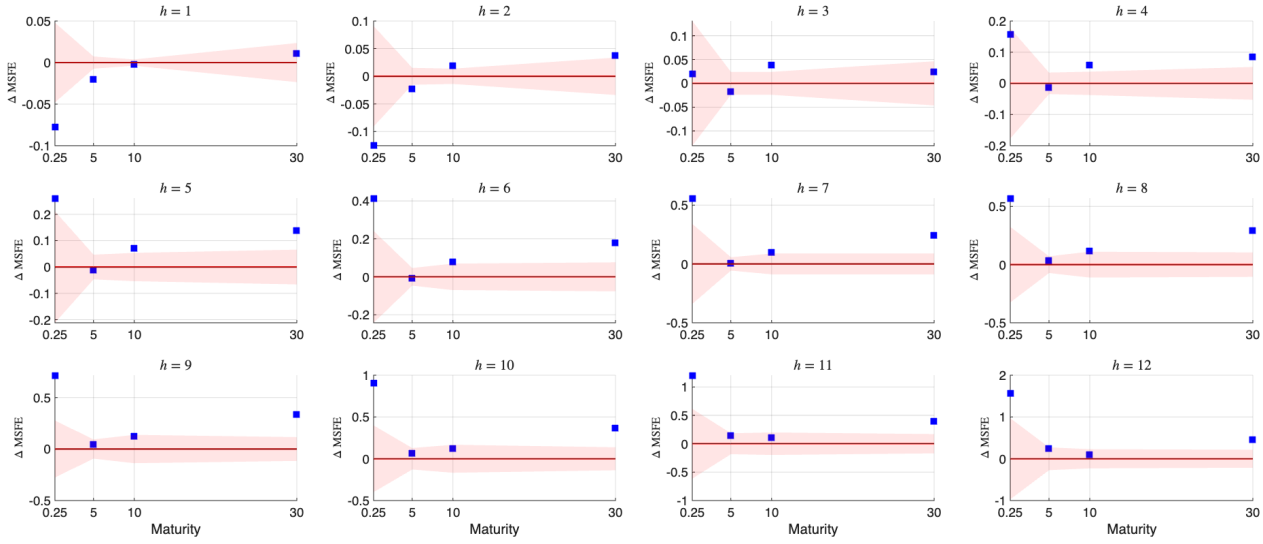

Note: Each panel reports the difference in MSFE between the FTS model and the VAR (blue squares), across four maturities: three months, five, ten, and thirty years. The shaded area denotes the one–standard deviation non-rejection region of the null hypothesis that the two approaches deliver equal predictive accuracy, as evaluated by the Diebold and Mariano (2002) statistic. Positive values of ΔMSFE imply that the VAR exhibits a larger MSFE than the FTS model.

Figure A-12: Mean Squared Forecast Error Difference: FTS *versus* VAR

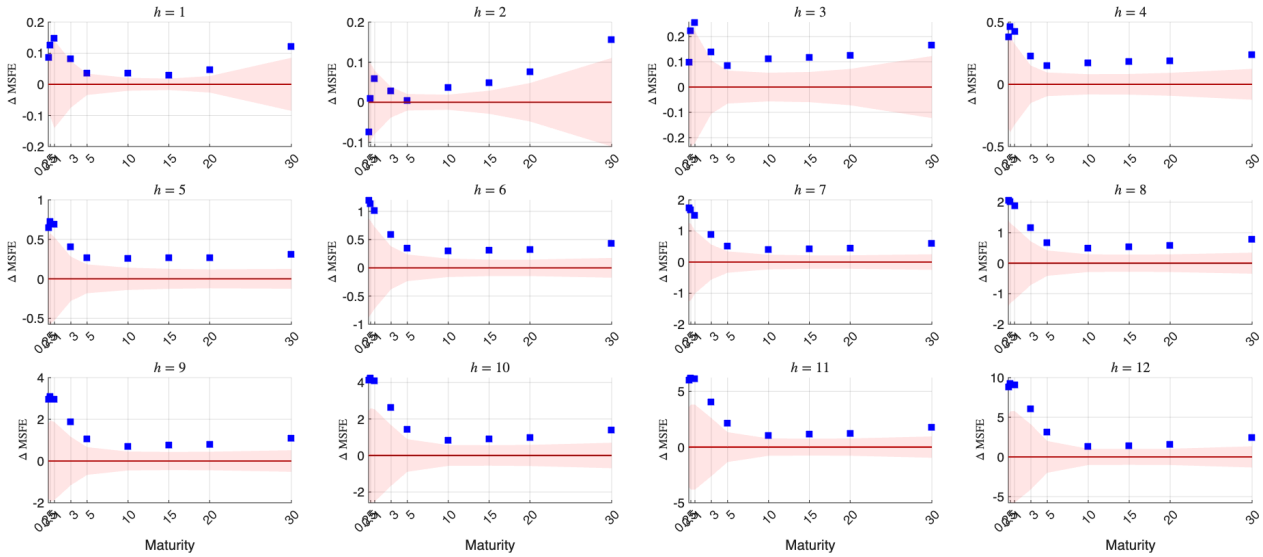

Note: Each panel reports the difference in MSFE between the FTS model and the VAR (blue squares), across nine maturities: three and six months, one, three, five, ten, fifteen, twenty, and thirty years. The shaded area denotes the one–standard deviation non-rejection region of the null hypothesis that the two approaches deliver equal predictive accuracy, as evaluated by the Diebold and Mariano (2002) statistic. Positive values of ΔMSFE imply that the VAR exhibits a larger MSFE than the FTS model.

(e.g., 1024 points between 0.25 and 30 years), PCA and the FTS approach would coincide. In practice, however, PCA is typically applied to coarser sets of yields, often at monthly intervals. To provide a fair comparison, we construct a PCA-based benchmark using yields at monthly maturities from 3 to 360 months, which corresponds to the support of the functions used in the main text: $I = [0.25, 30]$.

Figure A-13: Mean Square Forecast Error Difference: FTS *versus* PCA

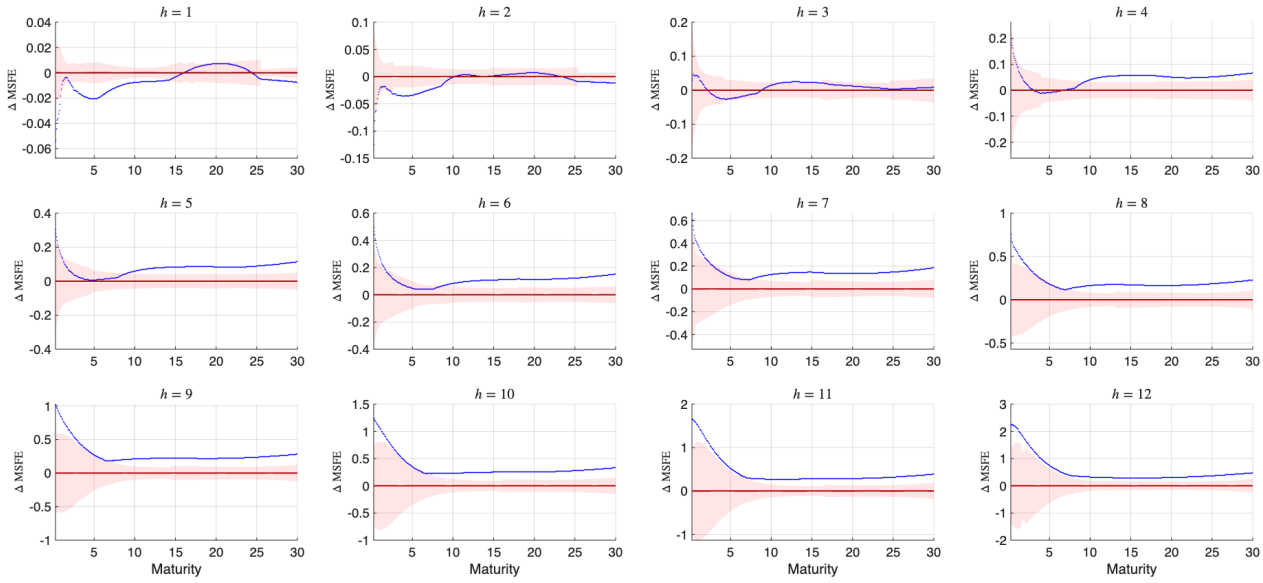

Note: Each panel shows the difference in MSFE between the FTS model and the PCA approach (blue line). We consider monthly maturities from 3 months to 30 years. The shaded area represents the one–standard deviation non-rejection region based on the Diebold and Mariano (2002) statistic. Positive values of ΔMSFE imply that the VAR exhibits a larger MSFE than the FTS model.

Figure A-13 reports the MSFE differences between the FTS model and this PCA benchmark. At short horizons, both approaches perform similarly, with differences generally inside the confidence bands. At longer horizons, however, the FTS model consistently delivers smaller forecast errors, with improvements that grow in size and significance across several maturities. While PCA is a useful dimension-reduction tool, these results show that the functional framework better captures the smooth dynamics of the yield curve, yielding more accurate forecasts at medium and long horizons.

These results confirm that the FTS specification provides forecasts that are, at the very least, competitive with these competitors and often superior, especially once we move past very short forecast horizons. Its strength lies in treating the yield curve as a continuous functional object, rather than a finite set of yields or factors, allowing it to exploit more of the information embedded in the term structure.