Bond pricing with a time-varying price of risk in an estimated medium-scale Bayesian DSGE model*

Ian Dew-Becker

Duke University

September 12, 2013

Abstract

A New-Keynesian model in which households have Epstein–Zin preferences with time-varying risk aversion and the central bank has a time-varying inflation target can match the dynamics of nominal bond prices in the US economy well. The model generates a large steady-state term spread and its fitting errors for bond yields are comparable to those obtained from a non-structural three-factor model, and one third smaller than in models with a constant inflation target or risk aversion. Including data on interest rates has large effects on variance decompositions, making investment technology shocks much less important than found in other recent papers.

1 Introduction

Non-structural and atheoretical models are widely used in both macroeconomics and the study of the term structure of interest rates. Recently, Smets and Wouters (2003) have shown that a structural New Keynesian model can match the dynamics of the macroeconomy as well as or better than a benchmark non-structural VAR. This paper extends that work by showing that a suitably augmented version of their model can also match the dynamics of the term structure of interest rates nearly as well as a standard non-structural model. In addition, including information from the term structure has substantial effects on the estimated sources of variation in the real economy.

*I appreciate helpful comments from Jason Beeler, John Campbell, Emmanuel Farhi, Andrea Tambalotti, and seminar participants at the Federal Reserve Bank of Chicago.

Whereas recent estimates of business cycle models have found shocks to investment productivity to be a dominant source of variation (Justiniano, Primiceri, and Tambalotti, 2010), I find that including data on long-term interest rates substantially weakens that conclusion. This paper thus not only shows that properly constructed DSGE models can match the behavior of asset prices, but that in fact data on asset prices is key in drawing proper conclusions about the behavior of the economy from these models.

Rudebusch and Swanson (2008) show that standard DSGE models are unable to generate the upward-sloping nominal term structure that we observe empirically. In a subsequent paper, Rudebusch and Swanson (2012) show that a calibrated model augmented with a time-varying inflation target (similar to Bekaert, Cho, and Moreno, 2010) can generate a realistically large term premium.[1] This paper extends that work in two directions. First, I fully estimate a medium-scale DSGE model, showing that it not only matches the steady-state properties of bond markets, but also fits observed dynamics. Second, I allow for time-variation in the term premium, helping match findings from the bond pricing literature that returns on long term bonds are predictable (e.g. Campbell and Shiller, 1991; Cochrane and Piazzesi, 2005).

The production side of the model I analyze is similar to other recent medium-scale DSGE models (e.g. Smets and Wouters, 2003). The central feature of the model estimated in this paper that differentiates it from the remainder of the literature is that it explicitly allows for first-order variation in risk premia over time driven by time-varying risk aversion as in Melino and Yang (2003) and Dew-Becker (2011a). While shifts in risk premia coming from various sources have been extensively studied in calibrated models, this paper is novel for estimating a full model of the economy with time-varying risk premia (i.e. time-varying expected excess returns on risky assets).[2] Shifts in risk premia are a central feature of asset markets (Cochrane, 2011), and there is strong reduced-form evidence that shifts in risk premia are important drivers of the business cycle (Gilchrist and Zakrajsek, 2012). This paper represents a first step towards estimating a full model of the economy that allows for time-varying risk premia.

Bekaert, Cho, and Moreno (2010) show that a log-linearized macro model naturally also delivers

[1] See also van Binsbergen et al. (2012).

[2] For calibrated models with time-varying risk premia, see, for example, Campanale, Castro, and Clementi (2010); Dew-Becker (2012); Gourio (2012, 2013); Guvenen (2009); Melino and Yang (2003); and Rudebusch and Swanson (2012), among others.

closed-form expressions for bond prices. Their approximation method, however, is not able to describe risk premia, and even if it could, the model assumes that risk premia are constant (because households have power utility and the fundamental shocks have constant variances). This paper builds on their work by using an approximation method that allows for positive and time-varying risk premia. In fact, I show that Epstein–Zin preferences with time-varying risk aversion naturally generate the essentially affine stochastic discount factor of Duffee (2002) that is widely studied in the bond pricing literature. With that result, the state variables of the economy follow a VAR and all bond prices are a linear function of the states, even when risk premia vary over time. The model is then naturally estimated using Kalman filter and Bayesian methods.

The estimated model fits interest rates with errors that are similar to those generated by a non-structural three-factor model. The errors in fitting annualized yields on bonds with maturities ranging from 1 quarter to 20 years have a standard deviation of 17 basis points, compared to a simple non-structural model that has fitting errors of 6–18 basis points. The steady-state term spread in the model represents the average risk premium on long-term bonds. It is estimated to be 191 basis points, similar in magnitude to the 207-basis-point average observed in the sample.

To understand why that risk premium would be large, we first need to understand what drives the variance of the pricing kernel. When the representative household has Epstein–Zin preferences with a coefficient of relative risk aversion that is substantially larger than the inverse of its EIS (preferring an early resolution of uncertainty), state prices are almost entirely driven by innovations to the household's lifetime utility, i.e. the value placed on its entire future stream of consumption and leisure. With a high EIS, transitory changes in consumption have a small effect on lifetime utility. Permanent technology shocks, though, will have large effects. Shifts in risk aversion also affect lifetime utility because they affect how much the household penalizes future uncertainty.

Even though there are nine shocks in the economy, only two of them turn out to be relevant for the pricing kernel—labor-neutral technology and risk aversion. Since all of the other shocks (e.g. monetary policy, markups, government spending) are purely transitory, they have trivial effects on permanent income and welfare, and thus they do not have a strong effect on state prices.

Following a positive innovation to the level of technology, nominal interest rates are estimated to fall, making long-term bonds risky and inducing a positive slope in the term structure. This result is common to a variety of New-Keynesian models, e.g. JPT, Smets and Wouters (2003), and Christiano, Trabandt, and Walentin (2011). In this paper, beyond the usual New-Keynesian effect that works in many papers, the central bank's inflation target also falls following positive technology shocks. Intuitively, a positive supply shock lowers inflationary pressure, which the central bank takes as an opportunity to drive inflation lower for an extended period. The fact that the negative correlation between technology shocks and interest rates is obtained in numerous other models that assume a constant inflation target suggests that this is in fact a robust feature of the data. The effect is compounded here by the shifts in the inflation target, which I find are necessary for being able to obtain a realistically large average term spread.

Variation in risk aversion also makes an important contribution to the model's ability to the term structure of interest rates, though. Standard statistical tests easily reject a model with constant risk aversion in favor of one with time-varying risk aversion. The pricing errors for bonds are smaller by a third when risk aversion is allowed to vary over time. Movements in risk aversion account for a large fraction of the variance of the term spread, particularly outside of recessions.

While the variance decompositions imply that the pricing kernel is driven entirely by the labor-neutral technology and risk aversion shocks, I find that those two shocks have only minor effects on the dynamics of the real economy in the short-run. Risk aversion explains less than 5 percent and the neutral technology shock less than 10 percent of the variance of output, consumption, investment, and hours worked at business-cycle frequencies. The variance decompositions also differ substantially from the results found by JPT. Whereas JPT find that investment technology shocks are an important driver of the business cycle, I find that they explain only 20 percent of the variance of investment growth and even less of the variance of other variables.

When bond prices are excluded from the estimation, the investment shocks are estimated to be much larger, but they have very large effects on long-term interest rates. Long-term bond prices encode information about expectations, and implicitly impulse response functions. Since short-term interest rates are a key driving force in business cycle models, it is not surprising that adding information that helps pin down expectations can have large and important effects on inferences.

In addition to matching the behavior of the term structure, the estimated parameters imply reasonable behavior for equity prices. The steady-state annualized Hansen–Jagannathan bound is estimated to be 0.55, which is consistent with the observed Sharpe ratio for the stock market in the data sample, even though data on equity returns is not included in the estimation. Furthermore, the estimated degree of variation in risk aversion is similar to (though somewhat higher than) the value used in Dew-Becker (2012a), who calibrates a general-equilibrium model that can match the both the average Sharpe ratio on equities and also empirical stock return forecasting regressions.

This paper is related to a small but growing literature on bond pricing in production economies. Bekaert, Cho, and Moreno (2010) and Doh (2011) estimate New-Keynesian macro models, but they do not focus on the size and volatility of the term premium, whereas that is the feature of the term structure that this paper concentrates on. Andreasen (2012) estimates a model of the UK economy with a focus on term premia, but with roughly constant risk aversion and fixed volatility, which makes it difficult to generate the large movements in risk premia for both bonds and other assets that are generated here. Rudebusch and Swanson (2012) generate a large and volatile term premium in a calibrated model. This paper moves beyond them by considering a substantially more complex model and showing that it can be dynamically estimated through standard Bayesian methods using the Kalman filter.

The remainder of the paper is organized as follows. Section 2 describes household preferences and derives the pricing kernel. Section 3 describes the remainder of the economy including the production process, price setting, and monetary and fiscal policy.

Next, section 4 explains how the model is solved. If we used perturbation methods, a third-order approximation would be necessary to capture time-variation in risk premia. The estimation of the model turns out to be sufficiently difficult, however (due to numerous local extrema in the likelihood function, a common feature of models of the term structure), that the use of a nonlinear filter for calculating the model's marginal likelihood is infeasible. I therefore use the essentially affine solution method described in Dew-Becker (2011b). The method approximates the pricing kernel separately from the remainder of the model, allowing it to take the essentially affine form with a time-varying price of risk described in Duffee (2002). The essentially affine method is equivalent to a first-order perturbation local to the non-stochastic steady-state, but it includes corrections for volatility that allow it to substantially outperform first-order perturbation in stochastic simulations. The key feature of the essentially affine method is that risk premia may vary over time and affect real variables, not just asset prices.

Section 5 describes the Bayesian methods used to estimate the model. Sections 6 and 7 examine the implications of the estimates for asset prices and the dynamics of the real economy, respectively.

Finally, section 8 concludes.

2 Household preferences

2.1 Objective function and budget constraint

I assume the household has recursive preferences over consumption and leisure

$$V_t = \left\{ (1-\beta) U \left(D_t, C_t, \bar{C}_{t-1}, N_t, Z_t \right) + \beta \left(E_t V_{t+1}^{1-\alpha_t} \right)^{\frac{1-\rho}{1-\alpha_t}} \right\}^{\frac{1}{1-\rho}} \tag{1}$$

where D_t is the household's cash holdings C_t is consumption, $\bar{C}_t$ is aggregate consumption, N_t is the number of hours worked outside the home, and E_t denotes the expectation operator conditional on information available at date t. The term $\bar{C}_{t-1}$ allows the period utility function to potentially include external habit formation. The level of technology, Z_t, may also affect household utility in order to ensure balanced growth (as in Rudebusch and Swanson, 2010).

The household's coefficient of relative risk aversion, α_t, is allowed to vary over time. Dew-Becker (2011a) motivates variation in α_t by considering adding a time-varying benchmark to the standard Epstein–Zin certainty equivalent, $E_t \left(V_{t+1} - H_t \right)^{1-\alpha}$. When V_{t+1} is close to H_t, the household's effective risk aversion over shocks to V_{t+1} rises. The formulation (1) has the advantage that it is log-linear and we do not have to worry about the possibility that V_{t+1} falls below H_t. In Dew-Becker (2011a), movements in α_t are connected to movements in the household's welfare. I loosen that constraint here and allow for independent shocks to risk aversion (equivalently, shocks to the habit). Melino and Yang (2003) study a similar specification, but without an explicit habit. Unlike an intertemporal preference shock, since α_t directly affects the level of welfare, shocks to α_t will be per se priced – that is, even if they have no effect on consumption or leisure, they will still affect the pricing kernel through their impact on the level of welfare.

The household's budget constraint is

$$P_t C_t + P_t I_t + H_t + D_t = (1 + i_t) H_{t-1} + W_t N_t + \Pi_t + R_{k,t} u_t K_{t-1} - P_t a (u_t) K_{t-1} + D_{t-1} \tag{2}$$

where P_t is the price of the consumption good, I_t is the expenditure on physical investment, H_t is holdings of one-period nominal bonds, i_t is the nominally riskless interest rate, W_t is the wage, and Π_t represents profits and other lump-sum transfers paid to the household. $R_{k,t}$ is the rental rate on capital and K_{t-1} the quantity of capital the household owns. The dynamics of investment and capital accumulation will be discussed in more detail below. For now it is sufficient to simply note that the household rents out its labor and capital and allocates the proceeds between consumption and saving.

I study the so-called cashless economy described in Woodford (2003). The monetary authority is able to control the interest rate because money enters the household's utility function, but the effect of money on total utility is sufficiently small that we can ignore it when writing V_t (i.e. we take the limit where the relative importance of money goes to zero). I do not discuss money any further and from now on drop D_t from the household's budget constraint and utility function.

The period utility function, $U\left(C_t, \bar{C}_{t-1}, N_t, Z_t\right)$ is motivated as a reduced form of a model of household production as in Rudebusch and Swanson (2010). Suppose households have power utility over both market goods and goods produced at home,

$$U_t = \frac{\left(C_t^{\eta} \bar{C}_{t-1}^{1-\eta}\right)^{1-\rho}}{1-\rho} + \varphi_1 \frac{C_{H,t}^{1-\rho}}{1-\rho} \tag{3}$$

where $C_{H,t}$ is consumption of the home good. Households do not derive utility directly from leisure, but rather from what they are able to produce in their non-market-work time (as in Campbell and Ludvigson, 2001). The home production function is $Z_t N_{H,t}^{\alpha_H}$, for hours worked at home $N_{H,t}$ and a coefficient $0 < \alpha_H < 1$. The level of labor-neutral technology in the economy is assumed to be equal (up to a constant of proportionality) in the home and market production sectors.[3]

The period utility function can then be written as

$$U_t \equiv \frac{\left(C_t^{\eta} \bar{C}_{t-1}^{1-\eta}\right)^{1-\rho}}{1-\rho} + Z_t^{1-\rho} \varphi_1 \frac{\left(\bar{H} - N_t\right)^{\alpha_H(1-\rho)}}{1-\rho} \tag{4}$$

[3]Note that in the household sector, an exogenous shift in Z_t, all else equal, raises output one-for-one, whereas below we will see that in the market sector it will raise output less than proportionally. The reason is that in the market sector, an increase in Z_t also leads to an identical increase in the size of the capital stock. So, ultimately, the marginal product of labor in both sectors is proportional to Z_t.

One way to rationalize this slight elision would be if the household accumulates durable goods at home that aid household production. That feature of the model is left out for simplicity.

$\bar{H}$ denotes the maximum number of hours that the household can work, either at home or in the market, and N_t is market labor. If sleep is part of home production, then $\bar{H}$ might equal 8760 hours for annual data. More generally, though, $\bar{H}$ might be smaller. As a practical matter, $\bar{H}$ affects both the elasticity of utility with respect to market labor and the Frisch elasticity. The three parameters φ_1, $\bar{H}$, and α_H jointly determine three primary features of household behavior: hours worked, the Frisch elasticity, and the elasticity of utility with respect to market labor.

The first term in (4) gives the utility that comes from consumption. The household has power utility over a Cobb–Douglas aggregate of current and (aggregate) past consumption. This formulation differs from the standard recent implementation in the macro literature in that I assume a multiplicative instead of additive habit. Campbell and Cochrane (1999) show that an additive habit can induce time-varying risk aversion, whereas the multiplicative habit will have no affect on risk aversion; the multiplicative habit here ensures that variation in risk preferences is driven purely by α_t.[4]

The key feature of the usual additive habit is simply that the marginal utility of current consumption is increasing in last period's consumption, which induces consumers to try to smooth consumption growth, as observed in the data. To obtain that result in this setting (assuming $0 < \eta < 1$), we need $\rho < 1$.

2.2 The stochastic discount factor

The marginal rate of substitution of consumption between neighboring dates is

$$\Lambda_{t+1} \equiv \frac{\partial V_t / \partial C_{t+1}}{\partial V_t / \partial C_t} = \beta \frac{U_{C,t+1}}{U_{C,t}} \frac{V_{t+1}^{\rho-\alpha_t}}{\left(E_t V_{t+1}^{1-\alpha_t}\right)^{\frac{\rho-\alpha_t}{1-\alpha_t}}} \tag{5}$$

where $U_{C,t} \equiv \partial U_t / \partial C_t$ is the marginal (period) utility of consumption. M_{t+1} denotes the SDF between dates t and $t+1$.

In the case where $U_t = C_t^{1-\rho}$, Λ_{t+1} reduces to the usual formula for the SDF when utility depends only on consumption (e.g. Epstein and Zin, 1991). If the (period) marginal utility of consumption depends on labor, then the SDF will be distorted in the usual ways through $\frac{U_{C,t+1}}{U_{C,t}}$. Even if U_C

[4] Jermann (1998) and Boldrin, Christiano, and Fisher (2001) find that models with additive habits have substantial difficulties in matching the dynamics of interest rates. Interest rates in their models are too volatile by an order of magnitude.

only depends on consumption, though (i.e. if period utility is separable between consumption and leisure), variation in labor will still affect the SDF through V_{t+1}: with recursive preferences, it is not generally possible to separate labor supply decisions from asset prices, unlike the case where preferences are separable between consumption and labor and over time.

2.2.1 Substituting in an asset return

Now consider an asset that pays $U_t U_{C,t}^{-1}$ as its dividend in each period. In the usual analysis of Epstein–Zin preferences, one substitutes the return on an asset that pays consumption as its dividend into the SDF. In the present case, dividing period utility, U_t, by the marginal utility of consumption intuitively converts U_t from utility units into consumption units.

We now derive the price of a claim to $U_t U_{C,t}^{-1}$. Denote the cum-dividend price of this asset as $W_{U,t}$. The appendix confirms that

$$W_{U,t} = V_t^{1-\rho} U_{C,t}^{-1} / (1-\beta) \tag{6}$$

and that we can substitute the return on this asset into the SDF to obtain

$$\Lambda_{t+1} = \beta^{\frac{1-\alpha_t}{1-\rho}} \left(\frac{U_{C,t+1}}{U_{C,t}} \right)^{\frac{1-\alpha_t}{1-\rho}} R_{U,t+1}^{\frac{\rho-\alpha_t}{1-\rho}} \tag{7}$$

$$\text{where } R_{U,t+1} \equiv \frac{W_{U,t+1}}{W_{U,t} - U_t U_{C,t}^{-1}} \tag{8}$$

The expression for the SDF in terms of an asset return is useful for two reasons. First, it helps show how the SDF is changed from the usual form when labor supply is included. For a general period utility function U_t, instead of the standard setup where only consumption matters, we see that the relevant return now depends on the entire evolution of future utility (scaled by marginal utility to convert it into consumption units) instead of just the evolution of consumption. Second, expressing the SDF in terms of an asset return will be important in the implementation of the approximation method for the model.

2.2.2 The market pricing kernel

In order to allow for an unexplained residual in interest rates, I assume that the market's pricing kernel is equal to the household's pricing kernel multiplied by a predetermined (but time-varying) shock,

$$M_{t+1} = \Lambda_{t+1} \exp(b_t) \tag{9}$$

The shock b_t induces variation in interest rates conditional on fundamentals. This type of residual is often treated as a simple shock to the rate of time preference. With Epstein–Zin preferences, though, a shock to the rate can have major effects on the behavior of the pricing kernel (specifically, a time-discount shock can end up accounting for the majority of the variance of the pricing kernel). The specification in (9) has the advantage that it allows for a residual in the short-term interest rate without having any further effect on the sources of risk premia.[5]

3 Aggregate supply

For the supply side of the model, I follow almost exactly the setup in Justiniano, Primiceri, and Tambalotti (JPT; 2010). JPT is a standard medium-scale New-Keynesian model. It has 7 fundamental shocks—price and wage markups, labor-augmenting technical change, investment-specific productivity, monetary policy, short-term interest rates, and government spending. In JPT's formulation, the monetary authority's inflation target is constant. I allow it to vary to help match the movements in the long end of the yield curve. Other than that and the preference specification, my model is nearly identical to theirs.

The model is also highly similar to Smets and Wouters (SW; 2003). The critical difference between the present setup and SW is that technology is difference-stationary rather than trend-

[5] Consider the Epstein–Zin preferences with constant risk aversion and no labor,

$$V_t = \left\{ B_1 C_t^{1-\rho} + B_2 \left(E_t \left[V_{t+1}^{1-\alpha} \right] \right)^{\frac{1-\rho}{1-\alpha}} \right\}^{\frac{1}{1-\rho}}$$

where in the usual specification, $B_1 = (1-\beta)$ and $B_2 = \beta$. There are two ways to affect the pure rate of time preference in such a way as to raise interest rates: either B_2 could fall or B_1 could rise. However, those two shifts have opposite effects on the level of V_t. We also know that any shock that affects V_t will be priced. So not only are shocks to the rate of time preference necessarily priced, but the sign of the price can take on different values for innocuous changes in the specification. I therefore skirt all these issues by assuming that the shocks to short-term interest rates are not driven by shifts in household preferences.

stationary, where the former is standard in the production-based asset pricing literature.[6] The difference-stationarity assumption helps generate large risk premia: when technology is trend-stationary, there is very little overall risk in the economy, so households must have an implausibly high coefficient of relative risk aversion in order to generate realistic asset prices.[7]

Since the model is standard and laid out in JPT and the main contribution of this paper is the preference specification and bond pricing, the remainder of this section gives a relatively short description of the production setup. The reader is referred to JPT for a more detailed analysis. My description follows theirs closely.

3.1 Producers of physical goods

Final-good producers are competitive in both input and output markets and have a CES production function,

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\lambda_{p,t}}} \, di \right]^{1+\lambda_{p,t}} \tag{10}$$

where i indexes types of intermediate goods, Y_t is output of the final good, which can be used for either consumption or investment, $Y_t(i)$ is the use of intermediate of type i, and the elasticity of substitution across the intermediates, which determines markups in the intermediate-goods sector, varies over time.

Intermediate-good producers are monopolists for their own goods with production function

$$Y_t(i) = \max \left\{ K_t(i)^{\gamma} Z_t^{1-\gamma} N_t(i)^{1-\gamma} - Z_t \bar{F}, 0 \right\} \tag{11}$$

where $\bar{F}$ is a fixed cost of production that ensures that profits are zero in steady state. $K_t(i)$ and $N_t(i)$ are intermediate-good producer's i purchases of capital and labor services, and Z_t is the level of labor-augmenting technology.

[6] A difference-stationary process has first-differences that follow a stationary process, so it is integrated of order one. A trend-stationary process, on the other hand, is a process that has random stationary deviations around a non-stochastic trend (where the trend is generally unmodeled and taken as exogenous).

[7] Below, I estimate average risk aversion to be 18.7 (ignoring the correction from Swanson, 2011). Rudebusch and Swanson (2011), who use stationary technology (with a slightly different preference specification) choose an analogous parameter to be 149.

3.2 Price setting

I assume Calvo pricing. In every period, a fraction $1 - \xi_p$ of intermediate good producers can change their prices, while the remainder index their prices following the rule,

$$P_t(i) = P_{t-1}(i)\,\pi_{t-1}^{\iota_p}\pi^{1-\iota_p} \tag{12}$$

where $P_t(i)$ is the price of good i in terms of the numeraire, $\pi_t \equiv P_t/P_{t-1}$ is aggregate inflation, and

$$P_t = \left[\int_0^1 P_t(i)^{\lambda_{p,t}^{-1}}\,di\right]^{\lambda_{p,t}} \tag{13}$$

is the aggregate price index (equal to the marginal cost of a unit of the final good). π is the steady-state inflation rate, and the parameter ι_p determines the degree of indexation to lagged inflation.

The firms that can choose their prices freely in a given period set them to maximize the present discounted value of profits over the period before they are allowed to choose a new price

$$E_t\left\{\sum_{s=0}^{\infty}\xi_p^s M_{t,t+s}\left[P_t(i)\left(\prod_{k=1}^{s}\pi_{t+k-1}^{\iota_p}\pi^{1-\iota_p}\right)Y_{t+s}(i) - W_{t+s}N_{t+s}(i) - R_{k,t+s}K_{t+s}\right]\right\} \tag{14}$$

where $M_{t,t+s} \equiv \prod_{j=1}^{s} M_{t+j}$, W_{t+s} is the wage rate, and R_{t+s}^k is the rental rate for capital.

3.3 Employment agencies and wage setting

Each household is a monopolistic supplier of specialized labor, $N_t(j)$. Competitive employment agencies aggregate labor supply into a homogeneous labor input (just as the final good producers aggregate intermediate goods) with the production function,

$$N_t = \left[\int_0^1 N_t(j)^{(1+\lambda_{w,t})^{-1}}\,dj\right]^{1+\lambda_{w,t}} \tag{15}$$

where, as with prices, $\lambda_{w,t}$ determines the elasticity of demand and hence markups in the labor market. $\lambda_{w,t}$ acts as a labor-supply shock. Since the employment agencies are competitive, the price of a unit of the homogeneous labor input is

$$W_t = \left[\int_0^1 W_t(j)^{\lambda_{w,t}^{-1}} \, dj \right]^{\lambda_{w,t}} \tag{16}$$

The labor demand function is then

$$N_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\frac{1+\lambda_{w,t}}{\lambda_{w,t}}} N_t \tag{17}$$

As with prices, wages can only be changed intermittently, with probability $(1 - \xi_w)$. If a household cannot change its wage, it indexes according to the rule

$$W_t(j) = W_{t-1}(j) \left(\pi_{t-1} \frac{Z_{t-1}}{Z_{t-2}} \right)^{\iota_w} (\pi \exp(\bar{g}))^{1-\iota_w} \tag{18}$$

where $\bar{g}$ is the average growth rate of technology. The household will choose its wage in a manner similar to how the intermediate-good firms set prices: it maximizes expected utility over the period that the wage will remain unchanged.

3.4 Capital and investment

Households accumulate capital according to the rule,

$$K_t = (1 - \delta) K_{t-1} + \mu_t \left(1 - S \left(\frac{I_t}{I_{t-1}} \right) \right) I_t \tag{19}$$

where δ is the depreciation rate and the function S incorporates adjustment costs in the rate of investment. In steady state, $S = S' = 0$ and $S'' > 0$. μ_t is a shock to the cost of investment at date t.

3.5 Government policy

The central bank follows a Taylor rule taking the form

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left[\pi_t^* \left(\frac{\pi_t}{\pi_t^*} \right)^{\phi_\pi} \left(\frac{Y_t}{Z_t} \right)^{\phi_Y} \right]^{1-\rho_R} \left[\frac{Y_t / Y_{t-1}}{Z_t / Z_{t-1}} \right]^{\phi_{dY}} \eta_{mp,t} \tag{20}$$

where R_t is the gross nominal interest rate, R is its steady-state value, and π_t^* is the inflation target at date t. Y_t denotes aggregate output. Y_t is scaled by Z_t so that the central bank responds to deviations of output from its stochastic trend (the level of technology).[8] The central bank is allowed to respond to both the level and change in the output gap. This flexibility helps ensure the model can match the dynamics of short-term interest rates, which is obviously critical for capturing the dynamics of the term structure. $\eta_{mp,t}$ is an exogenous monetary policy shock.

π_t^* is a time-varying inflation target, which can potentially help match the high inflation and long-term interest rates seen in the early part of the sample. More generally, π_t^* induces a level factor in the term structure. I take the inflation target as exogenous.

The government finances public spending by selling single-period bonds. Government expenditures, G_t, are a time-varying fraction of total output,

$$G_t = \left(1 - \frac{1}{g_t}\right) Y_t \tag{21}$$

where g_t follows an exogenous process defined below. Households receive no utility from government expenditures. As long as the share of output consumed by the government is stationary, that assumption will have minimal effects on asset prices.

3.6 Market clearing

The aggregate resource constraint is

$$C_t + I_t + G_t + K_{t-1} = Y_t \tag{22}$$

[8] To be more rigorous, the stochastic trend of output, $\tilde{Y}_t$ is defined as

$$\tilde{Y}_t \equiv \lim_{n \to \infty} E_t \left[Y_{t+n} / \exp\left(n\bar{g}\right) \right]$$

That is, it is the level to which output is expected to eventually return, where $\exp\left(n\bar{g}\right)$ takes into account expected technology growth. Since output and technology are cointegrated under balanced growth, $\tilde{Y}_t$ is proportional to Z_t.

3.7 Exogenous processes

The price and wage markup shocks follow ARMA(1,1) processes,

$$\log(1+\lambda_{p,t}) = \left(1-\rho_p\right)\log(1+\lambda_p) + \rho_p \log(1+\lambda_{p,t-1}) + \varepsilon_{p,t} - \theta_p \varepsilon_{p,t-1} \tag{23}$$

$$\log(1+\lambda_{w,t}) = (1-\rho_w)\log(1+\lambda_w) + \rho_w \log(1+\lambda_{w,t-1}) + \varepsilon_{w,t} - \theta_w \varepsilon_{w,t-1} \tag{24}$$

where $\varepsilon_{p,t} \sim N\left(0, \sigma_p^2\right)$ and $\varepsilon_{w,t} \sim N\left(0, \sigma_w^2\right)$. The ARMA(1,1) form potentially helps match both the high and low-frequency features of inflation.

Productivity has a unit root and its growth rate follows an MA(1) process,

$$\Delta z_t = \bar{z} + \varepsilon_{z,t} - \theta_z \varepsilon_{z,t-1} \tag{25}$$

where $\varepsilon_{z,t} \sim N\left(0, \sigma_z^2\right)$. While many recent models have studied AR(1) processes, I find that the MA(1) fits the data better (the estimate of θ_z is near zero and the behavior of the model is nearly unaffected by the choice of an AR or MA(1)).

The level of investment-specific productivity is assumed to be a stationary AR(1) process,

$$\log \mu_t = \rho_\mu \log \mu_{t-1} + \varepsilon_{\mu,t} \tag{26}$$

where $\varepsilon_{\mu,t} \sim N\left(0, \sigma_\mu^2\right)$. Note that μ_t simply determines the efficiency of the transformation of the final output good into the investment good, so investment still benefits from the unit-root innovations to Z_t.

The government's share of output, the monetary policy shock, the shock to the risk-free rate, and the shock to risk aversion follow AR(1) processes,

$$\log g_t = \left(1-\rho_g\right)\log g + \rho_g \log g_{t-1} + \varepsilon_{g,t} \tag{27}$$

$$\eta_{mp,t} = \rho_{mp}\eta_{mp,t-1} + \varepsilon_{mp,t} \tag{28}$$

$$b_t = (1-\rho_b)\,\bar{b} + \rho_b b_{t-1} + \varepsilon_{b,t} \tag{29}$$

$$\alpha_t = (1-\rho_\alpha)\,\bar{\alpha} + \rho_\alpha \alpha_{t-1} + \varepsilon_{\alpha,t} \tag{30}$$

where $\varepsilon_{g,t} \sim N\left(0, \sigma_g^2\right)$, $\varepsilon_{mp,t} \sim N\left(0, \sigma_{mp}^2\right)$, $\varepsilon_{\alpha,t} \sim N\left(0, \sigma_\alpha^2\right)$, $\varepsilon_{b,t} \sim N\left(0, \sigma_b^2\right)$.

While a number of recent papers have studied models with time-varying inflation targets (e.g. Gurkaynak, Sack, and Swanson, 2005; Doh, 2010), there is little understanding of what actually drives the inflation target. Because the inflation target has a very strong impact on long-term bond prices, the relationship between the inflation target and the other innovations is a key determinant of the prices of long-term bonds. I show below that the key shock that drives the pricing kernel is the innovation to labor-neutral technology (because technology is the only variable with a unit root). I therefore allow the innovation to the inflation target to be correlated with the innovation to labor-neutral technology,

$$\log \pi_t^* = (1 - \rho_\pi) \log \pi + \rho_\pi \log \pi_{t-1}^* + \varepsilon_{\pi*,t} + \sigma_{\pi,z} \varepsilon_{z,t}$$

with $\varepsilon_{\pi*,t} \sim N\left(0, \hat{\sigma}_{\pi*}^2\right)$. The parameter $\sigma_{\pi,z}$ determines how shocks to labor-neutral technology affect the inflation target. In robustness tests discussed below, I also allow for the shocks to investment technology and risk aversion to affect the inflation target. All of the shocks are otherwise assumed to be uncorrelated.

4 Model solution

The standard method for approximating models of the form studied here is perturbation. The drawback of perturbation methods for our purposes is that if we want time-variation in risk aversion to have any effect on the dynamics of the model, we need to take a third-order approximation to the model. Since the solution would be non-linear, we would have to use the particle filter or some other nonlinear method in order to calculate the marginal likelihood of the model. I have found, though, that it is in general very difficult to find the peak of the likelihood function for this model, and it would be infeasible with a method as slow as the particle filter. This is a common problem in models of the term structure (e.g. Ang and Piazzesi, 2003; Hamilton and Wu, 2011).

I therefore use the essentially affine approximation method described in Dew-Becker (2011b). The essentially affine method delivers an approximation to the equilibrium dynamics of the model that is linear in the state variables but still allows time-varying risk aversion to affect the behavior of the endogenous variables. Dew-Becker (2011b) describes the method in detail and show that Euler equation errors in simulated models are competitive with third-order perturbations. Local to the non-stochastic steady-state, the essentially affine approximation is as accurate as a first-order perturbation (in a Taylor sense), and hence less accurate than higher-order perturbations. However, in a stochastic setting, it performs well. This section gives a short overview of the method, and the appendix provides further details.

Denote the vector of the variables in the model (including the exogenous processes) as X_t and the vector of fundamental shocks as $\varepsilon_t \equiv [\varepsilon_{mp,t}, \varepsilon_{z,t}, \varepsilon_{b,t}, \varepsilon_{\mu,t}, \varepsilon_{g,t}, \varepsilon_{p,t}, \varepsilon_{w,t}, \varepsilon_{\alpha,t}, \varepsilon_{\pi*,t}]$. The equations determining the equilibrium of the model take the form

$$0 = G\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) \tag{31}$$

where the expectation operator may appear in the function G. There is one equation for each variable. σ is the perturbation parameter controlling the variance of the shocks. We will approximate around the point $\sigma = 0$, with the non-stochastic steady-state defined as the point $\bar{X}$ such that

$$0 = G\left(\bar{X}, \bar{X}, 0\right)$$

The equations G can be divided into two types: those that do not involve taking expectations over the SDF and those that do.

$$G\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) = \begin{bmatrix} D\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) \\ E_t\left[M\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) F\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right)\right] \end{bmatrix} \tag{32}$$

where D and F are vector-valued functions and M is the (scalar-valued) stochastic discount factor.[9]

For the equations that do not involve the SDF, I use standard perturbation methods and simply

[9]Note that this formulation does not actually restrict F. Specifically, suppose there were a set of equilibrium conditions $1 = E_t h\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right)$, i.e. that do not involve the SDF. We could simply say that $F\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) \equiv h\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) / M\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right)$.

take a log-linear approximation. We thus approximate D as

$$\begin{aligned} 0 &= \log\left(D\left(\exp\left(x_t\right), \exp\left(x_{t+1}\right), \sigma\varepsilon_{t+1}\right) + 1\right) & (33) \\ 0 &\approx d_0 + d_x\hat{x}_t + d_{x'}\hat{x}_{t+1} + d_\varepsilon \sigma\varepsilon_{t+1} & (34) \end{aligned}$$

where the terms d_0, d_x, $d_{x'}$, and d_ε are coefficients from a Taylor approximation and

$$\begin{aligned} x_t &\equiv \log X_t & (35) \\ \hat{x}_t &\equiv \log X_t - \log \bar{X} & (36) \end{aligned}$$

D will include equations such as budget constraints, laws of motion for exogenous processes, and optimization conditions for purely intratemporal decisions.

The second set of equations is dynamic and involves expectations. In many economic models, including the present one, equations involving expectations take the form

$$1 = E_t\left[M\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right) F\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right)\right] \tag{37}$$

The key source of non-linearity in the model is the time-variation in risk aversion, which induces heteroskedasticity in the SDF. It is therefore natural to deal with M and F separately to isolate the relevant non-linearity.

I now show that if we log-linearize F, we can transform (37) into a linear condition that can be solved alongside the remaining equations. $M\left(X_t, X_{t+1}, \sigma\varepsilon_{t+1}\right)$ will not even be log-linear in the state variables, but we will be able to state the equilibrium conditions in as a set of linear expectational difference equations.

First, guess that the approximated equilibrium dynamics of the model take the form

$$\hat{x}_{t+1} = C + \Phi\hat{x}_t + \Psi\varepsilon_{t+1} \tag{38}$$

We confirm in the end that the solution is actually in this form.

The next step then is to take approximations to M and F separately. Log-linearizing F is straightforward, and we obtain,

$$\log F\left(x_t, x_{t+1}, \sigma\varepsilon_{t+1}\right) \approx f_0 + f_x x_t + f_{x'} x_{t+1} + f_\varepsilon \sigma\varepsilon_{t+1} \tag{39}$$

For M, in the case of the preferences laid out in section 2, the appendix shows that is it is possible to derive a first-order accurate expression of the form

$$m_{t+1}^{(1)} = m_0 + m_x \hat{x}_t + \left(\kappa_0 + \alpha_t \kappa_1\right) \sigma\varepsilon_{t+1} - \frac{1}{2}\sigma^2 \alpha_t^2 \kappa_1 \Sigma \kappa_1' \tag{40}$$

where Σ is the variance matrix of ε_t. The superscript (1) indicates that $m_{t+1}^{(1)}$ is first-order accurate for the true SDF. (40) is the essentially affine form from Duffee (2002) that is widely used in the bond pricing literature.

Taking the expectation of the approximated Euler equation yields,

$$\begin{aligned}
0 &= \log E_t \exp \begin{pmatrix} m_0 + m_x \hat{x}_t + \left(\kappa_0 + \alpha_t \kappa_1\right) \sigma\varepsilon_{t+1} - \frac{1}{2}\alpha_t^2 \sigma^2 \kappa_1 \Sigma \kappa_1' \\ + f_0 + f_x x_t + f_{x'} x_{t+1} + f_\varepsilon \sigma\varepsilon_{t+1} \end{pmatrix} && (41) \\
0 &= m_0 + m_x \hat{x}_t + f_0 + f_x x_t + f_{x'}\left(C + \Phi \hat{x}_t\right) \\
&\quad + \frac{1}{2}\sigma^2 \left(f_{x'} + f_\varepsilon\right) \Psi \Sigma \Psi' \left(f_{x'}' + f_\varepsilon'\right) + \alpha_t \sigma^2 \kappa_1 \Sigma \Psi' \left(f_{x'}' + f_\varepsilon'\right) && (42)
\end{aligned}$$

Since every equation in the system is now linear in the variables of the model, we can solve the system for the parameters Φ and Ψ from (38). Specifically, we solve the following system,

$$\begin{aligned}
0 &= d_0 + d_x \hat{x}_t + d_{x'} \hat{x}_{t+1} + d_\varepsilon \sigma\varepsilon_{t+1} && (43) \\
0 &= m_0 + m_x \hat{x}_t + f_0 + f_x x_t + f_{x'}\left(C + \Phi \hat{x}_t\right) \\
&\quad + \frac{1}{2}\sigma^2 \left(f_{x'} + f_\varepsilon\right) \Psi \Sigma \Psi' \left(f_{x'}' + f_\varepsilon'\right) + \alpha_t \sigma^2 \kappa_1 \Sigma \Psi' \left(f_{x'}' + f_\varepsilon'\right) && (44)
\end{aligned}$$

at the point $\sigma = 1$. The reason that the essentially affine SDF is useful is that the expectation in (44) will be linear in the state variables, so we have a simple linear system to solve. This system can be solved through, for example, Sims' (2001) Gensys algorithm.

Dew-Becker (2011b) shows that the transition function for the model obtained through the essentially affine method is first-order accurate for the true transition function and first-order equivalent to a first-order perturbation. Clearly, though, the approximation includes higher-order terms that account for movements in risk aversion. α_t will affect not only asset prices but also the dynamics of real variables. Dew-Becker (2011b) calibrates a simple version of the RBC model with time-varying risk aversion and finds that the essentially affine approximation has accuracy between that of second and third-order perturbations.

Standard results derived in the appendix also deliver prices of real and nominal zero-coupon bonds.

5 Empirics

I estimate the model using standard Bayesian methods. The observable data is the same as in JPT, but with bond prices added. Both real variables and bond prices are linear functions of the underlying state variables contained in the vector x_t, so we can write the model in state-space form and measure the likelihood using the Kalman filter. I proceed by finding the posterior mode and running a Monte Carlo chain from that point to sample from full posterior distribution. The appendix describes further details of the estimation.

5.1 Data

The sample is 1983q1 to 2004q4. I do not include the financial crisis in the sample because the zero lower bound on nominal interest rates becomes binding, a phenomenon that the model is not designed to capture. The sample is cut off in 1983 due to the evidence for breaks in monetary policy at earlier dates (e.g. the shifts in the Federal Reserve chairmanship between Martin, Burns, Miller, and Volcker; see Clarida, Gali, and Gertler, 1999).

The observable variables are real GDP, consumption, and investment growth, hours worked per capita, wage and price inflation, and yields on three-month, 1, 2, 3, 5, 10, and 20-year Treasury bonds. All yields are from Gurkaynak, Sack, and Wright (2006) except for the three-month yield, which is from the Fama risk-free rate CRSP file. The bond yields and inflation rates are always reported in annualized percentage points, unless otherwise noted. The real variables are all obtained from the BEA and the BLS. Consumption is defined as expenditures on non-durables and services, while investment is the sum of residential and non-residential fixed investment and consumer durables expenditures. Real wages are calculated as nominal compensation per hour in the non-farm business sector (from the BLS) divided by the GDP deflator. The change in the log GDP deflator is the measure of inflation. Hours worked per capita in the non-farm business sector are obtained from Francis and Ramey (2009) as updated on Valerie Ramey's website. None of the variables are detrended.

Figure 1 plots the data used in the estimation (with the exception of the intermediate-term bond yields). Output, consumption, and investment growth all look stationary over the sample and relatively homoskedastic. Hours worked per capita has a strong upward trend in this sample. Interest rates decline significantly over the sample, even though inflation only declines marginally. The short-term interest rate is substantially more volatile than the long-term rate, and the term spread is clearly countercyclical.

The model has 9 fundamental shocks, but we have 14 observable variables. I follow JPT and other macro papers in assuming that the 6 macro variables plus the short-term interest rate are observed without error. For the remaining bonds, I assume that the yields have independent measurement errors with identical standard deviations.

5.2 Calibrated parameters

I calibrate a number of parameters following Christiano, Motto, and Rostagno (2012). The parameters that are calibrated are those that are expected to be more difficult to estimate, such as steady-state values.

Capital's steady-state share of income is set to 1/3. The steady-state growth rate of per-capita output is 1.88 percent per year. Steady-state price and wage markups are 0.22 and 0.05, respectively. Steady-state hours worked per capita is 675. Steady-state annual inflation is 2.5 percent. The price and wage stickiness parameters, ξ_w and ξ_p are set to 0.66.

For preferences, I focus on estimating the risk aversion process, though I also allow the habit parameter to be estimated to help match the persistence of consumption growth. I calibrate the EIS, as is common in the New Keynesian literature, selecting a value of 1.5, which is commonly used in the asset pricing literature. The inverse Frisch elasticity is set to 1, and the pure rate of time preference, β, is set to 0.9962, which implies an annualized discount rate of 1.5 percent.

Finally, I assume the inflation target follows nearly a random walk with $\rho_{\pi *} = 0.99$, consistent with the idea that the target is highly persistent. The assumption that $\rho_{\pi*} < 1$ ensures that inflation is stationary so that there is a steady-state around which we can approximate. The remainder of the parameters are estimated

5.3 Priors

Table 1 lists the estimated parameters and their priors in the benchmark estimation. The priors for the parameters shared with other recent studies are generally the same as in those papers. For the parameters unique to this paper, e.g. the variance of risk aversion, I use relatively flat priors. For the volatility of risk aversion, I choose a beta distribution over the ratio of the unconditional standard deviation of risk aversion to its mean. This means that average risk aversion is forced to be at least one standard deviation above zero.

In addition to the prior on the parameters, I also incorporate prior information about the steady-state of the model. Specifically, I add a penalty to the likelihood function for the deviation of the steady-state term spread from its sample average of 2 percent using a normal distribution with standard deviation of 0.1. This penalty helps eliminate local extrema in the likelihood that imply that the steady-state term structure is strongly downward sloping.

5.4 Posterior mode

Table 1 lists the posterior modes for the parameters along with the standard deviation and 2.5th and 97.5th percentiles of the posterior distribution. Many of the posterior modes are reasonably close to the corresponding prior means, and the standard deviations of the posteriors for each parameter are all substantially smaller than those of the priors, implying that the parameters are well identified. I focus my discussion mainly on those parameters that are unique to this model or where the posterior mode differs notably from the prior.

The prior for the standard deviation of the innovations to the inflation target favors a reasonably low standard deviation, 0.2 percent (with a wide prior), and the posterior is consistent with that—the estimated standard deviation of the innovations to the annualized inflation target in each quarter is 0.24 percent. This helps the model capture the observed volatility of the level factor in bond yields, but it is perhaps somewhat high. Intuitively, we will see below that the model ascribes shifts in the level factor interest rates to shifts in the inflation target. So the model needs a large degree of volatility in the inflation target to generate the observed volatility in the level factor. Bekaert, Cho, and Moreno (2009) obtain a similar result.

The shock to the level of labor-neutral technology has an important effect on the inflation target, accounting for roughly half of the variance of its innovations. Following a positive innovation to technology, the central bank is estimated to lower its inflation target, consistent with the idea that following beneficial supply shocks that drive inflation downward, the central bank takes the opportunity to drive inflation lower persistently (e.g. Gurkaynak, Sack, and Swanson, 2005). This mechanism will turn out to be critical to the model's ability to generate a strongly upward sloping term structure.

The labor-neutral technology shock has a standard deviation of 1.15 percent and an autocorrelation of roughly zero. The permanent component of the technology process (the Beveridge–Nelson trend) thus has a standard deviation of 1.11, which is similar to the values often calibrated in the production-based asset pricing literature (e.g. Tallarini, 2000, and Gourio, 2012). The estimated long-run variance of technology growth is far smaller than the values calibrated in the long-run risks literature (e.g. Bansal and Yaron, 2004, and Kaltenbrunner and Lochstoer, 2010), but it is consistent with estimates obtained in JPT and SW and with simple univariate estimates from consumption and output data (Dew-Becker, 2012).

The standard deviation of the investment technology shock is relatively small, especially compared to JPT. Investment adjustment costs are also relatively small. There are two reasons we obtain this result. First, one of the strongest pieces of evidence for high investment adjustment costs is the behavior of equity prices, but they are not included here. Second, we will see below that shocks to investment technology can have strong effects on interest rate dynamics. The inclusion of long-term bond prices forces the model to keep shocks to investment technology relatively subdued, because they otherwise imply that long-term interest rates are counterfactually volatile.

As in Smets and Wouters (2007), JPT, and Christiano, Motto, and Rostagno (2012), the government spending shock is estimated to follow nearly a unit root, explaining the trend in the consumption-output ratio over the sample.

Risk aversion is estimated to have an average value of 23.46, which is relatively high. It also has shocks with a quarterly autocorrelation of 0.9, though, which makes it volatile – its standard deviation is estimated to be 1/3 of its mean.

6 Asset pricing

This section studies the asset-pricing implications of the model. I first analyze the fit of the model to the term structure and show that it is competitive with a non-structural model. Next, I decompose the variance of the SDF to understand the source of the positive term premium in the model. I then analyze the prices of other assets, including the aggregate capital stock and a claim to aggregate profits.

6.1 Bond prices

6.1.1 Fitted yields

Figure 2 plots the deviations of the fitted yields from their actual values for the five yields that are assumed to be measured with error (reported in annualized basis points). The estimated standard deviation of these fitting errors is 17 basis points, which is economically small compared to the overall variation of the yields that is on the order of hundreds of basis points. The errors are all centered around zero, meaning that the model can capture the shape of the term structure on average. The volatility of the errors looks somewhat higher for the 1 and 5-year yields and in the earlier part of the sample. There is clearly some autocorrelation in the errors; the fitted value for the 3-year yield is consistently too high in the first half of the sample, and the 4-year fitted yield is consistently too low in the second half, for example. And there is also some cross-correlation in the errors; the first principal component explains 37 percent of the total variance of the errors (twice what it would if the errors were orthogonal). These are thus clearly not classical (i.i.d.) measurement errors, but their small mean and volatility shows that the model does a reasonable job of fitting the data, and they are not disturbingly far from white noise.

While there are nine unobservable shock processes that can help us match the data, the model is asked to fit 14 data series, so obtaining a good fit for the bond yields is not trivial. Loosely, we have 6 macro variables that identify 6 shock processes, plus three extra processes (the monetary policy shock, the inflation target, and risk aversion) that can be used to fit the bond yields. The degrees of freedom here are thus comparable to a non-structural bond-pricing model with three unobservable factors, but we also have numerous constraints on dynamics and risk prices.

Table 2 compares the measurement errors to the fitting errors obtained from simple non-structural models. Specifically, I regress interest rates on two or three factors extracted from the data and report the standard deviation of the residuals. In each case one of the factors is the 1-quarter interest rate so that the non-structural models match the feature of the structural models that that yield is measured without error. The remaining one or two factors are estimated as principal components from the residuals from regressions of the remaining yields on the 1-quarter yield.

The first row of table 2 lists the standard deviations of the yield errors in basis points obtained from a simple three-factor non-structural model, while the second row uses two factors. The standard deviations with three factors are all between 5.8 and 18.1 basis points, compared with a standard deviation of 16.72 basis points for the benchmark model. Though the fit of the structural model is clearly somewhat poorer than the non-structural regressions, it is certainly competitive, and the economic significance of the difference in the goodness of fit is minimal.

The third and fourth rows of table 2 report the measurement errors in constrained models that assume constant relative risk aversion and a constant inflation target (the other parameters are reestimated). In both cases, the measurement errors have standard deviations roughly 50 percent larger than the unrestricted model, and they are larger than the errors from both the two and three-factor non-structural models.

An alternative way to measure the importance of allowing risk aversion and the inflation target to vary over time is to look at the log marginal likelihood of the three models. Table 3 reports the marginal likelihoods of the benchmark model and the two constrained models. The difference, i.e. the Bayes factor, is 52 log points for the model with constant risk aversion and 89 log points for the model with constant risk aversion, implying that we would need to have an extraordinarily strong prior that risk aversion or the inflation target is fixed in order to prefer one of the restricted models over the benchmark. Comparing the two restricted models, the version with the constant risk aversion is substantially preferred over the model with the constant inflation target.

6.1.2 Steady-state yields

Another way to evaluate the fit of the model is to ask whether the steady state of the model matches the average term structure in the data. Looking at the steady state keeps the Kalman filter from using large deviations in the unobservable state variables to fit the term structure. Figure 3 plots the average term structure in the sample along with its model-implied steady state. The solid black line gives the steady-state term structure in the model, renormalized so that the ten-year yield matches the empirical ten-year yield.[10] To capture the uncertainty in the empirical term structure, the grey area gives the 95-percent confidence intervals for the means of the empirical yields relative to the ten-year yield (i.e. the confidence intervals for the spreads; the intervals are calculated using the Newey–West method with lag a 6-quarter lag window). What figure 3 shows is that the model matches the spread between the 10- and 2-year yields very well, and the model's steady-state 10-year/1-quarter term spread is well within the 95-percent confidence interval. However, the model does have some trouble matching the degree of curvature of the average term structure below two years, which could potentially be driven by unmodeled liquidity premia.

Two features of the model are critical for generating the large steady-state term premium: first, following a positive shock to technology, the central bank's inflation target falls; second, variation in risk aversion raises the premia on risky assets. To see the prima facie evidence that these two effects are key, figure 3 includes two lines giving the steady-state term structure in constrained models. The first line assumes that innovations to the inflation target are uncorrelated with the permanent technology shock, while the second line assumes that risk aversion is fixed. Neither line reestimates the other parameters, so they simply isolate the effects of those two features of the model.

The line exiting the top of the chart is for the model when shocks to technology are assumed to have no impact on the inflation target. We then obtain the usual result that the term structure is downward-sloping, and the steady-state term spread is -137 basis points. Time-varying risk aversion also turns out to be important, though. When risk aversion is fixed, the term structure is still upward-sloping, but the spread is much smaller—only 120 basis points in steady-state, compared to 207 in the data and 193 in the benchmark model.

Models with Epstein–Zin preferences often have trouble matching the behavior of the real yield curve (see the discussion in Bansal, Kiku, and Yaron, 2012). In this model, the real yield curve is slightly downward sloping on average – the real 10/1-year term spread is -0.6 in steady-state. Following a positive technology shock, real interest rates rise, both because the marginal product of

[10] I use this normalization because the estimated inflation target is above zero through most of the sample. The unconditional variance of the inflation target is sufficiently high that its average level is not well identified.

capital is high and because the central bank wants to hold inflation down. This causes real bonds to be hedges against technology shocks. As to whether this is an empirically plausible result, Evans (1998) finds that in the UK, which has the longest currently available series of inflation-indexed bonds, the real 10/1 year term spread averaged -0.88 percent between 1984 and 1995. Bansal, Kiku, and Yaron (2012) extend his results and find an average spread of -1.92 percent between 1996 and 2008.

6.1.3 Term premia

The size of the steady-state term spread in the model can be interpreted as the average term premium—it is the excess return (in logs) that an investor earns in expectation by buying a long-term bond and holding it to maturity instead of buying short-term bonds and rolling them over for the same amount of time. An important feature of this model is that risk aversion varies over time, which should make the term premium also vary over time.

The top panel of figure 4 plots the expected annualized excess return on holding a ten-year nominal bond (over a one-quarter bond) from the benchmark model against the expected excess return obtained from a simple VAR(1) in the 1-quarter, 5-year, and 10-year yields.[11] The realized annual returns are plotted in light gray.

The structural forecast is highly correlated (34 percent) with the fitted value from the VAR, though it is substantially smoother. Both series rise in the recessions in 1991 and 2001, though the VAR implies that risk premia on long-term bonds were substantially higher in the late 1980's than does the structural model. The high variance of the forecast from the VAR is somewhat implausible and suggests a degree of overfitting.

The bottom panel of figure 4 plots the term premium from the benchmark model against the term spread. The term premium is defined as the spread between the 10-year yield and the average of the expected 1-quarter yields over the life of a 10-year bond. The variance of the term premium is non-trivial in comparison to the term spread. In the two recessions in the sample, the increases in the term spread are substantially larger than the movements in the term premium, but the term premium does rise in both episodes. Interestingly, the movements in the term spread outside of the

[11] Where, as in Campbell and Shiller (1991), the return on the 10-year bond is measured as $-39y_{10,t}/400 + 40y_{10,t}/400$, where $y_{10,t}$ is the annualized 10-year yield measured in percentage points. The implied dynamics of the ten-year yield from the VAR are used to obtain a forecast of expected excess returns.

two recessions seem almost entirely driven by movements in the term premium. In particular, the rises in the term spread in 1985, 1987, 1997, and 1999 are all associated with increases in the term premium of similar magnitudes. On the other hand, the inversions of the yield curve in 1989 and 2000, both just prior to recessions, are associated with only minor declines in risk aversion, and the subsequent rises in the term spread with similarly small rises in risk aversion.

6.1.4 Variation in interest rates

To show how the bond yields respond to the various shocks, figure 5 plots responses of a level, slope, and curvature factor to the 9 fundamental shocks. Following Bekaert, Cho, and Moreno (2010), the level factor is defined as the average of the 1-quarter and 5 and 10-year yields; the slope factor is the 10-year/1-quarter term spread; and curvature is the sum of the 5 and 1-year yields minus twice the 3-year yield. The shocks are orthogonalized in the sense that the interactions between the inflation target and risk aversion and the other shocks are switched off. So figure 5 shows, for example, the effect of a pure increase in the level of technology, holding the inflation target fixed. The shocks are all unit standard deviation impulses.

For the level factor, most shocks have only small effects. The monetary policy and risk aversion shocks both reduce the level factor, but at relatively high frequencies. The low-frequency movements, as we would expect, are driven by the inflation target shock. Since shocks to labor-neutral technology are estimated to reduce the inflation target, they lead to persistently low inflation. The model's explanation of shifts in the level factor is thus simple: the level of the term structure is determined by the central bank's inflation target, which is in turn partly determined by shocks to technology.

As we would expect, an increase in risk aversion raises the term spread. High risk aversion is associated with a high term premium, and hence high interest rates on long-term bonds relative to the short-term interest rate. This result fits with results on bond return forecasting (Campbell and Shiller, 1988) and the fact that the term spread forecasts high equity returns (Fama and French, 1989).

6.1.5 Effect of bond price data on estimates

The majority of the macro literature does not include the term structure of interest rates in estimation. A key question, then, is how that addition here affects the estimates. The second column of table 6 shows estimates we obtain when bond prices are not included in the estimation. Because we now have fewer observables, I eliminate the shocks to risk aversion and the inflation target from the model. Those two processes are largely identified from the behavior of long-term bond prices. The first column replicates the benchmark estimates from table 1.

Many of the parameters estimates are similar between the benchmark estimates and those without bond prices. There are two sets of parameters that are noticeably different, though. First, as we might expect, the monetary policy parameters change somewhat. Since these parameters directly affect the behavior of interest rates, it would be surprising if they did not change. More interestingly, we see large changes in the estimates of the parameters related to investment. The size of investment technology shocks is much larger when bond prices are excluded, the shocks are more persistent, and investment adjustment costs are multiplied by a factor of 7 – rising from 0.14 to 1.06. Without bond prices in the estimation, we come closer to the results of JPT that investment technology shocks are the dominant drivers of investment behavior.

To help see why adding bond prices causes the estimation to place so much less weight on investment shocks, figure 6 plots impulse responses of the one-quarter, 5-year, and 10-year (annualized) interest rates to an investment shock at the two sets of estimates. Without bond prices included in the estimation, the investment shock has far larger effects on interest rates than with bond prices. On the impact of the shock, the five-year rate rises by three times as much when bond prices are excluded than when they are not. The long-term yields also decline substantially at longer horizons. On the other hand, when bond prices are included, the investment shock has a small initial impact on long-term interest rates which rapidly falls to zero.

6.2 Determinants of asset prices

6.2.1 The variance of the SDF

An asset's expected excess return over the real riskless interest rate is determined by its covariance with the stochastic discount factor. One of the more interesting outputs of a model as rich as this one is the variance decomposition for the SDF. Table 4 reports a variance decomposition for the SDF at the one-quarter horizon. The variance of the SDF is essentially entirely driven by the neutral technology and risk aversion shocks. The bar chart in the bottom panel of table 4 decomposes the variance of the SDF into components coming from the neutral technology shock, the risk aversion shock, and the remaining shocks combined. The lines at the top of each bar give the 2.5 and 97.5 percentiles of the posterior distribution. The 97.5 percentile for the variance share in the SDF for non-technology and non-risk aversion shocks is less than 2 percent

On first glance this result might be somewhat surprising, but it is in fact a deep characteristic of models with Epstein–Zin preferences with a high EIS and high risk aversion. The results are relatively straightforward to see in a log-linearized version of the model. The appendix shows that we obtain the following approximation to the pricing kernel,

$$\Delta E_{t+1} \log \Lambda_{t+1} = \Delta E_{t+1} \log U_{C,t+1} \tag{45}$$

$$+\frac{1-\alpha_t}{1-\rho} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1} \tag{46}$$

$$+\frac{\rho-\alpha_t}{1-\rho} \psi_{rp} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \alpha_{t+j} \tag{47}$$

where $\Delta E_{t+1} \equiv E_{t+1} - E_t$ denotes the change in expectations as of date $t+1$ and $\Delta \equiv (1-L)$ is the first-difference operator. ψ_{rp} is a function of the other parameters, and determines the relationship between expected returns on the utility portfolio and risk aversion.[12] This result is similar to one obtained by Tanaka (2012).

The innovation to the SDF has three components: the innovation to current marginal utility, the innovation to future expected growth in utility, and the innovation to risk aversion. When β is equal to 1, a shock that has no long-run effect on the level of utility will induce zero change in $\Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$. A permanent shock to utility, though, will raises $\Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$ by a positive amount.[13] The price of risk for a particular shock, which is simply the amount that it affects the pricing kernel, depends on the sum of its effects on $\Delta E_{t+1} \log U_{C,t+1}$, $\frac{1-\alpha_t}{1-\rho} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$, and risk aversion.

[12] That is, $E_t r_{U,t+1}$ is equal to a constant plus $\psi_{rp} \alpha_t$.

[13] Dew-Becker and Giglio (2013) use this insight to show that when household utility depends only on consumption, risk prices under Epstein–Zin preferences depend almost purely on the very long-run effects of shocks. A shock with a purely transitory impact on consumption will tend to have a very small risk price under standard parameterizations.

When $\eta = 0$, a one-percent temporary increase in consumption will reduce $\log U_{C,t+1}$ by ρ percent, while it will have roughly zero effect on $\Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$ (precisely zero when $\beta = 1$). A one-percent permanent increase in consumption, though, will both reduce $\log U_{C,t+1}$ by ρ percent and also raise $\sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$ by 1 percent, thus shifting the SDF down by $-\rho + \frac{1-\alpha_t}{1-\rho}$. Given the estimated parameters, and for standard calibrations where $\alpha > \rho^{-1}$, $-\rho + \frac{1-\alpha_t}{1-\rho}$ is much more negative than $-\rho$, which is why the temporary shocks have very small risk prices.

The above discussion focuses on just consumption. The reason is that under balanced growth, labor supply is stationary. That means that the changes in labor supply tend to induce minimal effects on $\Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$ (again, zero if $\beta = 1$). The specification of the period utility function also implies that U_C does not depend on labor supply. So to understand the pricing kernel, it is sufficient to understand consumption dynamics.

To help elucidate the results further, table 5 lists the marginal effects of each of the shocks on the pricing kernel and the three components of the decomposition above, $\Delta E_{t+1} \log U_{C,t+1}$, $\frac{1-\bar{\alpha}}{1-\rho} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$, and $\frac{\rho-\bar{\alpha}}{1-\rho} \psi_{rp} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \alpha_{t+j}$. That is, the first column is equal to the risk prices ($\kappa_0 + \bar{\alpha}\kappa_1$) from equation (40), and the remaining columns decompose the risk prices. While we can see that the various shocks have effects on $\Delta E_{t+1} \log U_{C,t+1}$ of a similar order of magnitude, those effects are dwarfed by the effects of risk aversion on $\frac{1-\bar{\alpha}}{1-\rho} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta \log U_{t+j+1}$ and of risk aversion on $\frac{\rho-\bar{\alpha}}{1-\rho} \psi_{rp} \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \alpha_{t+j}$.

6.2.2 Impulse responses

Since the technology and risk aversion shocks are the key to understanding the pricing kernel, it is natural to ask how they affect the economy. Figure 7 plots impulse responses to an increase in labor-neutral technology and a decrease in risk aversion. These impulse responses are different from those in figure 5 because they do not turn off the interactions between the inflation target and the neutral technology shock. The idea is that we want to see what happens on average following the technology and risk aversion shocks because that behavior is what is relevant for understanding the correlations with the SDF.

Following the technology shock, inflation falls, while output and real interest rates rise: a standard positive supply shock. The declines in inflation and nominal interest rates are especially pronounced and persistent. The long-run decline in inflation is obviously driven by the shock to the inflation target. It is somewhat surprising how rapidly inflation falls, though. There are at least two reasons for this result. First, since output is well below trend, marginal costs are low. Second, though, because current inflation depends substantially on expected future inflation, when the inflation target falls, the expectations channel causes firms to want to immediately reduce inflation.

A decline in inflation on the impact of a technology shock is obtained in JPT and SW, even though they do not have time-varying inflation targets. In all of these models, because prices are sticky, when there is a positive supply shock, rather than raising output, firms simply produce the same quantity as previously, thereby reducing employment and hence demand. Positive technology shocks are thus associated with small increases in output and large declines in the output gap. There is also a small empirical literature that provides more reduced-form evidence on effects of this sort using direct measures of technology (e.g. Basu, Fernald, and Kimball, 2006).

Figure 7 also reports impulse responses for a decline in risk aversion. The effects are very small compared to those for the technology shock. A unit standard deviation decline in risk aversion lowers output by less than one tenth of one percent, and has similar effects on the other variables plotted in figure 7. Figures 4 through 7 show, then, that while shifts in risk aversion have important effects in helping the model match asset prices, they have only minimal effects on the real economy.

6.2.3 Robustness

Given how important the correlation between shocks to the inflation target and risk aversion is, an obvious question is what happens if we allow more of the shocks to be correlated with the inflation target. I allow the investment technology to also play a role, since the conceptual difference between investment and labor-neutral technology shocks is small. I also allow the risk aversion shock to be correlated with the inflation target since we saw above that shocks to risk aversion are priced. The third pair of columns in table 6 reports results from this experiment. The parameter estimates are largely unchanged, and risk aversion and investment technology are estimated to have only relatively small effects on the inflation target compared to labor-neutral technology.

Another good question is why labor-neutral technology follows a unit root but investment-specific technology and the other shocks do not. The last pair of columns in table 6 uses a uniform prior on the persistence of the investment-specific technology shock and the shocks to price and wage markups. The model remains close to the original interior solution, suggesting that the assumption that these shocks do not have unit roots is not restrictive.[14]

6.2.4 Other asset prices

After the SDF, table 4 reports variance decompositions for the returns on a number of assets. The bottom panel of table 4 reports the fraction of the variance of the one-quarter innovation to each return coming from the neutral technology shock, the risk aversion shock, and all other shocks combined.

Column 2 reports the variance decomposition for the return on the utility portfolio. 53 percent of its variance comes from the discount rate shock. The reason is simply that the utility claim has a relatively long duration, like that of a consol with a coupon that grows at the average rate of the economy, so shifts in real interest rates have a large effect on its price. The discount rate shock mainly affects real interest rates, so it drives the variance of the utility claim. Row 11 shows that the correlation of the utility return with the SDF is -0.39.

More interestingly, the third and fourth columns of table 4 report variance decompositions for a claim on the profits earned by capital and the same claim levered two to one on short-term nominal debt.[15] Once again, little of the variance of the return is driven by the technology or risk aversion shocks. Column 5 shows that if capital claims are levered on real debt, then they become more highly correlated with the pricing kernel, and the model can actually generate a large equity premium (though one that is still too small by half). The reason is that the real term structure is downward sloping, so a portfolio that is short real bonds will earn a positive return, whereas being short nominal bonds drives the return downward.

6.2.5 The Hansen–Jagannathan bound

Figure 8 plots the fitted annualized Hansen–Jagannathan (HJ; 1991) bound against the nominal term spread. The Hansen–Jagannathan bound says that the maximum Sharpe ratio of any asset must be less than the ratio of the standard deviation of the pricing kernel to its mean. The estimated steady-state level of the HJ bound – 0.55 – is high enough to account for the observed Sharpe ratio

[14] I also tried constraining the investment technology shock to have a unit root and then reestimated the model. The log-likeilhood at that point was 90 log points lower than at the optimum.

[15] The profits earned by capital are modeled as the revenues the household receives from renting out its capital.

on the aggregate stock market in this sample of 0.26. More importantly, though, the price of risk in this model is highly volatile. The estimated standard deviation of the Hansen–Jagannathan bound is 95 percent of its mean. The level of variability here is somewhat higher than but still similar to that used in Dew-Becker (2011a) to match the degree of predictability observed for aggregate stock returns in the post-war sample. Furthermore, we can see that the Hansen–Jagannathan bound is highly correlated with the term spread, consistent with the view that the term spread helps forecast returns on a range of assets (e.g. Fama and French, 1989)

7 The real economy

Up to now, the analysis has focused mainly on asset pricing. But the model gives a rich description of the real side of the economy. While I leave a deeper analysis of New Keynesian models to papers focused on those models for their own sake, the interaction of the real side of the economy with asset prices is important to this paper.

Figure 9 gives a variance decomposition for the variables used in the estimation. The figure decomposes the variance of each variable at frequencies of 6 to 32 quarters into components coming from each of the structural shocks.[16] Except for consumption growth, for which the interest rate shock is dominant, none of the other variables examined in figure 9 are dominated by any particular shock. Notably, the shock to risk aversion has almost no effect on the variance of any of the real variables at business-cycle frequencies. Its largest effect is on consumption growth, for which the variance share is still only 4 percent. The bar for the term spread, though, shows that risk aversion has a somewhat larger effect on the term spread, as we saw in figure 5; it explains 11 percent of the variance of the term spread at business-cycle frequencies. This value is surprisingly low, and suggests that even though shifts in risk aversion are clearly important for generating a good fit of the model in terms of the marginal likelihood, movements in the term spread over the business cycle come largely from shifts in monetary policy.

A second interesting result is that I find that shifts in investment-specific technology are far less

[16] The variance decomposition is calculated using a spectral decomposition of the state-space form of the model. Specifically, since the structural shocks are orthogonal, the spectral density of the endogenous variables is equal to the sum of the densities obtained when each shock is turned on individually. Calculating variance shares over certain frequencies then simply requires integrating the density over those frequencies. I numerically integrate by calculating the spectral density at 100 increments between wavelengths of 6 and 32 quarters.

important for driving investment than JPT did. They estimated that 83 percent of the variance in investment growth at business-cycle frequencies is driven by the investment technology shock, whereas I estimate a value of only 21 percent. Since our models are essentially identical other than the preferences I use, and the time-varying risk aversion has minimal real effects, the difference must come from the fact that I include interest rates in the estimation.

My estimate for the volatility of the investment technology shock is smaller than that of JPT by a factor of 6. Looking at figure 5, if the investment technology shocks were that more volatile by a factor of 6, they would clearly have the largest effects on the slope and curvature factors of any of the shocks. It is the large effect of the investment-technology shock on interest rates that makes the model place less weight on it.

To test whether it is actually the inclusion of the interest rate data as opposed to the other modifications of the model that explain the difference with JPT, I reestimate the model excluding data on interest rates (other than the one-quarter yield). The bar on the far right in figure 9 shows the variance decomposition for investment growth from the reestimated model. I replicate JPT's result that the majority of the variance of investment growth at business-cycle frequencies is driven by the investment technology shock.[17] In other words, when the model is augmented to match the behavior of bond prices, JPT's primary result about the importance of the shock to investment technology no longer holds.

Figure 9 shows that including asset prices in the estimation of DSGE models can have major effects on estimated variance decompositions. Asset prices are particularly useful in helping to estimate structural models because they encode investors' expectations. In this case, they help reject the hypothesis that shocks to investment technology are the key driver of the business cycle.

Table 4 gives one-quarter-ahead variance decompositions for output, consumption, and investment growth. This decomposition is useful for understanding whether any of these variables would be powerful asset pricing factors. Specifically, in a world where consumption followed a random walk and households had Epstein–Zin preferences with constant relative risk aversion, consumption

[17]While I do not replicate JPT's results exactly, I do also find in the reestimated model that the investment technology shock is important for more variables than just investment itself. It accounts for 30 percent of the variation in output growth and 20 percent of consumption growth. Moreover, in the reestimated model the investment shock accounts for 41 and 48 percent of the variance of the one-quarter nominal interest rate and the term spread, respectively. It is this latter result that explains the divergence between the results when interest rates are included and excluded.

growth would be perfectly correlated with the SDF, so it would price assets in the economy.

What table 4 shows is that consumption, output, and investment growth are all only weakly correlated with the SDF at the one-quarter horizon—they all have correlations less than 16 percent. So asset pricing with only consumption growth would work poorly in this economy. Many asset-pricing studies with Epstein–Zin preferences include both consumption growth and the return on the stock market as pricing factors. If we believe that the stock market is a claim on aggregate capital, then table 3 shows that it will do little to help with asset pricing as it is also only weakly correlated with the SDF.

8 Conclusion

This paper studies bond pricing in a medium-scale New-Keynesian model with a time-varying price of risk. I show that the model can generate a large and volatile term premium. The term premium is driven by the combination of two factors—a negative response of interest rates to positive technology shocks and variation in risk aversion. Removing either of these effects eliminates the model's ability to match the magnitude of the term premium.

While shocks to risk aversion and technology determine average asset returns, they have only weak effects on real variables at business-cycle frequencies. The covariance of asset returns with real variables over the business cycle is therefore unimportant for determining average returns. Furthermore, while risk aversion is estimated to be highly volatile and to be an important determinant of the dynamics of the term spread, it has almost no effects on the real economy. This model thus suggests that there is a separation between the price of risk in financial markets and the real economy.

On the other hand, including asset prices in estimation has large effects on estimated parameters. Bond prices encode expected future interest rates, so they are very informative about parameters that affect how interest rates respond to shocks. This paper thus argues that even though shocks that affect financial risk premia may have limited real effects, including financial data in the estimation is still important for correctly identifying the underlying drivers of the economy.

References

Ang, Andrew and Monika Piazzesi, "No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables," *Journal of Monetary Economics*, 2003, *50*, 745–787.

Bansal, Ravi, Dana Kiku, and Amir Yaron, "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 2012, *1(1)*, 183–221.

Basu, Susanto, John G. Fernald, and Miles S. Kimball, "Are Technology Improvements Contractionary?," *American Economic Review*, 2006, *96(5)*, 1418–1448.

Bekaert, Geert, Seonghoon Cho, and Antonio Moreno, "new Keynesian Macroeconomics and the Term Structure," *Journal of Money, Credit ans Banking*, 2010, *42(1)*, 33–62.

Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher, "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 2001, *91(1)*, 149–166.

Campanale, Claudio, Rui Castro, and Gian Luca Clementi, "Asset Pricing in a Production Economy with Chew-Dekel Preferences," *Review of Economic Dynamics*, 2010, *13(2)*, 379–402.

Campbell, John Y. and Robert J. Shiller, "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, 1991, *58(3)*, 495–514.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin, "DSGE Models for Monetary Policy Analysi," in "Handbook of Monetary Economics," North-Holland, 2011.

Clarida, Richard, Jordi Galí, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 1999, *37(4)*, 1661–1707.

Cochrane, John H., "Discount Rates," *Journal of Finance*, 2011, *66(4)*, 1047–1108. AFA Presidential Address.

—— **and Monika Piazzesi**, "Bond Risk Premia," *American Economic Review*, 2005, *95* (1), 138–160.

Dew-Becker, Ian, "A Model of Time-Varying Risk Premia with Habits and Production," 2011. Working paper.

—, "Essentially affine approximations for economic models," 2012. Working paper.

— **and Stefano Giglio**, "Asset Pricing in the Frequency Domain: Theory and Empirics," 2013. Working paper.

Doh, Taeyoung, "Yield Curve in an Estimated Nonlinear Macro Model," *Journal of Economic Dynamics & Control*, 2011.

Duffee, Gregory R., "Term Premia and Interest Rate Forecasts in Affine Models," *Journal of Finance*, 2002, *57(1)*, 405–443.

Evans, Martin, "Real Rates, Expected Inflation and Inflation Risk Premia," *Journal of Finance*, 1998, *53(1)*, 187–218.

Fama, Eugene F. and Kenneth R. French, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 1989, *25(1)*, 23–49.

Francis, Neville and Valerie A. Ramey, "Measures of Per Capita Hours and Their Implications for the Technology-Hours Debate," *Journal of Money, Credit, and Banking*, 2009, *41(6)*, 1071–1097.

Gilchrist, Simon and Egon Zakrajsek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102(4)*, 1692–1720.

Gourio, Francois, "Disaster Risk and Business Cycles," *American Economic Review*, 2012, *102(6)*, 2734–2766. Working Paper.

—, "Credit Risk and Disaster Risk," *American Economic Journal: Macroeconomics*, 2013, *5(3)*, 1–34. Working paper.

Gurkaynak, Refet S., Brian Sack, and Eric Swanson, "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models," *American Economic Review*, 2005, *95(1)*, 425–436.

Guvenen, Fatih, "A Parsimonious Macroeconomic Model for Asset Pricing," *Econometrica*, 2009, *77(6)*, 1711–1750.

Hamilton, James D. and Cynthia Wu, "Identification and Estimation of Gaussian Affine-Term-Structure Models," *Journal of Econometrics*, 2012. Working Paper.

Hansen, Lars Peter and Ravi Jagannathan, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 1991, *99(2)*, 225–262.

Jermann, Urban J., "Asset Pricing in Production Economies," *Journal of Monetary Economics*, 1998, *41* (2), 257–275.

Justiniano, Alejandro, Giorgio Primiceri, and Andrea Tambalotti, "Investment Shocks and Business Cycles," *Journal of Monetary Economics*, 2010, *57(2)*, 132–145.

Melino, Angelo and Alan X. Yang, "State-Dependent Preferences can Explain the Equity Premium Puzzle," *Review of Economic Dynamics*, 2003, *6(4)*, 806–830.

Rudebusch, Glenn D. and Eric T. Swanson, "Examining the Bond Premium Puzzle with a DSGE Model," *Journal of Monetary Economics*, 2008, *55(Supplement 1)*, S111–S126.

____ **and** ____ , "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks," *American Economic Journal: Macroeconomics*, 2012, *4(1)*, 105–143. Federal Reserve Bank of San Francisco Working Paper 2008-31.

Smets, Frank and Raf Wouters, "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, 2003, *1(5)*, 1123–1175.

Swanson, Eric T., "Risk Aversion and the Labor Margin in Dynamic Equilibrium Models," *American Economic Review*, 2011.

Tallarini, Thomas D., "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 2000, *45(3)*, 507–532.

Tanaka, Hiroatsu, "Monetary Policy Regimes and the Term Structure of Interest Rates with Recursive UItility," 2012. Working paper.

van Binsbergen, Jules H., Jesus Fernandez-Villaverde, Ralph S.J. Koijen, and Juan F. Rubio-Ramirez, "The term structure of interest rates in a DSGE model with recursive preferences," *Journal of Monetary Economics*, 2012, *59(3)*, 634–648.

Woodford, Michael, *Interest and Prices*, Princeton University Press, Princeton, NJ, 2003.

Table 1. Priors and posterior modes

			Priors			Posterior			
		Description	Distribution	Mean	Std. Dev.	Mode	Std. Dev.	5%	95%
1	ι_p	Price indexation	Beta	0.3	0.1	0.47	0.03	0.42	0.51
2	ι_w	Wage indexation	Beta	0.3	0.1	0.30	0.02	0.28	0.33
3	η	Habit parameter	Beta	0.5	0.2	0.13	0.01	0.11	0.15
4	S	Investment adjustment costs	Gamma	5	3	0.14	0.02	0.11	0.17
5	f_π	Taylor rule inflation	Normal	1.6	0.3	1.71	0.04	1.65	1.77
6	f_y	Taylor rule output gap	Gamma	0.15	0.03	0.12	0.00	0.12	0.12
7	$f_{\Delta y}$	Taylor rule output gap growth	Normal	0.15	0.1	0.24	0.01	0.23	0.26
8	ϱ_R	Monetary policy AR	Beta	0.75	0.1	0.97	0.00	0.97	0.97
9	ϱ_z	Technology MA	U[-1,1]	0.5	0.29	-0.04	0.01	-0.05	-0.03
10	ϱ_g	Government spending AR	Beta	0.6	0.2	0.95	0.00	0.95	0.96
11	ϱ_μ	Investment technology AR	Beta	0.6	0.2	0.50	0.02	0.48	0.53
12	$\varrho_{\lambda p}$	Price markup AR	Beta	0.6	0.2	0.89	0.01	0.87	0.91
13	$\varrho_{\lambda w}$	Wage markup AR	Beta	0.6	0.2	0.41	0.01	0.39	0.43
14	ϱ_b	Consumption demand AR	Beta*0.85	0.43	0.085	0.78	0.01	0.77	0.80
15	ϱ_α	Risk aversion AR	Beta	0.8	0.1	0.90	0.00	0.89	0.90
16	σ_R	MP shock vol.	IG(1)	0.1	1	0.14	0.00	0.14	0.14
17	σ_z	Neutral tech. shock vol.	IG(1)	1.5	2	1.15	0.03	1.10	1.20
18	σ_g	Gov't spending vol.	IG(1)	0.5	1	0.68	0.01	0.66	0.69
19	σ_μ	Investment tech. vol.	IG(1)	0.5	0.5	0.94	0.02	0.90	0.98
20	$\sigma_{\lambda p}$	Price markup vol.	Normal	0.14	0.05	0.14	0.01	0.13	0.15
21	$\sigma_{\lambda w}$	Wage markup vol.	Normal	0.2	0.05	0.32	0.01	0.31	0.32
22	σ_b	Demand shock vol.	IG(1)	0.001	0.003	0.0031	0.0001	0.0029	0.0033
23	$\sigma_{\pi*}$	Inflation target vol.	IG(1)	0.2	0.2	0.19	0.01	0.18	0.20
24	σ_α/α	RRA volatility/RRA mean	U[0,1]	0.5	0.29	0.34	0.01	0.32	0.36
25	$\sigma_{\pi,z}$	corr(pi*,tech.)	U[-1,1]	0	0.68	-0.15	0.01	-0.16	-0.14
26	α_{SS}	Mean risk aversion	Normal	7	5	23.46	0.74	22.30	24.71
27	σ_{yields}	Bond measurement errors (bp)	IG(1)	20	20	16.72	0.63	15.73	17.78
28	R*	Steady-state risk-free rate	U[-1,1]	0	0.58	0.00	0.00	0.00	0.00
29	θ_p	Price markup MA	Beta	0.7	0.15	0.37	0.01	0.35	0.39
30	θ_w	Wage markup MA	Beta	0.7	0.15	0.36	0.02	0.33	0.38

Note: Priors, posterior mode, and percentiles of the posterior distribution from the benchmark model.

Table 2. Fitting errors

	1-quarter	1-year	2-year	3-year	4-year	5-year	10-year	20-year
3-factor model	0	9.44	5.90	5.75	10.60	18.13	8.73	14.38
2-factor model	0	23.27	25.07	21.01	19.56	18.65	15.95	26.34
Benchmark model	0	16.72	16.72	16.72	16.72	16.72	16.72	16.72
Constant RRA	0	24.18	24.18	24.18	24.18	24.18	24.18	24.18
Constant π*	0	24.21	24.21	24.21	24.21	24.21	24.21	24.21

Note: Fitting errors measured in annualized basis points. The model-based estimates use the posterior modal estimate for the standard deviation. The 1-quarter and 10-year errors are constrained to equal zero in the structural model. The errors from PCA are the standard deviations of the residuals from regressions on the bond yields on their first three principal components.

Table 3. Log marginal likelihoods

	Likelihood	Bayes Factor
Benchmark	3796.8	
Constant RRA	3745.3	51.50
Constant π*	3707.6	-89.20

Table 4. One-quarter ahead variance decompositions

	1	2	3	4	5	6	7	8	9
	SDF	Utility return	Cap. claim	Nom. lev. cap. claim	Real lev. cap. claim	10-year bond return	Output growth	Consumption growth	Investment growth
1 Monetary policy	0.01	0.11	0.29	0.01	0.11	0.02	0.35	0.12	0.30
2 Neutral tech.	0.77	0.02	0.06	0.30	0.31	0.26	0.01	0.00	0.01
3 Gov't spending	0.01	0.05	0.00	0.01	0.04	0.00	0.01	0.05	0.01
4 Investment tech.	0.00	0.04	0.22	0.06	0.08	0.00	0.10	0.05	0.25
5 Price markup	0.00	0.01	0.02	0.03	0.03	0.00	0.01	0.02	0.00
6 Wage markup	0.00	0.01	0.06	0.05	0.02	0.00	0.03	0.01	0.03
7 Time preference	0.00	0.53	0.00	0.05	0.01	0.00	0.05	0.57	0.01
8 Inflation target	0.01	0.15	0.36	0.49	0.12	0.56	0.44	0.15	0.38
9 Risk aversion	0.20	0.07	0.00	0.00	0.28	0.15	0.00	0.04	0.01
Annualized Moments:									
10 Standard deviation	0.27	3.45	0.02	0.39	2.69	21.07	5.12	3.41	10.67
11 Correl. w/ SDF	1.00	-0.39	0.12	0.40	-0.65	-0.53	-0.06	-0.23	0.01
12 Expected return	N/A	0.72	0.00	-0.09	0.95	6.02	N/A	N/A	N/A

Variance decompositions and 95% credible intervals

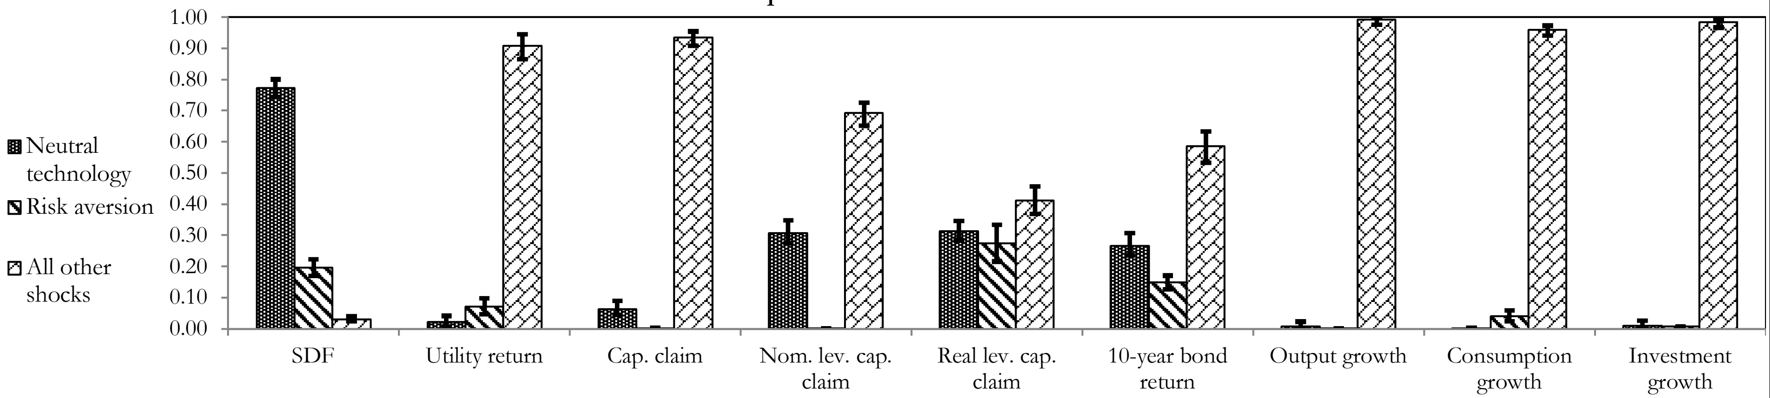

Note: decompositions of one-quarter ahead forecast error. Levered returns for capital and consumption claims assume that the investor finances half the purchase price of the given claim with a 10-year nominally riskless bond. The 10-year return is the one-quarter return from holding a 10-year nominally riskless bond. The moments in rows 10–12 are annualized. The black bars in the figure give the 95 percent credible region based on random draws from the posterior density.

Table 5. Decomposition of innovation to pricing kernel

	SDF	U_C	$\Sigma\Delta U$	$\Sigma\alpha$
Monetary policy	0.021	0.006	0.015	0.000
Neutral tech.	-0.237	0.000	-0.237	0.000
Gov't spending	0.023	0.004	0.019	0.000
Investment tech.	0.002	0.004	-0.002	0.000
Price markup	0.001	-0.002	0.003	0.000
Wage markup	0.002	-0.002	0.003	0.000
Time preference	0.017	0.013	0.004	0.000
Inflation target	-0.031	-0.006	-0.024	0.000
Risk aversion	0.120	0.003	0.001	0.115

Notes: The first column is the effect of each of the shocks on the market pricing kernel. The next three columns decompose the effects into parts coming from news about the marginal utility of consumption, news about future growth in period utility, and news about future risk aversion.

Table 6. Posterior modes for alternative models

		Description	Benchmark		No bonds in estimation		π^* correlated with μ and α shocks		π^* correlated with μ and α shocks; flat prior on ϱ_μ, $\varrho_{\lambda w}$, and $\varrho_{\lambda p}$	
1	ι_p	Price indexation	0.47	0.02	0.4558	0.07	0.45	0.02	0.4524	0.00
2	ι_w	Wage indexation	0.30	0.03	0.225	0.04	0.32	0.03	0.31	0.05
3	η	Habit parameter	0.13	0.03	0.0707	0.02	0.14	0.02	0.13	0.08
4	S	Investment adjustment costs	0.14	0.01	1.0593	0.13	0.11	0.01	0.11	0.02
5	f_π	Taylor rule inflation	1.71	0.03	2.5308	0.09	1.74	0.05	1.79	0.16
6	f_y	Taylor rule output gap	0.12	0.00	0.0605	0.00	0.12	0.01	0.12	0.00
7	$f_{\Delta y}$	Taylor rule output gap growth	0.24	0.00	0.15	0.01	0.24	0.01	0.24	0.00
8	ϱ_R	Monetary policy AR	0.97	0.00	0.8277	0.01	0.97	0.00	0.97	0.00
9	ϱ_z	Technology MA	-0.04	0.00	-0.0076	0.01	-0.06	0.01	-0.06	0.00
10	ϱ_g	Government spending AR	0.95	0.00	0.9502	0.00	0.95	0.00	0.95	0.00
11	ϱ_μ	Investment technology AR	0.50	0.03	0.8171	0.01	0.51	0.03	0.51	0.05
12	$\varrho_{\lambda p}$	Price markup AR	0.89	0.01	0.8404	0.03	0.89	0.01	0.89	0.01
13	$\varrho_{\lambda w}$	Wage markup AR	0.41	0.02	0.7183	0.02	0.41	0.01	0.38	0.00
14	ϱ_b	Consumption demand AR	0.78	0.01	0.8128	0.01	0.78	0.01	0.78	0.01
15	ϱ_α	Risk aversion AR	0.90	0.04	0.8461	0.01	0.89	0.00	0.89	0.00
16	σ_R	MP shock vol.	0.14	0.00	0.1479	0.00	0.14	0.00	0.14	0.00
17	σ_z	Neutral tech. shock vol.	1.15	0.02	0.8795	0.02	1.22	0.02	1.22	0.00
18	σ_g	Gov't spending vol.	0.68	0.03	0.6741	0.01	0.68	0.01	0.68	0.00
19	σ_μ	Investment tech. vol.	0.94	0.02	1.844	0.02	0.87	0.01	0.87	0.00
20	$\sigma_{\lambda p}$	Price markup vol.	0.14	0.00	0.1552	0.01	0.14	0.00	0.14	0.00
21	$\sigma_{\lambda w}$	Wage markup vol.	0.32	0.01	0.2183	0.00	0.32	0.01	0.32	0.00
22	σ_b	Demand shock vol.	0.0031	0.0003	0.0025	0.0002	0.0032	0.0001	0.0032	0.0000
23	σ_{π^*}	Inflation target vol.	0.19	0.01	0.0011	0.00	0.18	0.00	0.18	0.00
24	σ_α/α	RRA volatility/RRA mean	0.34	0.01	0.2845	0.03	0.36	0.01	0.37	0.02
25	$\sigma_{\pi,z}$	corr(pi*,tech.)	-0.15	0.00	N/A	N/A	-0.18	0.00	-0.19	0.00
26	α_{SS}	Mean risk aversion	23.46	0.34	6.9292	2.14	21.28	0.78	20.92	0.87
27	σ_{yields}	Bond measurement errors (bp)	16.72	0.64	N/A	N/A	16.09	0.56	17.11	0.00
28	R*	Steady-state risk-free rate	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
29	θ_p	Price markup MA	0.37	0.01	0.3501	0.01	0.36	0.01	0.36	0.00
30	θ_w	Wage markup MA	0.36	0.01	0.49	0.02	0.37	0.01	0.34	0.00
31	$\sigma_{\pi,\mu}$	corr(pi*,inv.)	N/A	N/A	N/A	N/A	-0.05	0.01	-0.05	0.00
32	$\sigma_{\pi,\alpha}$	corr(pi*,RRA)	N/A	N/A	N/A	N/A	0.06	0.01	0.06	0.00

Note: Priors, posterior mode, and percentiles of the posterior distribution from the benchmark model.

Figure 1. Data series for estimation

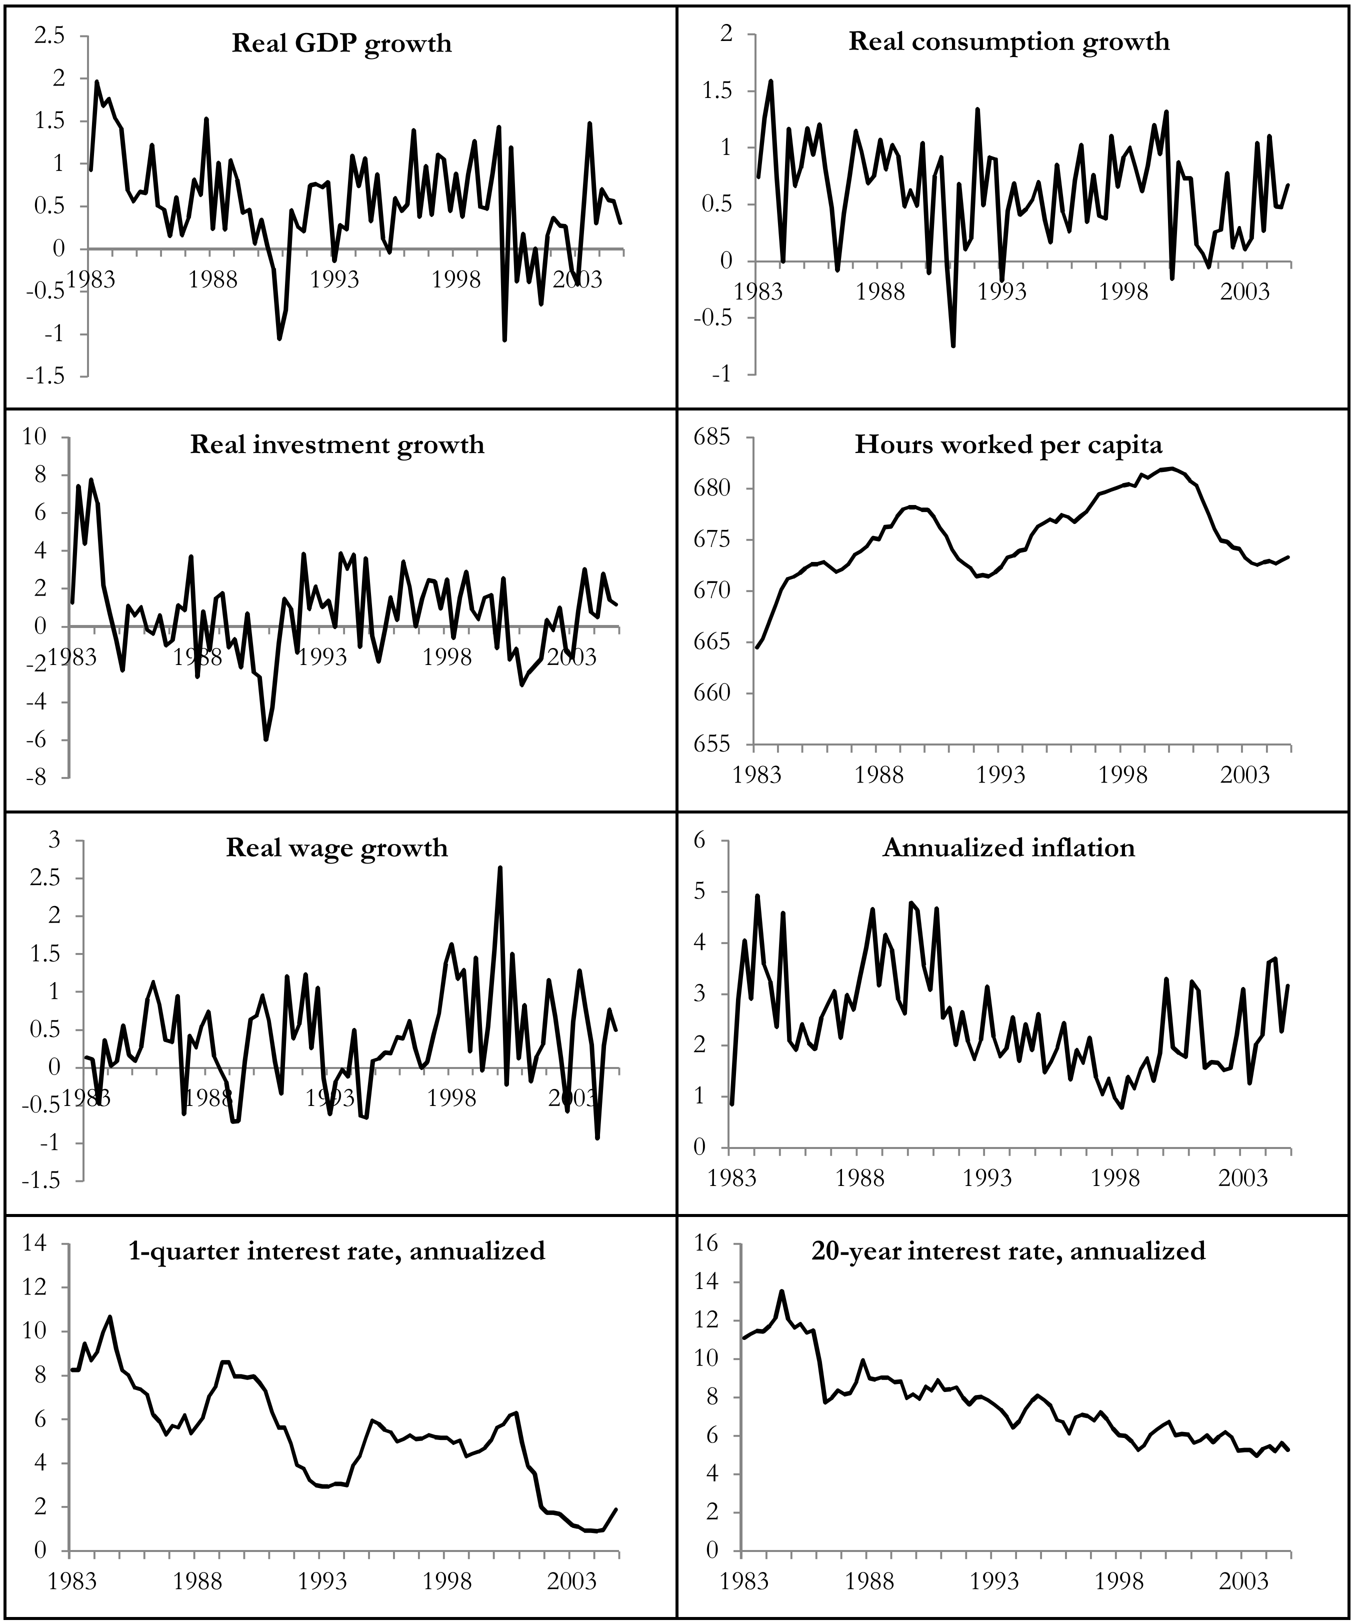

Note: No variables are detrended. GDP, consumption, and investment are obtained from the BEA. Compensation per hour, and inflation are obtained from the BLS. Hours worked is obtained from Valerie Ramey's website. The one-quarter yield is the Fama risk-free rate. The ten-year yield is from Gurkaynak, Sack, and Wright (2006).

Figure 2. Bond yield errors

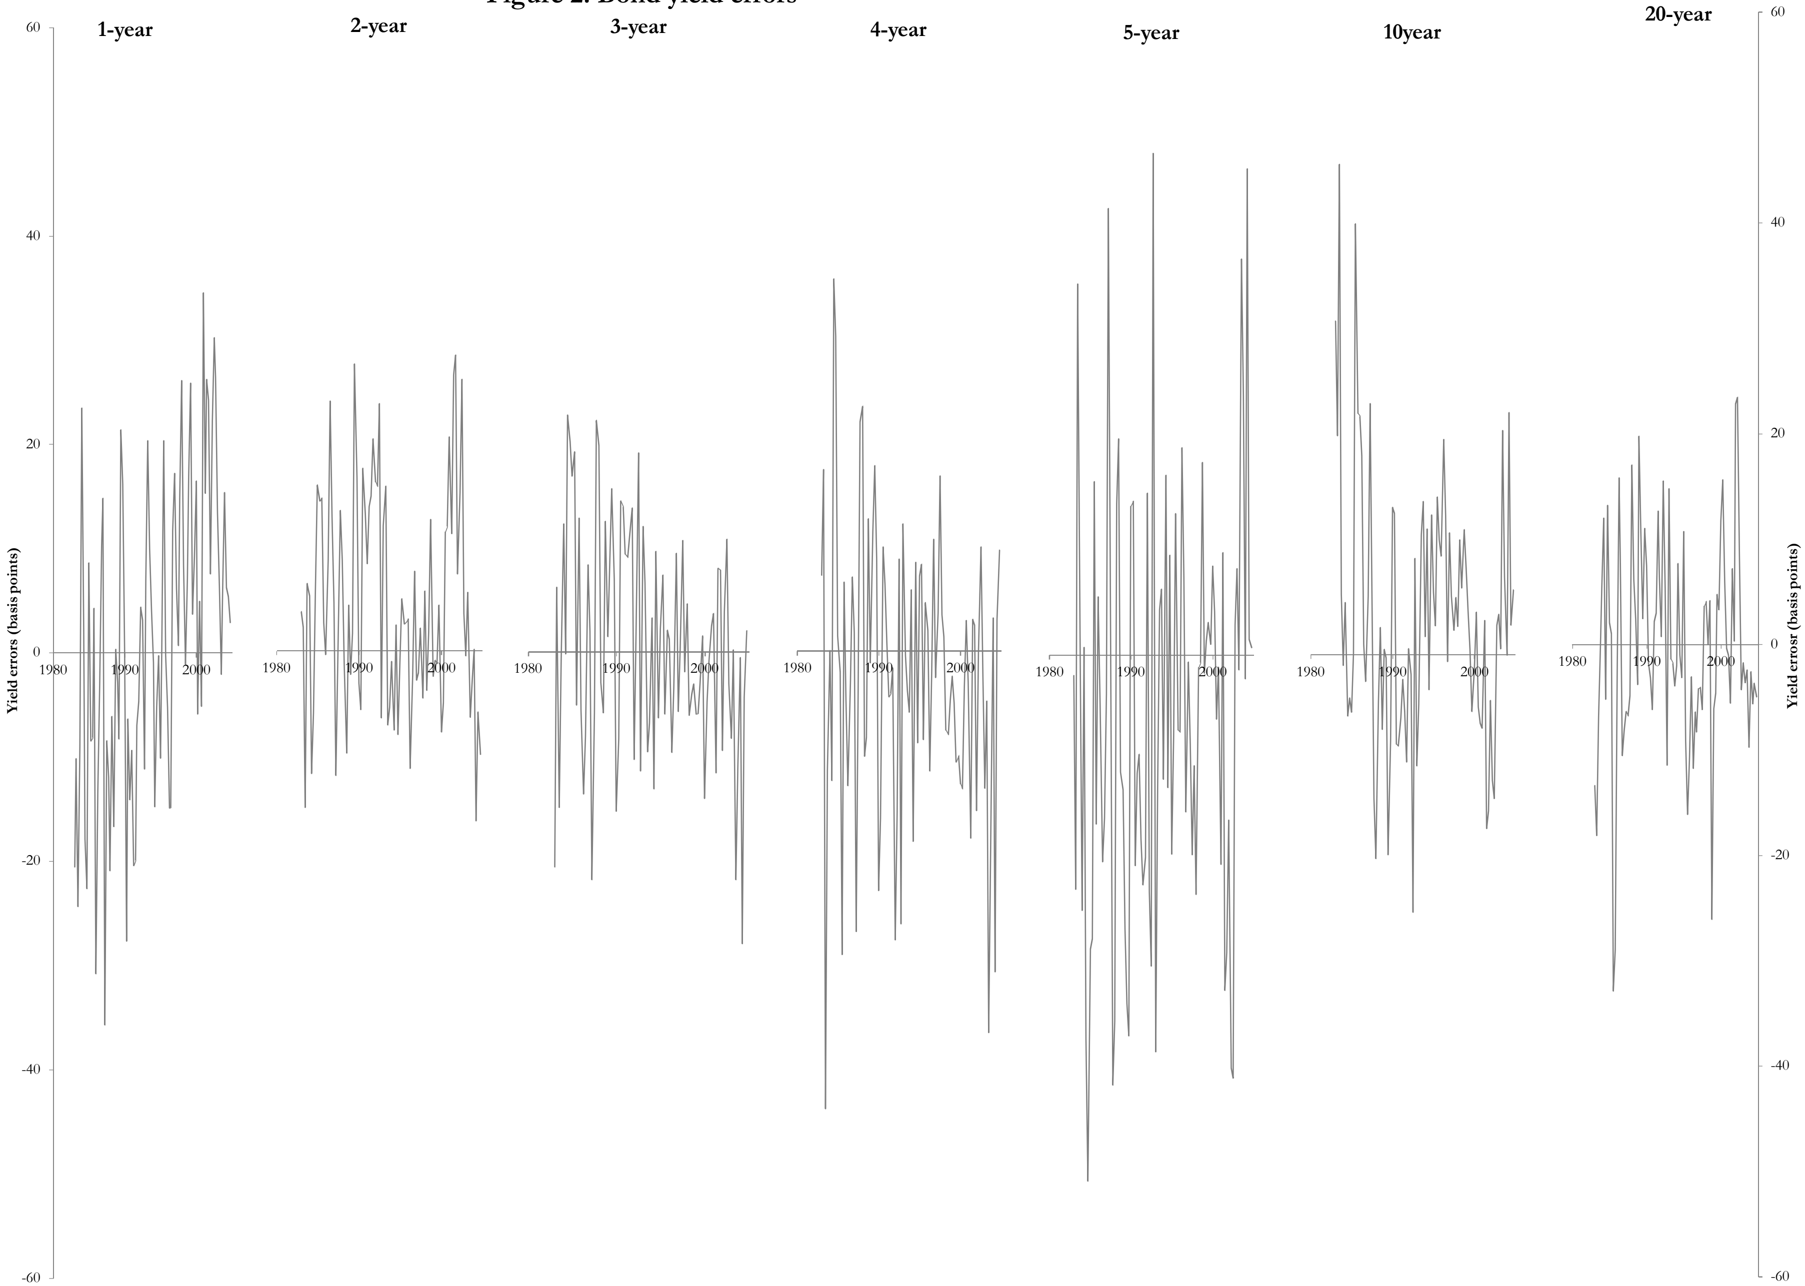

Note: each axis plots the measurement errors in basis points for one of the bond yields. Errors are measured from the Kalman-filtered estimates at the posterior mode.

Figure 3. Steady-state nominal bond yields

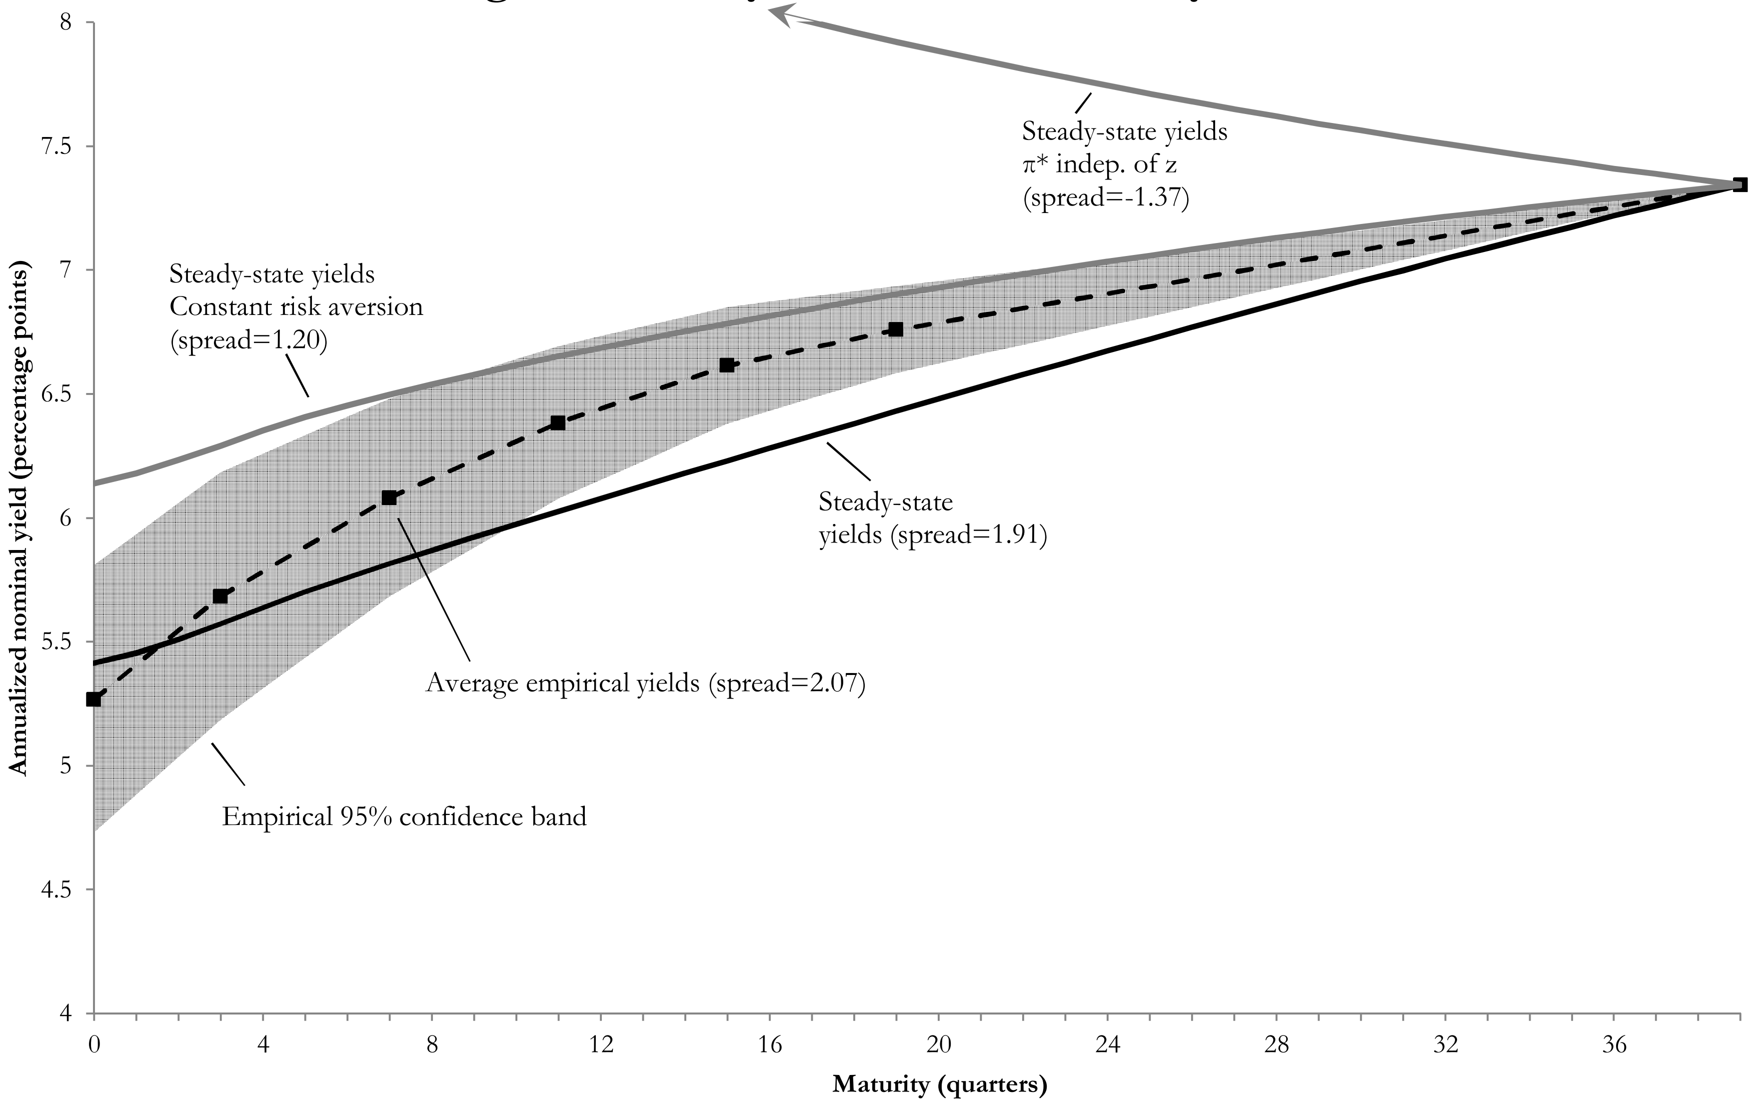

Note: The solid black line gives the yield curve at the model's steady state; the grey lines are for the model with constant risk aversion and where the inflation target is unaffected by shocks to labor-neutral technology. The other parameters are not reestimated. All the lines are normalized to match the 10-year yield exactly, so the plot measures steady-state spreads. Boxes are average sample yields. The grey area is the 95% confidence band for the average yields relative to the 10-year yield, calculated using the Newey–West method with 6 lags. The solid black line gives the yield curve at the model's steady state, normalized to match the 10-year yield exactly.

Figure 4. Expected returns and the term premium

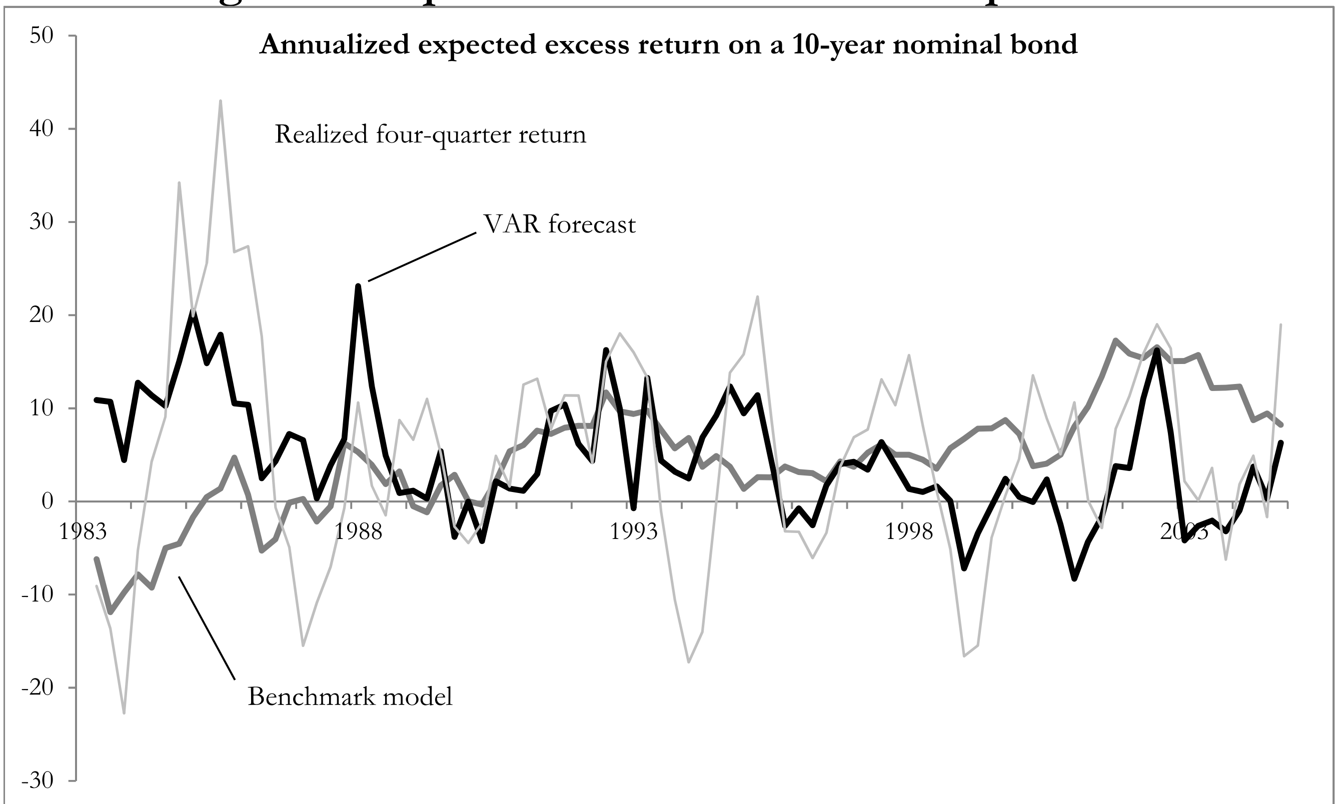

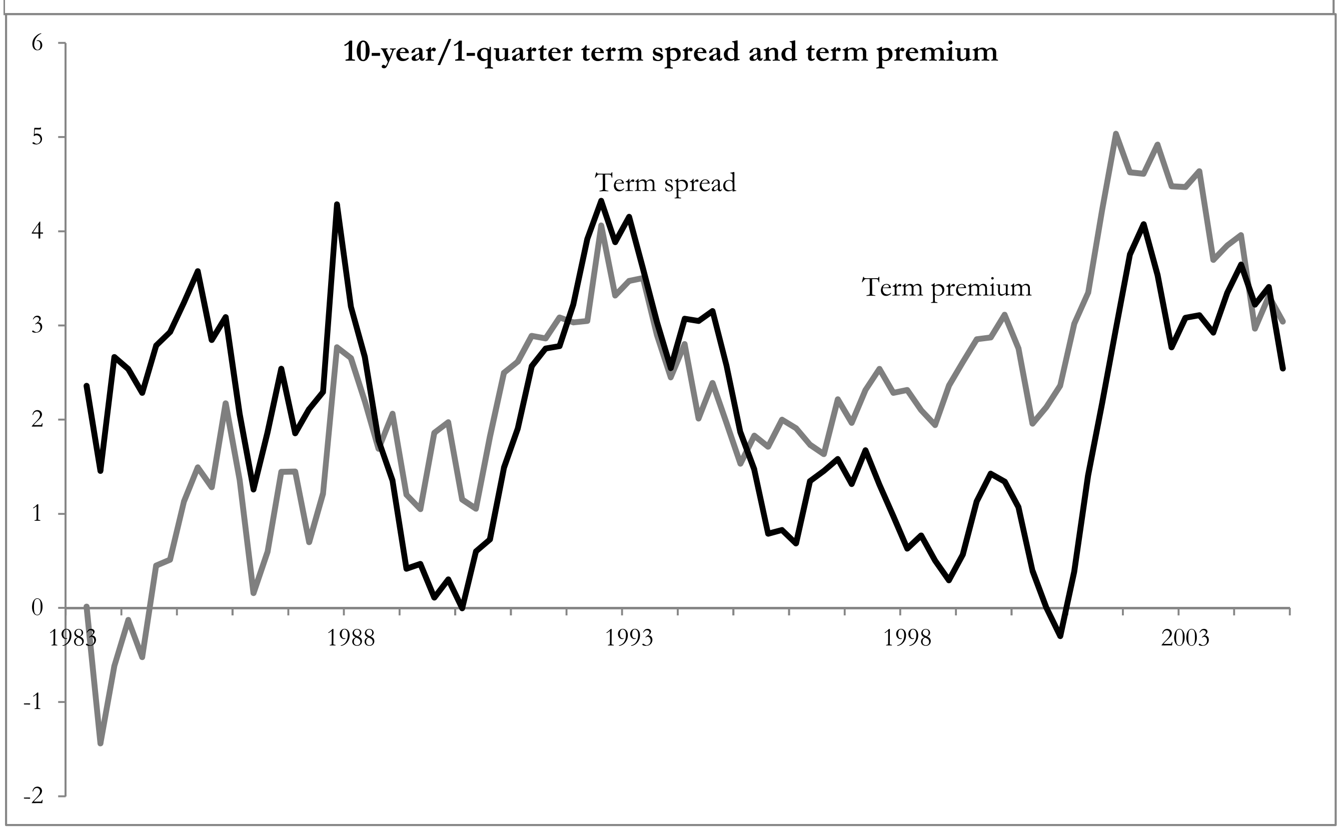

Note: The top panel plots expected returns on a 10-year nominal bond over the following quarter in excess of the one-quarter nominal interest rate, annualized. The non-structural forecast is from a VAR(1) using the 1-quarter, 5-year, and 10-year bond yields. The thin gray line is the realized four-quarter return on a 10-year Treasury bond. The term premium in the bottom panel is defined as the gap between the 10-year yield and the average of expected 1-quarter yields over the next ten years.

Figure 5. Responses of term structure factors to orthogonalized shocks

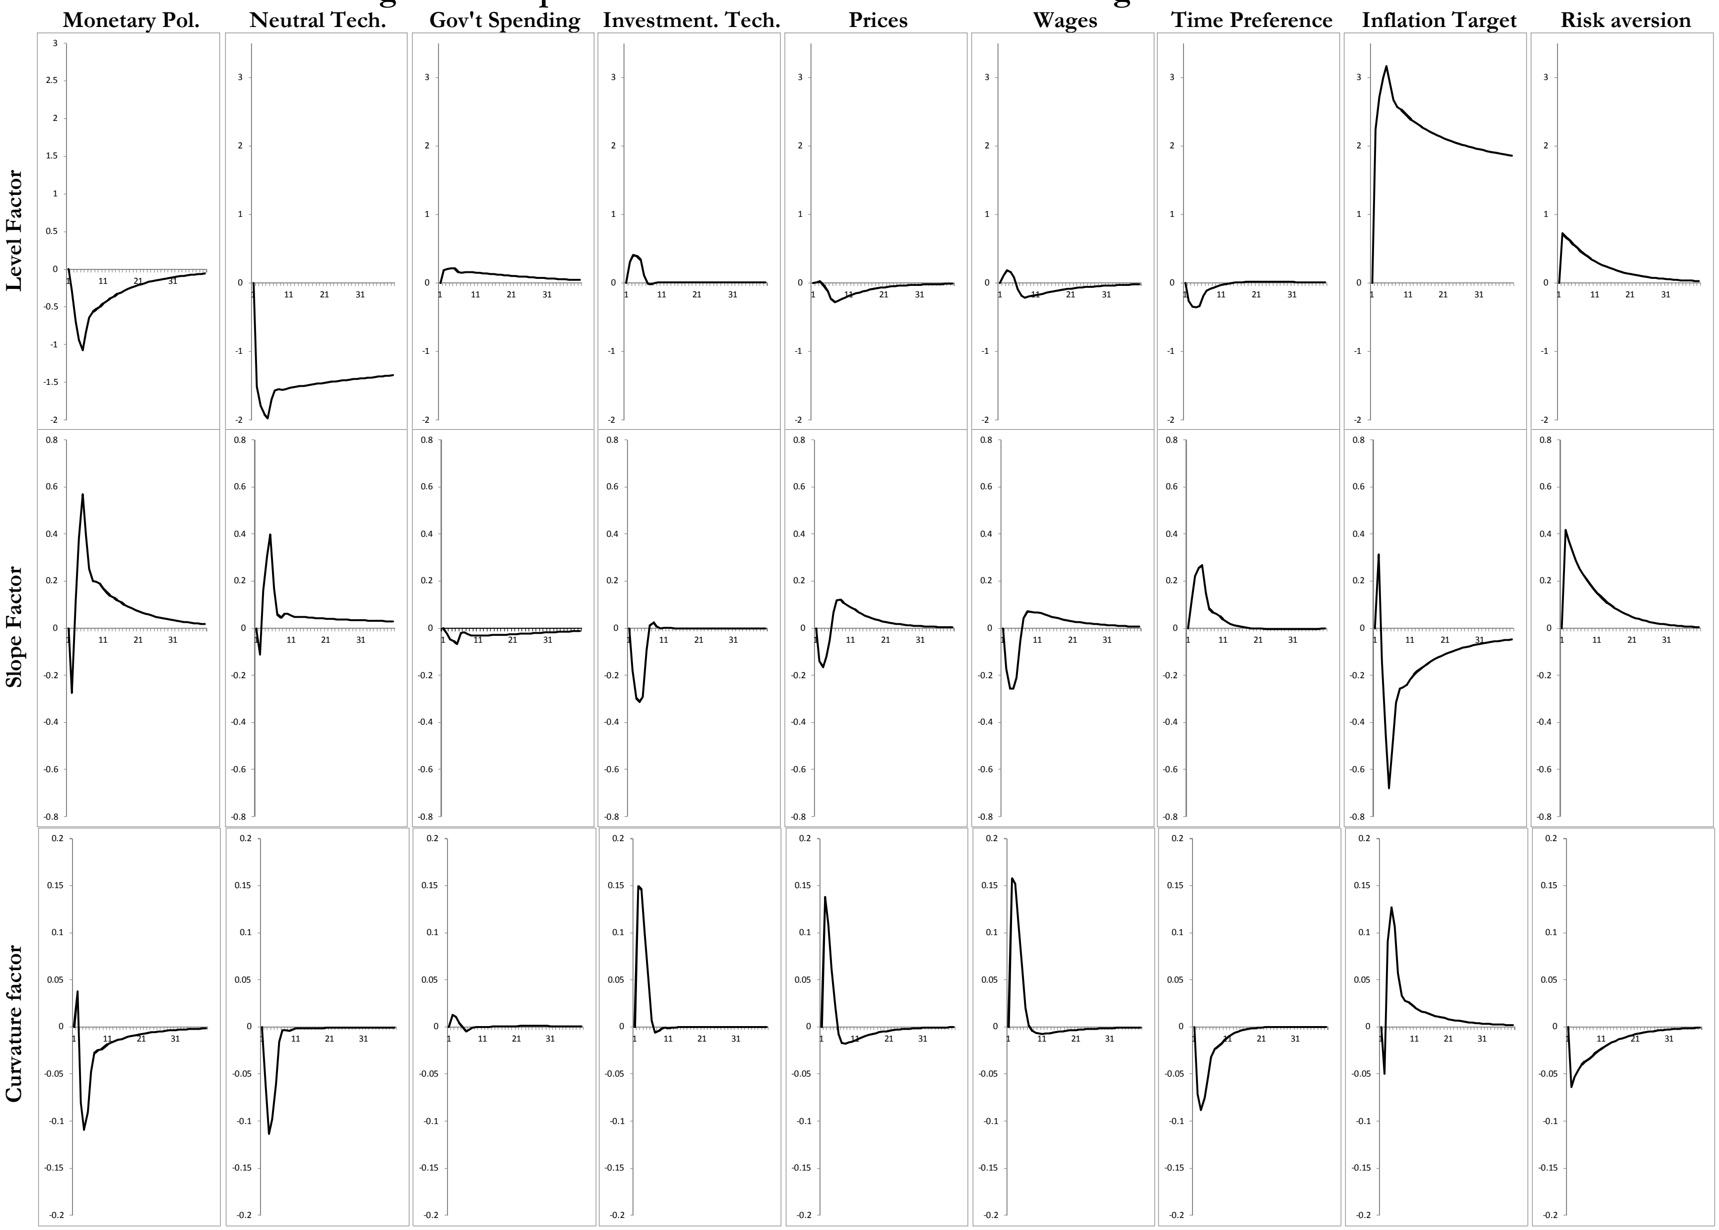

Note: responses of each of the term structure factors to the orthogonal structural shocks. Specifically, the inflation target in this plot is not affected by shocks to labor-neutral technology. The level factor is the average the 1-quarter, 5-year, and 10-year yields. The slope factor is the gap between the 10-year and 1-quarter yields. Curvature is the sum of the 5-year and 1-year yields minus twice the 3-year yield. The shocks are all unit standard deviations. All scales in each row are identical and are measured in annualized percentage points.

Figure 6. Responses of interest rates to an investment technology shock

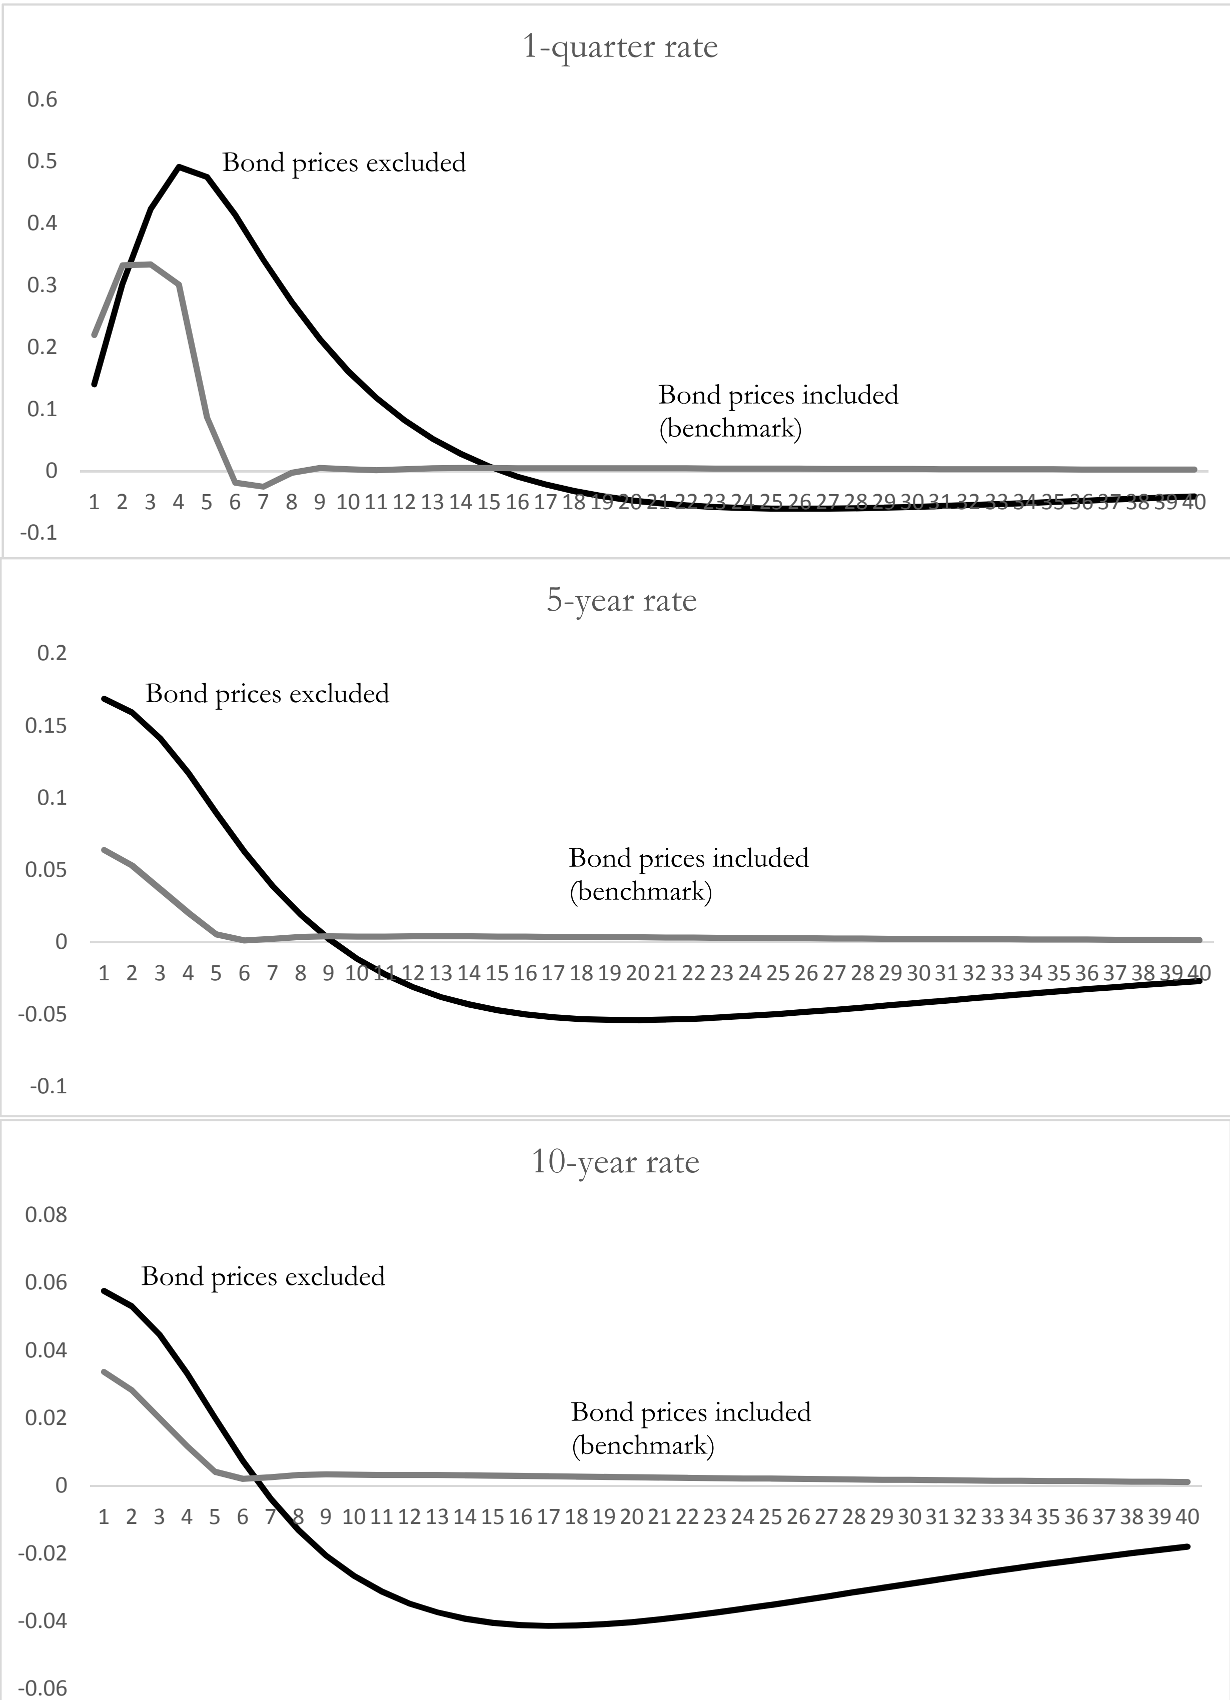

Figure 7. Responses to technology and risk-aversion shocks

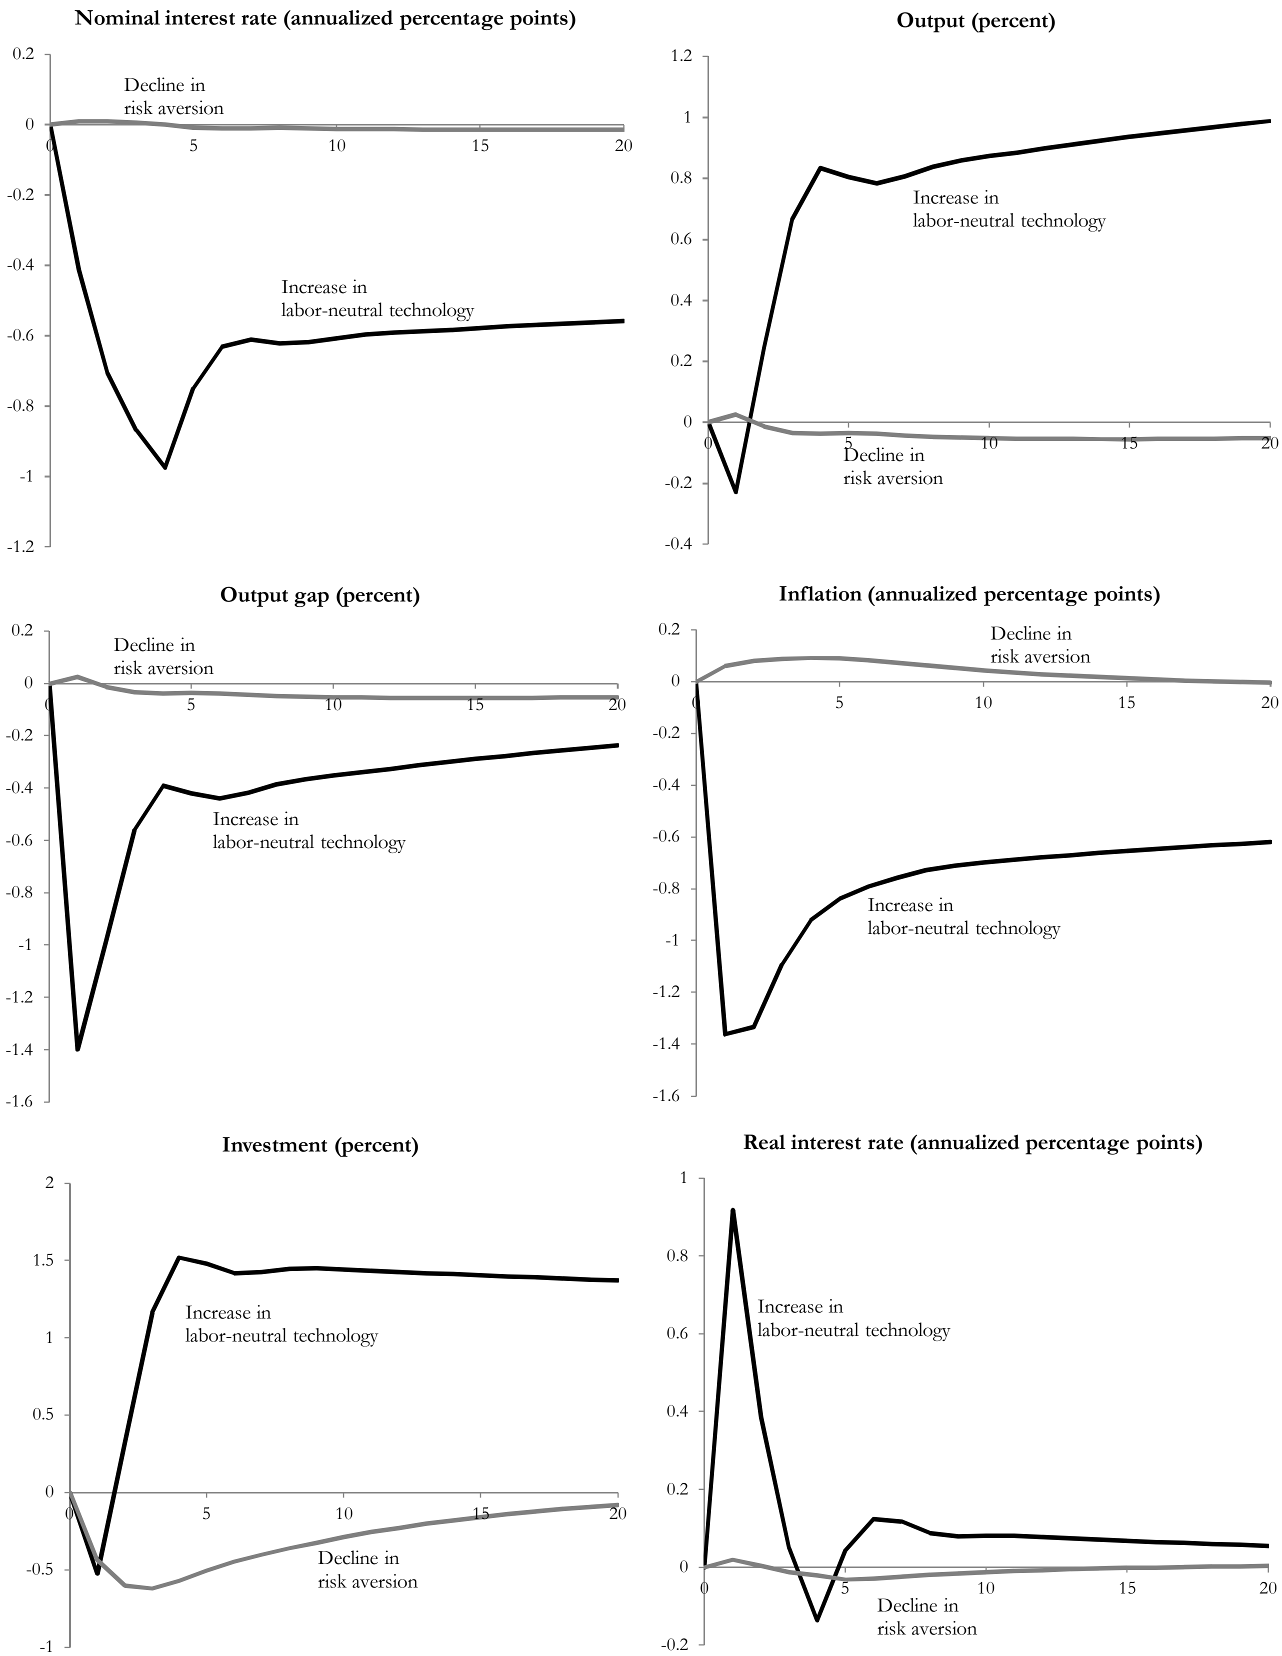

Note: responses in a unit standard deviation increase in labor-neutral technology and a decrease in risk aversion.

Figure 8. The Hansen–Jagannathan bound and the term spread

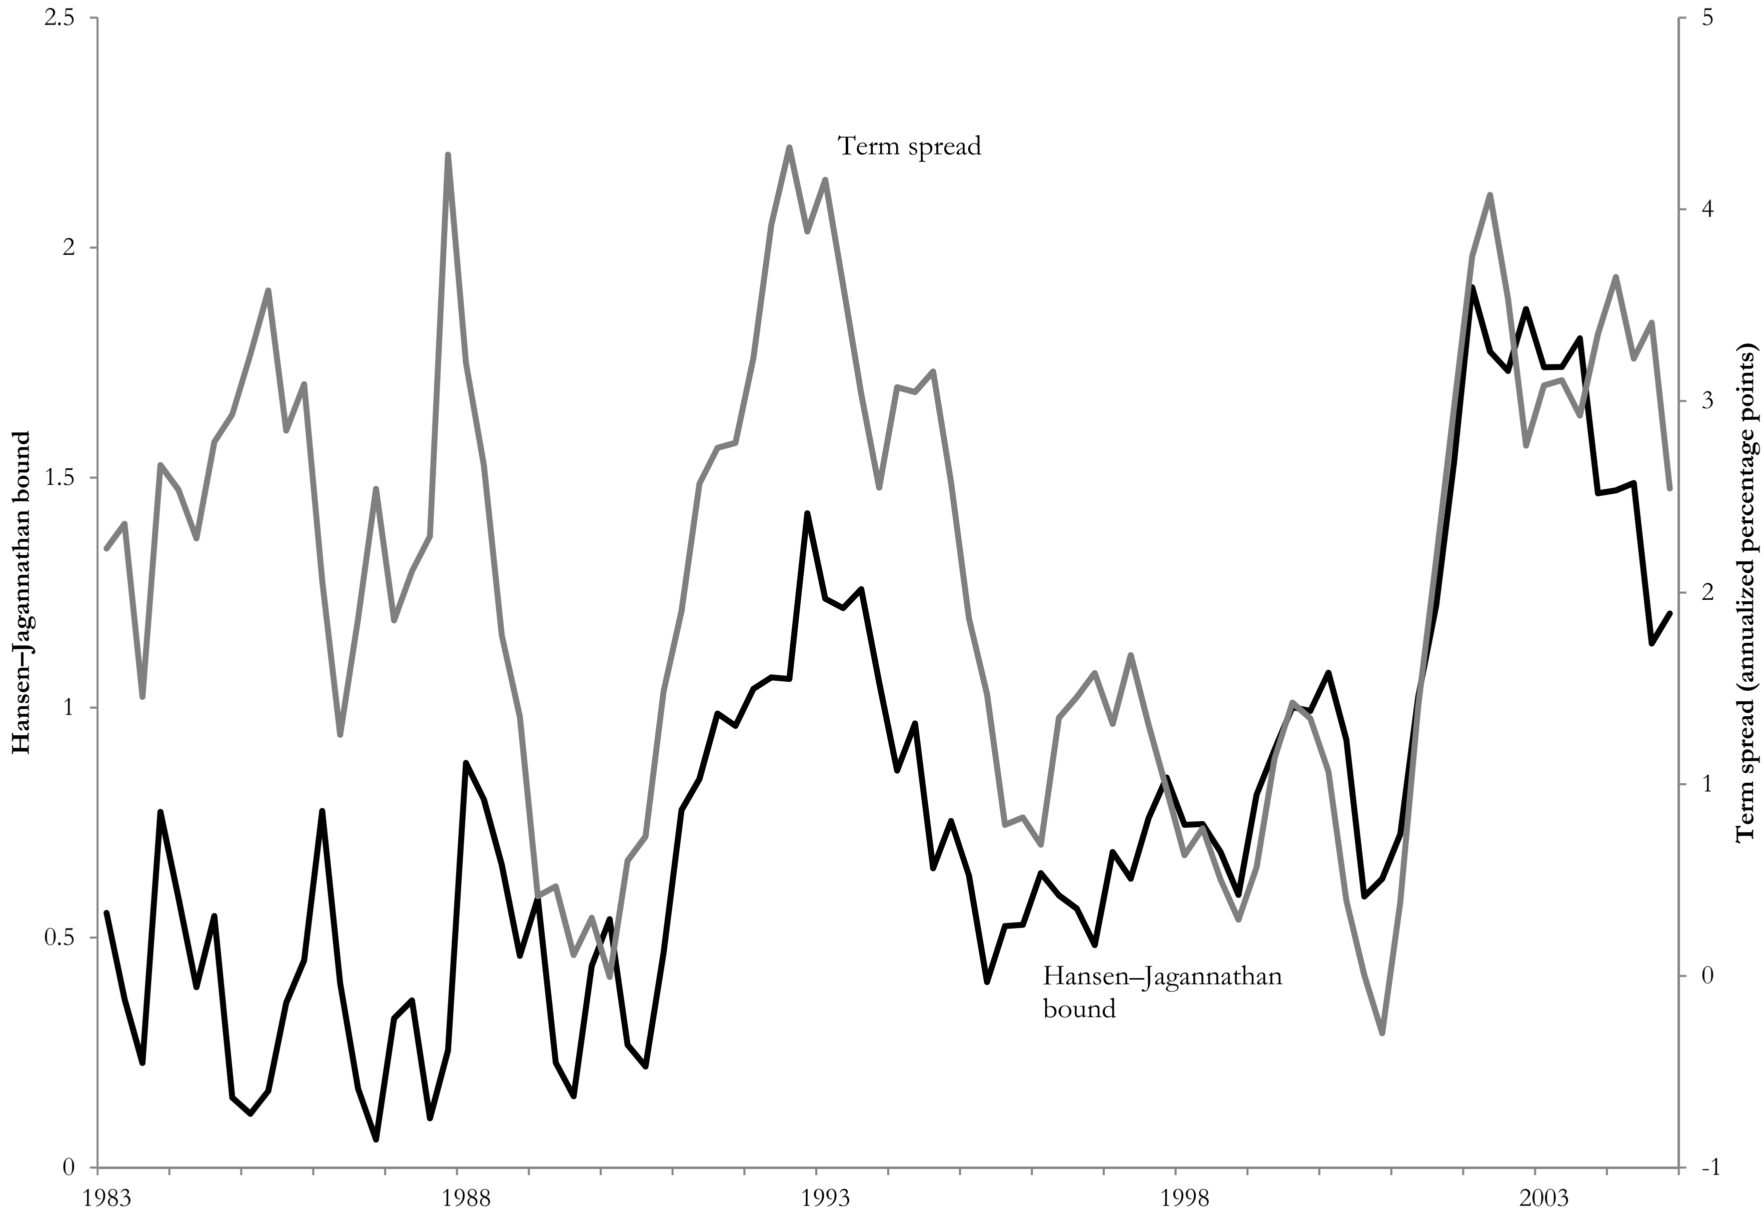

Figure 9. Variance decompositions at business-cycle frequencies

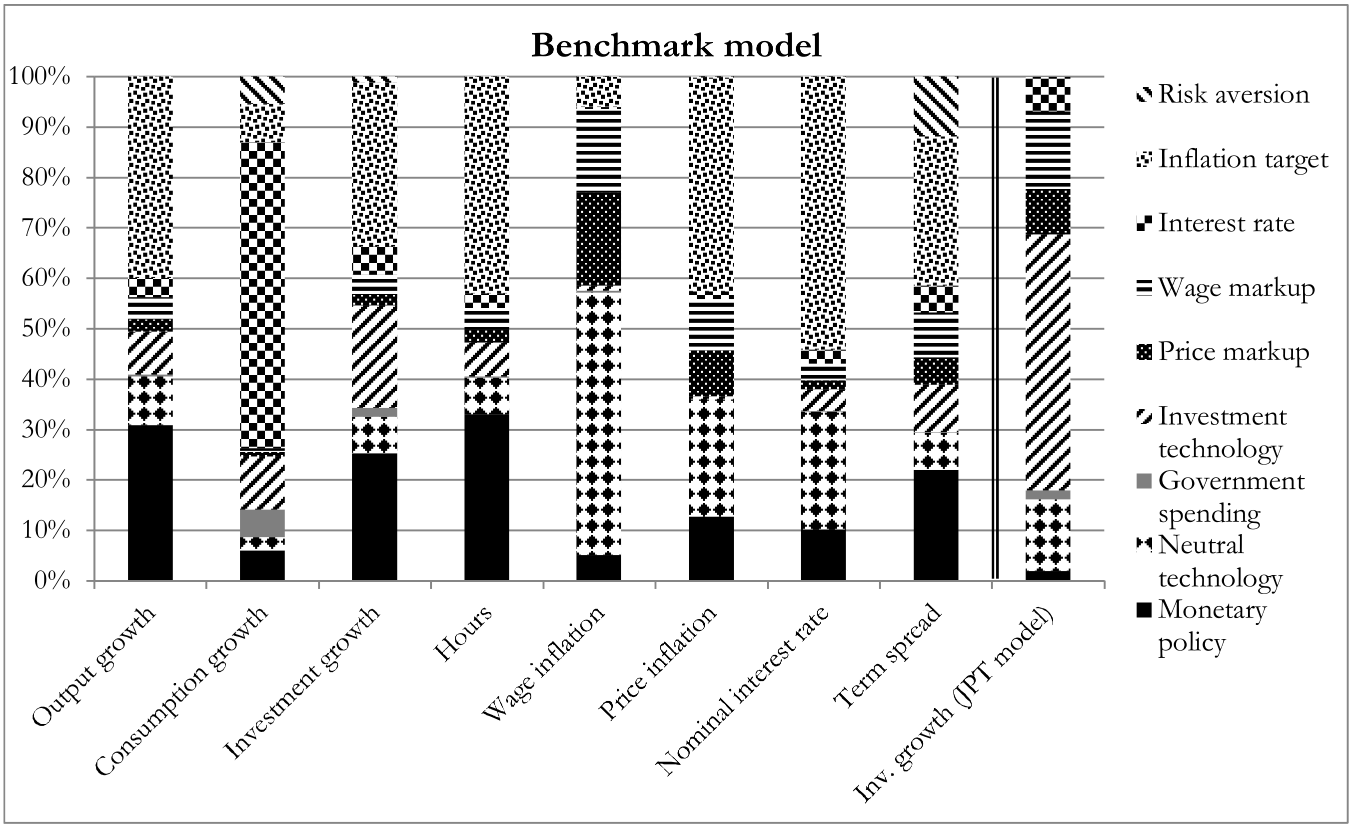

Note: each section of a bar represents a fraction of the variance of one of the series at frequencies between 6 and 32 quarters. The far-right column gives a variance decomposition for investment growth in the JPT model (bonds excluded from the estimation; constant risk aversion and inflation target) instead of the benchmark model used in the left-hand columns.