



Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmeExamining the bond premium puzzle with a DSGE model[☆]Glenn D. Rudebusch, Eric T. Swanson^{*}

Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 19105, USA

ARTICLE INFO

Article history:

Received 28 November 2007

Received in revised form

3 July 2008

Accepted 12 July 2008

Available online 29 July 2008

JEL classification:

E43

G12

Keywords:

Yield curve

Term premium

Bond pricing

ABSTRACT

The basic inability of standard theoretical models to generate a sufficiently large and variable nominal bond risk premium has been termed the “bond premium puzzle.” We show that the term premium on long-term bonds in the canonical dynamic stochastic general equilibrium (DSGE) model used in macroeconomics is far too small and stable relative to the data. We find that introducing long-memory habits in consumption as well as labor market frictions can help fit the term premium, but only by seriously distorting the DSGE model’s ability to fit other macroeconomic variables, such as the real wage; therefore, the bond premium puzzle remains.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Understanding the risk premium on long-term bonds is of clear practical importance. For example, central banks around the world use the yield curve to help assess market expectations about future interest rates and inflation as well as to evaluate the overall stance of monetary policy, but they have long recognized that such information can be obscured by time-varying risk premiums.¹ Increasingly, central banks also use dynamic stochastic general equilibrium (DSGE) models with interest rate policy rules to think about the consequences of alternative policy actions in a rational expectations setting. These DSGE models can make predictions about the expectational and term premium components in bond yields; thus, it is natural to examine the extent to which these predictions are consistent with the observed data. Accordingly, we attempt to account for the observed size and volatility of the risk premium on long-term nominal bonds using a fairly wide variety of alternative DSGE model specifications that have been proposed in the literature.

Early work on bond pricing by Backus et al. (1989) examined the bond premium using a consumption-based asset pricing model of an endowment economy. They found that “the representative agent model with additively separable preferences fails to account for the sign or the magnitude of risk premiums” and “cannot account for the variability of risk

[☆] For helpful comments, we thank participants at the conference honoring John Taylor held at the Federal Reserve Bank of Dallas in 2007, especially our discussant Monika Piazzesi, and participants at the 2007 NBER Summer Institute session on DSGE modeling, the 2007 Dynare annual conference in Paris, the 2008 London School of Economics conference on macroeconomics, and the 2008 Western Finance Association Meetings, particularly our discussant Benjamin Croitoru, as well as many colleagues in the Federal Reserve System. Anjali Upadhyay provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the views of other individuals within the Federal Reserve System.

^{*} Corresponding author. Tel.: +1415 974 3172.

E-mail address: eric.swanson@sf.frb.org (E.T. Swanson).

¹ Notably, in 2004 and 2005, as long- and short-term interest rates diverged—the so-called bond yield “conundrum”—measures of the size of the risk premium on long-term nominal bonds attracted widespread attention (see Rudebusch et al., 2006; and Smith and Taylor, 2007).

premiums” (p. 397). This basic inability of a standard theoretical finance model to generate a sufficiently large and variable nominal bond risk premium has been termed the “bond premium puzzle.” Subsequently, Donaldson et al. (1990) and Den Haan (1995) showed that the bond premium puzzle is likewise present in standard real business cycle models with variable labor and capital and with or without simple nominal rigidities. Since these early studies, however, the “standard” theoretical model in macroeconomics has undergone dramatic changes and now includes a prominent role for habits in consumption and nominal rigidities that persist for several periods (such as staggered Taylor, 1980, or Calvo, 1983, price contracts), both of which should help the model to account for the term premium. Christiano et al. (2005) and Smets and Wouters (2003) have shown that DSGE models with these features can match the impulse responses of the economy to nominal shocks and technology shocks better than the earlier generation of models. We investigate whether these models are likewise better able to match the price and risk premium on a long-term nominal bond. To preview our results, we find that the bond premium puzzle remains in state-of-the-art macroeconomic DSGE models, even when these models are extended to include large and persistent habits as in Campbell and Cochrane (1999) and Wachter (2006) and real wage bargaining rigidities as in Blanchard and Galí (2005). That is, these models are still very far from matching the level and variability of the term premium, the slope of the yield curve, and the excess returns to holding long-term bonds that we see in the data.

The importance of jointly modeling both macroeconomic variables and asset prices within a DSGE framework is sometimes underappreciated. Indeed, a standard research tack has been to use DSGE models to explain the behavior of macroeconomic variables and latent-factor finance models to fit asset prices, but this dichotomous modeling approach suffers from at least two serious shortcomings.² First, as a theoretical matter, asset prices and the macroeconomy are inextricably linked, so a failure of the standard DSGE framework to explain asset prices suggests flaws in the model. As emphasized by Cochrane (2007), asset markets are the mechanism by which consumption and investment are allocated across time and states of nature, so asset prices, which equate marginal rates of substitution and transformation, are at the very foundation of the dynamics of macroeconomic quantities. If a DSGE model can match the data on macroeconomic quantities but not asset prices, then how does the model propose that marginal rates of substitution and transformation are being equated? Surely, such behavior is a sign that the model itself is flawed or at least incomplete. Second, from a practical point of view, policymakers and others are often very interested in the interaction between macroeconomic variables and asset prices—both the effects of asset prices on macro variables and the effects of interest rates and other macro variables on asset prices. For example, a question of recent interest is how does a seemingly very low term premium—the bond yield “conundrum”—affect the economy. As Rudebusch et al. (2007) discuss, this question cannot be addressed with a dichotomous macroeconomic and financial modeling approach; it requires a structural macro-finance model.

Although the bond premium puzzle has received far less attention in the literature than Mehra and Prescott's (1985) equity premium puzzle, it is in fact just as interesting and important. Indeed, as a practical matter, the value of long-term bonds outstanding in the U.S. is far larger than the value of equities. In addition, from a modeling perspective, the bond premium puzzle provides a very different metric for model performance. For example, Boldrin et al. (2001) can account for the equity premium puzzle in a two-sector DSGE model because capital immobility across the two sectors greatly increases the variance of the price of capital (and thus stock prices) and its covariance with consumption. However, this mechanism cannot explain the bond premium puzzle, which involves the valuation of a constant nominal coupon on a default-free government bond. In contrast to the equity premium puzzle, the bond premium puzzle is intimately related to the behavior of inflation, nominal rigidities, and nominal asset prices, which are crucial and still unresolved aspects of the current generation of DSGE models.

The bond premium puzzle has also attracted renewed interest in the finance and macro literatures. Wachter (2006) and Piazzesi and Schneider (2006) have had notable success in resolving this puzzle within an endowment economy by using preferences that have been modified to include either an important role for habit, as in Campbell and Cochrane (1999), or “recursive utility,” as in Epstein and Zin (1989). While such success in an endowment economy is encouraging, it is somewhat unsatisfying because, as noted above, the lack of structural relationships between the macroeconomic variables precludes studying many questions of interest. Accordingly, there has been interest in extending the endowment economy results to more fully specified DSGE models.³ Wu (2006), Bekaert et al. (2005), Hördahl et al. (2007), and Doh (2006) use the stochastic discount factor from a standard DSGE model to study the term premium, but to solve the model, these authors have essentially assumed that the term premium is constant over time—that is, they have essentially assumed the expectations hypothesis.⁴ Since we are interested in the variability as well as the level of the term premium, and in the

² These shortcomings have fostered a greater emphasis in the literature on the importance of macro-finance linkages, for example, Ang and Piazzesi (2003) and Diebold et al. (2005).

³ Hördahl et al. (2006), Rudebusch and Wu (2007), and other macro-finance researchers have examined term premiums with an affine no-arbitrage structure and a log-linearized version of a DSGE model. However, these models employ an exogenous stochastic pricing kernel that does not enforce a consistency between the asset pricing structure and the utility function underlying the macro-structure.

⁴ Wu (2006) and Bekaert et al. (2005) use a log-linear, log-normal approximation to solve the model, which allows some second- and higher-order terms from the log-normal distribution to remain in these models, although the implied term premium is constant. An additional drawback of their approach is that it treats some second-order terms as important while dropping other terms of similar magnitude. In contrast, Hördahl et al. (2006), compute a full second-order approximate solution to the model, which treats all second-order terms equally; however, the term premium is also a constant in this approach, as we discuss below. Doh (2006) does allow for a time-varying term premium, but does so by combining a full second-order

relationship between the term premium and the macroeconomy, a higher-order approximate solution method or a global nonlinear method is required, as in Ravenna and Seppälä (2006), Rudebusch et al. (2007), and Gallmeyer et al. (2005).⁵ Still, as we discuss in detail below, these last authors have had mixed success in solving the bond premium puzzle, and in particular, it remains unclear whether the size and volatility of the bond premium can be replicated in a DSGE model without distorting its macroeconomic fit and stochastic moments. Our analysis sheds light on this issue.

Our paper also has some similarities with the equity premium studies of Jermann (1998) and Lettau and Uhlig (2000). Just as those authors raise questions about the ability of Campbell and Cochrane's (1999) habit specification to match the equity premium in a real business cycle model, our results raise serious questions about the ability of standard macroeconomic habit specifications, and the Campbell and Cochrane (1999) and Wachter (2006) extensions of that specification, to match the nominal asset pricing facts. Our results build on the earlier work by considering nominal bond prices in a modern DSGE model with a central role for nominal rigidities, labor market frictions, and habits in consumption, such as Christiano et al. (2005). We show that even such state-of-the-art macro models fall egregiously short of being able to price nominal assets. Our results suggest that non-habit-based modifications of the model, such as Epstein and Zin (1989) and Weil (1989) recursive preferences, may be more promising extensions of DSGE models to asset pricing.

In the next section, we introduce our benchmark DSGE model, which is set well within the broad range of the literature, and show how to derive the term premium and other measures of long-term bond risk in the model. Section 3 compares the implications of the model to the data and shows that the term premium in the model is counterfactually small and stable. Section 4 explores whether long-memory habits, as in Campbell and Cochrane (1999), can help the model to explain the term premium, and Section 5 considers whether adding labor market frictions to the model might improve its fit. Section 6 concludes.

2. A benchmark DSGE model

We begin our investigation by outlining a standard benchmark DSGE model with nominal rigidities. We then define the term premium on a long-term bond and several other common measures that can be used to assess the bond market performance of a model. Finally, we describe our solution methods.

2.1. The benchmark model

The economy contains a continuum of households with a total mass of unity. Households are representative and seek to maximize utility over consumption and labor streams given by

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{(c_t - bh_t)^{1-\gamma}}{1-\gamma} - \chi_0 \frac{l_t^{1+\chi}}{1+\chi} \right), \tag{1}$$

where β denotes the household's discount factor, c_t denotes consumption in period t , l_t denotes labor, h_t denotes a predetermined stock of consumption habits, and γ , χ , χ_0 , and b are parameters. In our baseline specification, we will set $h_t = C_{t-1}$, the level of aggregate consumption in the previous period (so the habit stock is external to the household), although we will consider alternative formulations such as long-memory habits and internal habits later. The household's nominal stochastic discount factor from period t to $t+j$ in this model thus satisfies

$$m_{t,t+j} \equiv \beta^j \frac{(c_{t+j} - bC_{t+j-1})^{-\gamma} P_t}{(c_t - bC_{t-1})^{-\gamma} P_{t+j}}. \tag{2}$$

The economy also contains a continuum of monopolistically competitive intermediate goods firms indexed by $f \in [0, 1]$ that set prices according to Calvo contracts and hire labor competitively from households. Firms have Cobb–Douglas production functions:

$$y_t(f) = A_t \bar{k}^{(1-\alpha)} l_t(f)^\alpha, \tag{3}$$

where \bar{k} is a fixed, firm-specific capital stock (identical across firms) and where A_t denotes an aggregate technology shock that affects all firms.⁶ The level of aggregate technology follows an exogenous AR(1) process:

$$\log A_t = \rho_A \log A_{t-1} + \varepsilon_t^A, \tag{4}$$

where ε_t^A denotes an *i.i.d.* aggregate technology shock with mean zero and variance σ_A^2 .

(footnote continued)

solution with an ARCH process on one of the shocks, which again treats some third- and higher-order terms as being important while dropping other terms of similar magnitude.

⁵ Ravenna and Seppälä (2006) and Rudebusch et al. (2007) use a third-order approximate solution to the model, which allows for a time-varying term premium. Gallmeyer et al. (2005) are able to compute a closed-form solution for bond prices for a very special monetary policy reaction function, but their method does not apply to more general monetary policy reaction functions such as the Taylor policy rule we consider below.

⁶ Several authors, such as Woodford (2003) and Altig et al. (2004), have emphasized the importance of firm-specific fixed factors for generating a level of inflation persistence that is consistent with the data. With firm-specific capital stocks, the term premium is higher as well as inflation being more persistent.

Intermediate goods are purchased by a perfectly competitive final goods sector that produces the final good with a CES production technology:

$$Y_t = \left[\int_0^1 y_t(f)^{1/(1+\theta)} df \right]^{1+\theta}. \quad (5)$$

Each intermediate goods firm f thus faces a downward-sloping Dixit–Stiglitz demand curve for its product and the aggregate price level P_t is defined to be the Dixit–Stiglitz price aggregate.

Each firm sets its price $p_t(f)$ according to a Calvo contract that expires with probability $1 - \xi$ each period, with no indexation. Firms hire labor $l_t(f)$ from households in a competitive labor market, paying the nominal market wage w_t . Firms are collectively owned by households and distribute profits and losses back to the households. When a firm's price contract expires and it is able to set a new contract price, the firm maximizes the expected present discounted value of profits over the lifetime of the contract, using the representative household's stochastic discount factor (2) to value future profits. Firms' optimality conditions and the aggregate resource constraints in this model are standard and are described in the appendix to Rudebusch et al. (2007).

Although agents cannot invest in physical capital in the baseline version of the model, we do assume that an amount $\delta\bar{K}$ of output each period is devoted to maintaining the fixed capital stock. Households can also buy and sell one-period risk-free nominal bonds, subject to an individual borrowing constraint that is not binding but rules out Ponzi schemes. Optimizing behavior by households gives rise to the intratemporal condition

$$\frac{w_t}{P_t} = \frac{\chi_0 l_t^\chi}{(c_t - bC_{t-1})^{-\gamma}}, \quad (6)$$

and the intertemporal Euler equation

$$(c_t - bC_{t-1})^{-\gamma} = \beta e^{i_t} E_t(c_{t+1} - bC_t)^{-\gamma} P_t/P_{t+1}, \quad (7)$$

where i_t denotes the continuously compounded interest rate on the one-period risk-free nominal bond.

The government levies lump-sum taxes G_t on households and destroys the resources it collects. The aggregate resource constraint implies that

$$Y_t = C_t + \delta\bar{K} + G_t, \quad (8)$$

where $C_t = c_t$, the consumption of the representative household. Government consumption follows an exogenous AR(1) process:

$$\log G_t = \rho_G \log G_{t-1} + \varepsilon_t^G, \quad (9)$$

where ε_t^G denotes an *i.i.d.* government consumption shock with mean zero and variance σ_G^2 .

Finally, there is a monetary authority in the economy which sets the one-period nominal interest rate i_t according to a Taylor-type policy rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)[1/\beta + \bar{\pi}_t + g_y(Y_t - \bar{Y})/\bar{Y} + g_\pi(\bar{\pi}_t - \pi^*)] + \varepsilon_t^i, \quad (10)$$

where $1/\beta$ is the steady-state real interest rate in the model, \bar{Y} denotes the steady-state level of output, π^* denotes the steady-state rate of inflation, ε_t^i denotes an *i.i.d.* stochastic monetary policy shock with mean zero and variance σ_i^2 , and ρ_i , g_y , and g_π are parameters.⁷ The variable $\bar{\pi}_t$ denotes a geometric moving average of inflation:

$$\bar{\pi}_t = \theta_\pi \bar{\pi}_{t-1} + (1 - \theta_\pi)\pi_t, \quad (11)$$

where current-period inflation $\pi_t \equiv \log(P_t/P_{t-1})$ and we set $\theta_\pi = 0.7$ so that the geometric average in (11) has an effective duration of about four quarters, which is typical in estimates of the Taylor rule. The advantage of using (11) rather than the four-quarter average inflation rate is that (11) only requires keeping track of one lagged variable ($\bar{\pi}_{t-1}$) and hence one extra state variable in the model, while a four-quarter moving average would require keeping track of three (π_{t-1} , π_{t-2} , and π_{t-3}). All of our results below are very similar whether we use (11) or a more traditional four-quarter average inflation rate in the policy rule (10).

2.2. The term premium in the model

The price of any asset in the model economy must satisfy the standard stochastic discounting relationship in which the household's stochastic discount factor is used to value the state-contingent payoffs of the asset in period $t + 1$. For example,

⁷ In Eq. (10) (and Eq. (10) only), we express i_t , π_t , and $1/\beta$ in annualized terms, so that the coefficients g_π and g_y correspond directly to the estimates in the empirical literature. We also follow the literature by assuming an "inertial" policy rule with *i.i.d.* policy shocks, although there are a variety of reasons to be dissatisfied with the assumption of AR(1) processes for all stochastic disturbances except the one associated with short-term interest rates. Indeed, Rudebusch (2002, 2006) and Carrillo et al. (2007) provide strong evidence that an alternative policy specification with serially correlated shocks and little gradual adjustment is more consistent with the dynamic behavior of nominal interest rates.

the price of a default-free n -period zero-coupon bond that pays one dollar at maturity satisfies

$$p_t^{(n)} = E_t[m_{t+1}p_{t+1}^{(n-1)}], \tag{12}$$

where $m_{t+1} \equiv m_{t,t+1}$, $p_t^{(n)}$ denotes the price of the bond at time t , and $p_t^{(0)} \equiv 1$, i.e., the time- t price of one dollar delivered at time t is one dollar. The continuously compounded yield to maturity on the n -period zero-coupon bond is defined to be

$$i_t^{(n)} \equiv -\frac{1}{n} \log p_t^{(n)}. \tag{13}$$

In the U.S. data, the benchmark long-term bond is the 10-year Treasury note. Thus, we wish to model the term premium on a bond with a duration of about 10 years. Computationally, it is inconvenient to work with a zero-coupon bond that has more than a few periods to maturity; instead, it is much easier to work with an infinitely lived consol-style bond that has a time-invariant or time-symmetric structure. Thus, we assume that households in the model can buy and sell a long-term default-free nominal consol which pays a geometrically declining coupon in every period in perpetuity. The nominal consol's price per one dollar of coupon in period t , which we denote by $\tilde{p}_t^{(n)}$, then satisfies

$$\tilde{p}_t^{(n)} = 1 + \delta_c E_t m_{t+1} \tilde{p}_{t+1}^{(n)}, \tag{14}$$

where δ_c is the rate of decay of the coupon on the consol. By choosing an appropriate value for δ_c , we can thus model prices of a bond of any desired Macaulay duration or maturity n , such as the 10-year maturity that serves as our zero-coupon benchmark in the data.⁸ Finally, the continuously compounded yield to maturity on the consol, $\tilde{\gamma}_t^{(n)}$, is given by

$$\tilde{\gamma}_t^{(n)} \equiv \log \left(\frac{\delta_c \tilde{p}_t^{(n)}}{\tilde{p}_t^{(n)} - 1} \right). \tag{15}$$

Note that even though the nominal bond in our model is default-free, it is still risky in the sense that its price can covary with the household's marginal utility of consumption. For example, when inflation is expected to be higher in the future, then the price of the bond generally falls, because households discount its future nominal coupons more heavily. If times of high inflation are correlated with times of low output (as is the case for technology shocks in the model), then households regard the nominal bond as being very risky, because it loses value at exactly those times when the household values consumption the most. Alternatively, if inflation is not very correlated with output and consumption, then the bond is correspondingly less risky. In the former case, we would expect the bond to carry a substantial risk premium (its price would be lower than the risk-neutral price), while in the latter case we would expect the risk premium to be smaller.

In the literature, the risk premium or *term premium* on a long-term bond is typically expressed as the difference between the yield on the bond and the unobserved risk-neutral yield for that same bond. To define the term premium in our model, then, we first define the risk-neutral price of the consol, $\hat{p}_t^{(n)}$:

$$\hat{p}_t^{(n)} \equiv E_t \sum_{j=0}^{\infty} e^{-i_{t,t+j}} \delta_c^j, \tag{16}$$

where $i_{t,t+j} \equiv \sum_{n=0}^j i_{t+n}$. Eq. (16) is the expected present discounted value of the coupons of the consol, where the discounting is performed using the risk-free rate rather than the household's stochastic discount factor.⁹ Equivalently, Eq. (16) can be expressed in first-order recursive form as

$$\hat{p}_t^{(n)} = 1 + \delta_c e^{-i_t} E_t \hat{p}_{t+1}^{(n)}, \tag{17}$$

which directly parallels Eq. (14). The implied term premium on the consol is then given by

$$\psi_t^{(n)} \equiv \log \left(\frac{\delta_c \tilde{p}_t^{(n)}}{\tilde{p}_t^{(n)} - 1} \right) - \log \left(\frac{\delta_c \hat{p}_t^{(n)}}{\hat{p}_t^{(n)} - 1} \right), \tag{18}$$

which is the difference between the observed yield to maturity on the consol and the risk-neutral yield to maturity.

For a given set of structural parameters of the model, we will choose δ_c so that the bond has a Macaulay duration of $n = 40$ quarters, and we will multiply Eq. (18) by 400 in order to report the term premium in units of annualized percentage points rather than logs.

⁸ As δ_c approaches 0, the consol behaves more like cash—a zero-period zero-coupon bond. As δ_c approaches 1, the consol approaches a traditional consol with a fixed (nondepreciating) nominal coupon, which, under our baseline parameter values below, has a duration of about 25 years. By setting $\delta_c > 1$, the duration of the consol can be made even longer.

⁹ In computing the term premium, some authors take the expectation over *yields* rather than over prices (with the difference between the two approaches being a convexity term). Eq. (16) follows the no-arbitrage finance and macro-finance literatures (e.g., Ang and Piazzesi, 2003), which compute risk-neutral bond prices by setting the prices of risk to zero. An alternative definition of the risk-neutral bond price, suggested to us by Oreste Tristani, would consider what value a single hypothetical agent with utility function (1) and $\gamma = 0$ would assign to the bond. This definition is problematic, however, because the risk-neutral agent has an intertemporal elasticity of substitution that differs from that of the representative agent in the economy, which implies that the risk-neutral agent and representative agents have different one-period risk-free rates after a shock. Thus, even in a riskless world, this alternative definition would imply a time-varying “term premium”. Our definition appears more consistent with the finance and macro-finance literatures.

2.3. Alternative measures of long-term bond risk

Although the term premium is the cleanest conceptual measure of the riskiness of long-term bonds, it is not directly observed in the data and must be inferred using term structure models or other methods. Accordingly, the literature has also focused on three other empirical measures that are closely related to the term premium but are more easily observed: the slope of the yield curve, the excess return to holding the long-term bond for one period relative to the one-period short rate, and the slope coefficient from a Campbell–Shiller (1991) predictability regression of the change in the long-term yield on the yield curve slope.¹⁰ In the next section, we will compare the model's ability to fit the data using each of these measures, and here we define them in more detail.

The slope of the yield curve is simply the difference between the yield to maturity on the long-term bond and the one-period risk-free rate, i_t . The slope of the yield curve is an imperfect measure of the riskiness of the long-term bond because the yield curve slope can vary in response to shocks even if all investors in the model are risk-neutral. However, on average, the slope of the yield curve equals the term premium, and the volatility of the yield curve slope provides us with a noisy measure of the volatility of the term premium.

A second measure of the riskiness of long-term bonds is the excess one-period holding return—that is, the return to holding the bond for one period less the one-period risk-free rate. For the case of an n -period zero-coupon bond, this excess return is given by

$$x_t^{(n)} \equiv \frac{p_t^{(n-1)}}{p_{t-1}^{(n)}} - e^{i_{t-1}}. \tag{19}$$

The first term on the right-hand side of (19) is the gross return to holding the bond and the second term is the gross one-period risk-free rate. For the case of the consol in our model, the excess holding period return is a bit more complicated, since the consol pays a coupon in period $t - 1$ and then depreciates in value by the factor δ_c , so the excess holding period return is given by

$$\tilde{x}_t^{(n)} \equiv \frac{\delta_c \tilde{p}_t^{(n)} + e^{i_{t-1}}}{\tilde{p}_{t-1}^{(n)}} - e^{i_{t-1}}. \tag{20}$$

Again, the first term on the right-hand side of (20) is the gross return to holding the consol and includes the one-dollar coupon in period $t - 1$ that can be invested in the one-period security. As with the yield curve slope, the excess returns in (19) and (20) are imperfect measures of the term premium because they would vary in response to shocks even if investors were risk-neutral. However, the mean and standard deviation of the excess holding period return provide popular measures of the average term premium and the volatility of the term premium.

A third measure of bond risk is based on the “long-rate regression” popularized by Campbell and Shiller (1991):

$$i_{t+1}^{(n-1)} - i_t^{(n)} = \alpha_{CS}^{(n)} + \beta_{CS}^{(n)} \frac{i_t^{(n)} - i_t}{n-1} + \varepsilon_{t+1}^{(n)}, \tag{21}$$

where the dependent variable is the change in the n -period zero-coupon yield from period t to $t + 1$, the independent variable is the yield curve slope at time t (divided by $n - 1$), and $\alpha_{CS}^{(n)}$ and $\beta_{CS}^{(n)}$ are maturity-specific intercept and slope coefficients. Under the expectations hypothesis, expected excess returns (19) are zero, that is

$$\log p_{t+1}^{(n-1)} - \log p_t^{(n)} = i_t. \tag{22}$$

After substituting definition (13) and rearranging terms, (22) would imply that the coefficients $\beta_{CS}^{(n)} = 1$ and $\alpha_{CS}^{(n)} = 0$; that is, the yield curve slope is the optimal forecast of the future change in the long rate. Deviations from risk neutrality drive $\alpha_{CS}^{(n)}$ away from zero, and time-variation in the term premium pushes $\beta_{CS}^{(n)}$ away from unity. Note that (21) could equivalently be written and run as

$$\frac{\log p_t^{(n)}}{n} - \frac{\log p_{t+1}^{(n-1)}}{n-1} = \alpha_{CS}^{(n)} + \beta_{CS}^{(n)} \frac{i_t^{(n)} - i_t}{n-1} + \varepsilon_{t+1}^{(n)}, \tag{23}$$

which turns out to be the more useful format for the consol below.

For the consol in our model, the derivation of the Campbell–Shiller regression coefficient is a bit more complicated, reflecting the extra terms in (20) relative to (19). Instead of (22), the appropriate equation is

$$\log(\delta_c \tilde{p}_{t+1}^{(n)} + e^{i_t}) - \log \tilde{p}_t^{(n)} = i_t. \tag{24}$$

After substituting definition (15) and rearranging terms, this implies

$$\log \tilde{p}_t^{(n)} - \log \tilde{p}_{t+1}^{(n)} = \tilde{i}_t^{(n)} - i_t. \tag{25}$$

¹⁰ See, respectively, Piazzesi and Schneider (2006), Cochrane and Piazzesi (2005), and Rudebusch and Wu (2007).

Thus, the Campbell–Shiller regression for the consol in the model is most simply written as

$$\log \tilde{p}_t^{(n)} - \log \tilde{p}_{t+1}^{(n)} = \alpha_{CS}^{(n)} + \beta_{CS}^{(n)} (\tilde{r}_t^{(n)} - i_t) + \varepsilon_{t+1}^{(n)}, \quad (26)$$

where the coefficients $\alpha_{CS}^{(n)}$ and $\beta_{CS}^{(n)}$ in the model have exactly the same interpretation as in the data.

2.4. Model solution method

A technical issue in solving the model above arises from the relatively large number (nine) of state variables— C_{t-1} , A_{t-1} , G_{t-1} , i_{t-1} , A_{t-1} , $\bar{\pi}_{t-1}$, and the three shocks, ε_t^A , ε_t^G , and ε_t^i .¹¹ Because of such high dimensionality, value-function iteration-based methods such as projection methods (or, even worse, discretization methods) are computationally intractable. We instead solve the model using the standard macroeconomic technique of approximation around the nonstochastic steady state—so-called perturbation methods. However, a first-order approximation of the model (i.e., a linearization or log-linearization) eliminates the term premium entirely, because Eqs. (14) and (17) are identical to first order, a manifestation of the well-known property of certainty equivalence in linearized models. A second-order approximation to the solution of the model produces a term premium that is nonzero but constant (a weighted sum of the variances σ_A^2 , σ_G^2 , and σ_i^2). Since our interest in this paper is not just in the level of the term premium but also in its volatility and variation over time, we must compute a third-order approximate solution to the model around the nonstochastic steady state. We do so using the n th-order perturbation AIM algorithm of Swanson et al. (2006), which automatically and quickly computes n th-order approximate solutions to dynamic discrete-time rational expectations models of this type. For the baseline model above with nine state variables, a third-order accurate solution can be computed in about 10 minutes on a standard laptop computer. Additional details of this solution method are provided in Swanson et al. (2006) and Rudebusch et al. (2007).

Once we have computed an approximate solution to the model, we compare the model and the data using a standard set of macroeconomic and financial moments, including the standard deviations of consumption, labor, and other variables, and the means and standard deviations of the term premium and the alternative measures of long-term bond risk described above. Because our approximate solution to the model is nonlinear, we compute these moments from synthetic model data. (Namely, beginning from the nonstochastic steady state, we simulate the model forward 500,000 observations using normally distributed shocks.)

3. Comparing the model to the data

How well can the benchmark DSGE model fit the first and second empirical moments of macroeconomic and financial variables? We begin with a baseline parameterization of the DSGE model drawn from the literature. Since this standard parameterization is unable to fit the empirical facts, we then explore alternatives that may help the model fit the data.

3.1. Baseline model parameterization and sensitivity analysis

The baseline set of parameter values with which we begin our analysis are reported in the first column of Table 1 and are typical of those in the literature (see, e.g., Levin et al., 2005). We set the household's discount factor to 0.99 per quarter (implying a steady-state real interest rate of 4.02 percent per year), firms' output elasticity with respect to labor to 0.7, firms' steady-state markup to 0.2 (implying a price-elasticity of demand of 6), and the average price contract duration to four quarters. The importance of habits in the household's utility is set to 0.66, consistent with typical estimates in the macro literature. We set the utility curvature parameter γ to 2, which is a little on the high side of standard macroeconomic estimates, to give the model a better chance of generating an appreciable term premium. We set the utility curvature parameter on labor χ to 1.5 (implying a Frisch elasticity of about 0.7), which is again a little higher than typical macro estimates (but in line with estimates from the labor literature) but also gives the model a better chance of matching the term premium. The shock persistences ρ_A and ρ_G are set to 0.9, as is common, and the shock variances σ_A^2 and σ_G^2 are set to 0.01^2 and 0.004^2 , respectively, consistent with typical estimates in the literature. The monetary policy rule coefficients are taken from Rudebusch (2002) and are also typical of those in the literature. We set the steady-state capital–output ratio to 2.5, which is close to what is found in the data, and steady-state government spending to about 17 percent of output. As is standard, we set the baseline steady-state inflation rate in the model to 0 percent per year. The parameter χ_0 is chosen to normalize the steady-state quantity of labor to unity and, as discussed above, the parameter δ_c is chosen to set the Macaulay duration of the consol in the model to 10 years.

We compute the third-order approximate solution to the model with these baseline parameter values, and then compute various model-implied moments by simulation. The results of this exercise are reported in the first two columns of Table 2, along with the corresponding empirical moments for quarterly U.S. data from 1960 to 2007. For the empirical

¹¹ The number of state variables can be reduced a bit by noting that G_t and A_t are sufficient to incorporate all of the information from G_{t-1} , A_{t-1} , ε_t^G , and ε_t^A , but the basic point remains valid, namely, that the number of state variables in the model is large from a computational point of view.

Table 1
Baseline parameter values and sensitivity analysis

Parameter	Baseline case value	Low case		High case	
		Value	Mean[$\psi^{(10)}$]	Value	Mean[$\psi^{(10)}$]
α	0.7	0.5	0.013	0.85	0.015
β	0.99	0.97	0.014	0.995	0.014
θ	0.2	0.05	0.008	0.4	0.017
ξ	0.75	0.5	0.026	0.9	0.005
γ	2	0.5	−0.015	6	0.045
χ	1.5	0	0.006	5	0.029
b	0.66	0	0.010	0.9	0.026
ρ_A	0.9	0.7	0.004	0.95	0.039
ρ_G	0.9	0.7	0.014	0.95	0.014
σ_A^2	0.01 ²	0.005 ²	0.006	0.02 ²	0.047
σ_G^2	0.004 ²	0.002 ²	0.014	0.008 ²	0.015
ρ_i	0.73	0	0.038	0.9	0.007
g_π	0.53	0.05	−0.035	1	0.033
g_y	0.93	0	0.035	2	−0.010
σ_i^2	0.004 ²	0.002 ²	0.012	0.008 ²	0.022
$\bar{K}/(4\bar{Y})$	2.5	1.25	0.010	5	0.025
\bar{G}/\bar{Y}	0.17	0.1	0.011	0.3	0.021
π^*	0	–	–	0.02	0.021
Memo:					
χ_0	4.74				
δ_c	0.9848				

Note: Term premium means are measured in percentage points.

Table 2
Empirical and model-based unconditional moments

Variable	U.S. data 1960–2007	Parameterizations of DSGE model			
		Baseline	HTV	RS	Best fit
sd[C]	1.19	1.36	12.54	5.14	0.77
sd[Y]	1.50	0.86	7.90	3.24	0.50
sd[L]	1.71	2.81	9.73	5.14	1.75
sd[w ^f]	0.82	2.27	12.57	10.67	1.50
sd[π]	2.52	2.35	15.29	7.67	2.57
sd[i]	2.71	2.06	15.05	7.02	2.79
sd[r]	2.30	1.97	5.67	5.10	2.19
sd[i ⁽¹⁰⁾]	2.41	0.55	10.16	2.70	0.98
Mean[$\psi^{(10)}$]	1.06	0.014	0.686	0.197	0.106
sd[$\psi^{(10)}$]	0.54	0.001	1.51	0.081	0.013
Mean[i ⁽¹⁰⁾ − i]	1.43	−0.053	0.095	0.011	−0.089
sd[i ⁽¹⁰⁾ − i]	1.33	1.55	5.37	4.55	1.90
Mean[x ⁽¹⁰⁾]	1.76	0.014	10.94	1.38	0.067
sd[x ⁽¹⁰⁾]	23.43	6.98	62.26	28.21	10.96
$\beta_{CS}^{(10)}$	−3.49	0.97	0.98	1.00	0.99

Except for the Campbell–Shiller regression coefficient, $\beta_{CS}^{(10)}$, all variables are quarterly values expressed in percent. Inflation and interest rates, the term premium (ψ), and excess holding period returns (x) are expressed at an annual rate.

moments, consumption is real personal consumption expenditures from the U.S. national income and product accounts, labor is total hours of production workers from the Bureau of Labor Statistics, and the real wage is total wages and salaries of production workers from the BLS divided by total production worker hours and deflated by the GDP price index. The standard deviation was computed for logarithmic deviations of each series from a Hodrick–Prescott trend and reported in percentage points. Standard deviations for inflation, interest rates, and the term premium were computed for the raw series rather than for deviations from trend. Inflation is the annualized rate of change in the quarterly GDP price index from the Bureau of Economic Analysis. The short-term nominal interest rate i is the end-of-month federal funds rate from the Federal Reserve Board, in annualized percentage points. The short-term real interest rate r is the short-term nominal interest rate less the realized quarterly inflation rate. The 10-year zero-coupon bond yield is the end-of-month 10-year zero-coupon bond yield taken from Gurkaynak et al. (2007). The term premium on the 10-year zero-coupon bond

is the term premium computed by Kim and Wright (2005), in annualized percentage points.¹² The yield curve slope, one-period excess holding return, and Campbell–Shiller coefficient are from the authors' calculations based on the data above; the yield curve slope and one-period excess holding return are reported in annualized percentage points.

As can be seen in Table 2, for the benchmark model with standard parameter values, the average term premium is a bit less than 1.4 *basis points*, and the standard deviation of the term premium is around 0.1 basis point, both roughly two orders of magnitude smaller than the data.¹³ The results are basically no better for the yield curve slope or excess holding period return measures—although the standard deviations of these variables are greater than zero, that variation is due entirely to the risk-neutral components of those variables rather than any variation in the riskiness of the long-term bond. By any measure, the long-term bond in the model is priced essentially risk-neutrally, resulting in term premia and other measures of bond risk that are negligible.

Although the benchmark DSGE model that we have used to conduct this experiment is fairly simple, we have obtained similar results from more complicated DSGE models in the literature. For example, in the moderately sized DSGE model of Christiano et al. (2005), the mean term premium is just 1 basis point—even smaller than in our benchmark model.¹⁴ In the Levin et al. (2005) version of the Smets and Wouters (2003) model, which has a greater number of shocks with high persistence and variance, the mean term premium is just 2.1 basis points.

From the point of view of a second- or third-order approximation to a macroeconomic model, these results should not be too surprising. The shocks in the benchmark model have a standard deviation of only about 1 percent, so first-order terms in the model have a magnitude that is roughly proportional to (0.01), second-order terms have a magnitude that is roughly proportional to (0.01)², where the constant of proportionality is related to the curvature of the model, and third-order terms have a magnitude that is roughly proportional to (0.01)³. Because the shocks in a macro-model like our benchmark model are typically so small, the second-order terms should be expected to be roughly 100 times smaller than the first-order terms for a relatively flat model, and the third-order terms should be expected to be roughly 10,000 times smaller. Only for an extremely curved model, or for much larger shock standard deviations, could we reasonably expect the second- or third-order terms to matter very much.

This basic intuition is supported by the sensitivity analysis we conduct in the remaining columns of Table 1. In each row of the table, we vary each parameter in turn over a wide range that broadly covers the empirical estimates of the parameter in the literature, to see if any of these variations causes the mean term premium to change substantially. (To conserve space, we do not report the alternative measures of bond risk, but their behavior is always similar to that of the mean term premium.) The middle columns of Table 1 report the mean term premium that results from using the “low” value for each parameter, and the rightmost columns report the mean term premium that results from using the “high” value (as each parameter is varied, the other parameters are held fixed at their baseline values). For example, setting $\alpha = 0.5$ instead of 0.7 reduces the mean term premium to 1.3 basis points, while setting $\alpha = 0.85$ increases it to 1.5 basis points. Across all of these parameter variations, the mean term premium is always at least an order of magnitude too small relative to the data. Still, in line with the intuition above, some parameters are more important than others. In particular, the mean term premium appears most sensitive to the variance and persistence of the technology shock (σ_A^2 and ρ_A) and to the curvature of the utility function (γ and χ). These results foreshadow the two main approaches to increasing risk premiums in DSGE models, which we will discuss below.

3.2. Increased shock volatility

As suggested by the preceding discussion and Table 1, a simple way to increase the term premium in our model is to increase the size and persistence of the shocks. Indeed, two recent papers, by Hördahl et al. (2007) and Ravenna and Seppälä (2006), do exactly that. In order to generate a term premium that is in line with the data, however, both papers require extremely large shocks. For example, Hordahl et al. (HTV) assume that the technology shock has a quarterly standard deviation of 2.37 percent and a persistence of 0.986, compared with our baseline values of $\sigma_A = 1$ percent and $\rho_A = 0.9$. Adopting the two HTV parameter values in our model (while holding all other parameters fixed at their baseline values) increases the mean term premium from 1.4 to 69 basis points, and increases the standard deviation of the term premium from 0.1 to 151 basis points, both of which are much closer to the empirical estimates in Table 2. However, as shown in the third column of Table 2, the increased shock volatility also increases the volatility of output and the other macroeconomic variables in the model. For example, the unconditional standard deviation of labor and real wages are around 10 percent, far in excess of the data, and the unconditional standard deviation of inflation and the one-period

¹² Kim and Wright (2005) use an arbitrage-free, three-latent-factor affine model of the term structure to compute the term premium. Alternative measures of the term premium using a wide variety of methods produce qualitatively similar results in terms of the overall magnitude and variability—see Rudebusch et al. (2007) for a detailed discussion and comparison of several methods.

¹³ The lack of variation in the baseline model's estimate of the term premium is also illustrated by its impulse response to economic shocks. For example, the term premium moves less than *five one-hundredths* of one basis point on impact in response to a 1 percent technology shock and decays thereafter. See Rudebusch et al. (2007) for further discussion.

¹⁴ In contrast to Christiano et al. (2005), we assume that the central bank follows a Taylor-type reaction function for the short-term nominal interest rate (Eq. (18)) rather than a money growth rule. This modification to the model is standard practice in the large-scale DSGE models being put into practice at central banks and the IMF, among others.

nominal interest rate is over 15 percent.¹⁵ That is, the HTV parameterization can solve the bond premium puzzle, but only by sacrificing the model's fit to the macroeconomic variables.

The results of Ravenna and Seppala (RS) are very similar, as can be seen in the fourth column of Table 2. Instead of an unusually large technology shock, RS introduce a very large taste shock d_t into their model, where d_t is an AR(1) marginal rate of substitution shock assumed to have an out-sized quarterly standard deviation of 8 percent and a serial correlation of 0.95. Introducing a taste shock d_t of this size into our model leads to an average term premium of 19 basis points and a standard deviation of the term premium of 8 bp, closer to the empirical estimates. However, as in HTV, even this partial solution to the bond premium puzzle produces counterfactually large volatilities for all of the macroeconomic variables in the model.

3.3. Best-fit parameterization of the model

If the standard macroeconomic parameterization produces a term premium that is too small, and larger shock volatilities produce a reasonable term premium but destroy the fit of macroeconomic variables, is there some other parameterization of the model which might generate a reasonable fit along both dimensions? To address this question, we search over the wide range of parameter values reported in Table 1 to find the set of values that provides the best joint fit to both the macroeconomic and financial moments.

The computational time required to solve the model for each set of parameter values is a few minutes, so it is generally not feasible to estimate the model using maximum likelihood or Bayesian estimation procedures. Instead, we perform a grid search over the six parameters in Table 1 that are among the most uncertain and appeared to be the most important for the term premium—namely, γ , b , χ , ζ , ρ_A , and σ_A —and report the set of parameter values that best fits the macroeconomic and financial moments in Table 2.¹⁶ We define the “best fit” to be the set of parameters that matches the equally weighted standard deviations of consumption, labor, the real wage, inflation, the short-term nominal interest rate, short-term real interest rate, long-term bond yield, and the mean term premium as closely as possible.¹⁷

The last column of Table 2 presents the moments from the resulting best-fitting parameter values (which are $\gamma = 6$, $b = 0.9$, $\chi = 3$, $\zeta = 0.65$, $\rho_A = 0.95$, and $\sigma_A = 0.005$). With these parameter values, the mean term premium is about 11 basis points and the unconditional standard deviation of the term premium is about 1.3 basis points, a much better fit than the baseline model though still too small relative to the data. To achieve this better fit, the estimation procedure picks the highest possible curvature of the utility function with respect to consumption, $\gamma = 6$, and $b = 0.9$, and the highest possible technology shock persistence, $\rho_A = 0.95$. With these extreme parameter values, holding the technology shock standard deviation fixed at its baseline value would result in macroeconomic moments that are too volatile relative to the data, so the estimation chooses the lowest possible standard deviation, $\sigma_A = 0.005$. The values $\chi = 3$ and $\zeta = 0.65$ are intermediate, reflecting the compromise between better financial fit and worse macroeconomic fit.

Nevertheless, even for the best-fitting set of parameter values, our benchmark DSGE model is unable to match simultaneously both the term premium and the most basic macroeconomic moments in the data, let alone additional macroeconomic and financial moments. Although the best-fitting parameterization of the model improves the model's fit to many of the macroeconomic moments, the two moments that the model most fails to match are the mean term premium and the variability of the long-term bond yield. To model and eventually understand the behavior of these two variables clearly requires a more dramatic modification to the standard DSGE model.

4. Habit formation and the term premium

Several modifications to standard models have been suggested as explanations to the equity premium puzzle, including long-memory habit formation in consumption (Campbell and Cochrane, 1999), time-inseparable “recursive utility” preferences (Epstein and Zin, 1989), and heterogeneous agents (Constantinides and Duffie, 1996; Alvarez and Jermann,

¹⁵ Hör Dahl, Tristani, and Vestin use a monetary policy rule that differs from our Eq. (10) in that their specification assumes $\rho_i = 1$ and has no coefficient on output growth or the output gap, in contrast to standard estimates in the literature. Using their monetary policy rule instead of ours reduces inflation and interest rate volatility down to more reasonable levels, but increases the volatility of consumption, output, labor, and the real wage to levels that are even higher than we report in Table 2. Note that HTV report the standard deviation of consumption growth implied by their model but do not report the business-cycle variability of the output gap or labor or other variables such as the real wage. Our Table 2 makes it clear that these variables are indeed extremely volatile, far more so than in the data.

¹⁶ We conducted the grid search in two stages, first searching over a coarse grid: $\gamma \in \{0.5, 1, 1.5, 2, 2.5, 3, 4, 5, 6\}$, $b \in \{0, 0.2, 0.4, 0.5, 0.6, 0.66, 0.7, 0.8, 0.9\}$, $\chi \in \{0.1, 0.5, 1, 1.5, 2, 3, 4, 5\}$, $\zeta \in \{0.5, 0.6, 0.75, 0.9\}$, $\rho_A \in \{0.7, 0.8, 0.9, 0.95\}$, and $\sigma_A \in \{0.005, 0.0075, 0.01, 0.015, 0.02\}$. After finding a best fit at $\gamma = 6$, $b = 0.9$, $\chi = 3$, $\zeta = 0.6$, $\rho_A = 0.95$, and $\sigma_A = 0.005$, we then refined the grid near this parameter vector and searched over the finer grid: $\gamma \in \{5, 5.25, 5.75, 6\}$, $b \in \{0.75, 0.8, 0.85, 0.9\}$, $\chi \in \{2, 2.25, 2.5, 2.75, 3, 3.25, 3.5, 3.75, 4\}$, $\zeta \in \{0.5, 0.55, 0.6, 0.65, 0.7\}$, $\rho_A \in \{0.9, 0.95\}$, and $\sigma_A \in \{0.005, 0.0075, 0.01, 0.015, 0.02\}$.

¹⁷ Minimizing the equal-weighted distance to these six moments provides us with a consistent estimator of our parameters, though it is not efficient. We do not try to match both output and consumption because our benchmark DSGE model has a fixed capital stock and thus has nothing to say about investment, which is the primary difference between consumption and output. We do not try to match the standard deviation of the term premium because doing so requires a third-order approximation rather than a second-order approximation for every iteration of the model, which increases the computation required for each iteration. We do not try to match the other bond risk measures because they are very highly correlated with the term premium in the model, so they do not provide additional information for the model to match.

2001). Since these modifications have been relatively successful at producing substantial risk premiums in endowment model economies, it is natural to ask whether they might help a DSGE model match the level and volatility of the term premium. We focus on the long-memory habit specification of Campbell and Cochrane (1999) because the standard DSGE models in macroeconomics already include a prominent role for habit in consumption. Moving from the standard habit specification to the Campbell–Cochrane long-memory specification habit is a small variation that might allow the standard macroeconomic framework to fit the term premium without significantly degrading its fit to the macroeconomic data.

Campbell and Cochrane (1999) propose replacing the standard habit preferences (1) with a habit stock h_t that has a much longer memory over past consumption, and a parameter b that is much closer to unity, which increases the importance of habits in agents' utility. Moreover, to prevent current consumption from ever falling below habits (which would cross a singularity of the utility function (1)), Campbell and Cochrane define habits implicitly through surplus consumption s_t , as follows. First, define the household's surplus consumption ratio:

$$s_t \equiv \frac{c_t - bh_t}{c_t}. \tag{27}$$

The habit stock h_t is assumed to be external to the household ("keeping up with Joneses" habits), so letting capital letters denote aggregate quantities as above, household habits h_t are defined to equal the aggregate habit stock,

$$h_t \equiv \frac{C_t(1 - S_t)}{b}, \tag{28}$$

which is in turn defined implicitly by a process on S_t ,

$$\log S_t = (1 - \phi) \log \bar{S} + \phi \log S_{t-1} + \left(\frac{\sqrt{1 - 2 \log(S_t/\bar{S})}}{\bar{S}} - 1 \right) [\log(C_t/C_{t-1}) - E_{t-1} \log(C_t/C_{t-1})], \tag{29}$$

where ϕ and \bar{S} are parameters. The primary advantage of this complicated definition of habits is that it ensures household surplus consumption is always positive, which is important when the habit stock is a large fraction of current consumption. Campbell and Cochrane (1999) discuss the parameterization of (29) in detail, but surplus consumption and the habit stock must be persistent (ϕ close to 1) to match the persistence of risk premiums and \bar{S} must be very low (the habit stock must be very large relative to consumption) to match the level of the risk premium and keep the risk-free rate stable.

We investigate whether these long-memory habits can potentially explain the term premium by replacing the definition of h_t in our benchmark model with the definition of habits given by Eqs. (27)–(29). In all other respects, we keep the benchmark model the same.¹⁸ From the point of view of a Taylor series approximation, it is clear how these Campbell–Cochrane preferences could help make second- or even third-order terms more important. By increasing the size of habits relative to consumption (making \bar{S} small), this specification greatly increases the curvature of the household's utility function with respect to consumption—from a value of γ in the model with no habits or $\gamma/(1 - b)$ in our baseline DSGE model, to γ/\bar{S} in the model with long-memory habits. When \bar{S} is small, such as the value of 0.0588 calibrated by Campbell and Cochrane (1999), then the curvature of the utility function is magnified by a factor of more than 16 as compared to a model with no habits and by a factor of more than five relative to our baseline model. Such a large increase in the curvature of the model should be expected to increase the importance of higher-order terms in the Taylor series expansion.

Perhaps surprisingly, even with Campbell–Cochrane habits, our benchmark DSGE model is still unable to match the level and volatility of the term premium. The mean term premium implied by this model rises to 2.7 basis points, which is still far less than the 106-basis-point mean term premium estimated in the data. Moreover, this model does essentially no better at matching the term premium's volatility, as the unconditional standard deviation of the term premium remains less than 1 basis point. As was the case for our benchmark model in Table 1, this result is very robust—for example, it does not change if we vary each of the model's parameters over a wide range broadly covering the estimates in the literature, including variations in the habit importance \bar{S} as considered by Wachter (2006).¹⁹

This result stands in sharp contrast to Wachter (2006), who finds that Campbell–Cochrane habits can match the mean term premium in an endowment economy, where the exogenous process for consumption and inflation is estimated from the data. What can explain this dramatic difference in conclusions? The key, as also emphasized by Boldrin et al. (2001), Lettau and Uhlig (2000), and Jermann (1998), is that in a production-based model, households can endogenously choose their labor–consumption tradeoff; in contrast, in an endowment-based economy, households must consume whatever the endowment turns out to be.²⁰ If households are hit by a negative shock in a production-based model, they can compensate

¹⁸ As in the baseline model, we set χ_0 to normalize the quantity of labor $L = 1$ in steady state. However, because the marginal utility of consumption is so much higher with Campbell–Cochrane habits, the marginal disutility of labor must also be higher to arrive at the same steady-state quantity of labor, which produces $\chi_0 = 158.5$, much larger than in the baseline version of the model.

¹⁹ The results of the sensitivity analysis for each parameter are not reported in the interest of space, but are available in the working version of this paper (Rudebusch and Swanson, 2007). Wachter (2006), in contrast to Campbell and Cochrane (1999), allows the parameter \bar{S} to vary independently from the other parameters of the model, and we consider varying \bar{S} in this way as well. As with the other parameters in our model, variations in \bar{S} of the magnitude considered by Wachter have only a tiny, negligible effect on the term premium in our DSGE model.

²⁰ In Jermann (1998), households are unable to vary their labor supply but can vary investment instead, so the basic point is the same.

for the shock by increasing their labor supply and working more hours. As a result, they have the ability to insure themselves to some extent from the effects of the shock on consumption by endogenously varying their labor supply in response. Households in an endowment economy do not have this opportunity, so the consumption cost of shocks in an endowment economy is correspondingly greater and risky assets thus carry a larger risk premium. In the Campbell–Cochrane version of our benchmark model, this ability of households to self-insure is enough to almost completely offset the large effects that those habit preferences would otherwise have on the term premium.

This observation suggests that if labor in the model is not perfectly flexible or not completely within the household's control, then the ability of households to self-insure against shocks will be substantially diminished, and risk premiums increase toward the higher levels in the endowment economy case. Boldrin et al. (2001) and Uhlig (2007) have also emphasized the importance of labor market frictions for matching the equity premium in a production economy. We explore this case in the next section.

5. Labor market frictions and the term premium

Habits that are both very large and very persistent, like those in Campbell and Cochrane (1999), are unable to solve the bond premium puzzle in a DSGE model in which labor supply and production are endogenous. However, limiting the ability of households to vary their labor supply through labor market frictions should improve the model's ability to fit the term premium. There are many ways to introduce labor market rigidities into our benchmark DSGE model, and we consider three such frictions below. We begin with the simplest form of labor market friction, a quadratic adjustment cost, then we consider two frictions with more institutional realism: real wage rigidities and staggered nominal wage contracting.

5.1. Quadratic labor adjustment costs

In this subsection, we consider the effects on the term premium of a quadratic adjustment cost on changes in the quantity of labor from one period to the next. Specifically, in each period, households must pay an adjustment cost, $\kappa(\log(l_t/l_{t-1}))^2$, which is proportional to the squared log percentage change in labor from the previous quarter. Although this labor market friction is simplistic, it is particularly useful for gaining intuition because its size is so clearly parameterized by κ ; as κ increases, it becomes more expensive for households to insure themselves against a shock by varying their labor supply, so the magnitude of the term premium should increase. Convex adjustment costs to labor is also very similar to incorporating labor or leisure habits into the household's utility function, as in Uhlig (2007): under both specifications, changing labor is costly; with convex adjustment costs, these costs are taken out of income, while with labor habits these costs are taken directly out of utility, but there is a close correspondence between the two via the marginal utility of income.²¹

Fig. 1 displays the relationship between the mean term premium and κ for two versions of our model with quadratic adjustment costs: one with Campbell–Cochrane habits and one without. The horizontal axis measures κ in units of output cost for a 1 percent change in labor—that is, if $\kappa = 50\bar{Y}$, then a 1 percent change in labor from the previous quarter costs households 0.5 percent of quarterly steady-state output today. The unconditional standard deviation of labor in these models is about 2.5 percent, so a change in labor from one quarter to the next of about 1 percent is about the right order of magnitude for this representative-agent model. Fig. 1 shows that *both* a moderate level of adjustment costs and long-memory habits are necessary to generate a mean term premium that is roughly consistent with the data. Without such habits, even extreme levels of adjustment costs to labor do not have much of an effect on the term premium because variation in consumption is simply not that abhorrent to households. With Campbell–Cochrane habits, consumption variation is much more undesirable, so adjustment costs to labor quickly begin to generate a substantial aversion by households to risky assets. Indeed, with Campbell–Cochrane habits and adjustment costs of around 0.5 percent of output for a 1 percent change in labor ($\kappa = 50\bar{Y}$), the mean term premium in this model is 65 basis points, which is within range of the empirical estimate in the data.

Unfortunately, adding labor adjustment costs to the model comes at a cost in terms of fitting macroeconomic quantities. In the second column of Table 3, we report the unconditional standard deviations for other macroeconomic and financial variables for the model with Campbell–Cochrane habits and quadratic labor adjustment costs of $\kappa = 50\bar{Y}$. Although the volatility of the term premium is larger than under the baseline specification (in Table 2), so are the unconditional standard deviations of real wages, inflation, and short-term nominal interest rates. The volatility of the real wage in particular is over 220 log percentage points.²² Intuitively, the presence of labor adjustment costs along with Campbell–Cochrane habits

²¹ Jaccard (2007) considers an alternative habit formulation in which household utility is given by $(C_t v(L_t) - h_t)^{1-\gamma}/(1-\gamma)$ and habits, h_t , are an average of the lagged consumption–leisure composite $C_t v(L_t)$, so although consumption and labor can vary separately, households prefer a smooth composite. We embedded this habit specification in our DSGE model and found that the term premium remains quite small—on the order of a few basis points for the parameterizations in Table 1. In contrast, Jaccard reports a sizable equity premium without extreme macroeconomic fluctuations by imposing a high utility curvature ($\gamma = 10$) and steady-state importance of habits (\bar{h} equal to 99.7 percent of $\bar{C}v(\bar{L})$). However, Jaccard's parameterization implies an effective coefficient of relative risk aversion of about 3,000 for gambles over the consumption–leisure aggregate.

²² We have endeavored, without success, to find a parameterization that can deliver a large term premium and plausible real wage volatility. For example, even after allowing for parameter variation of the type shown in Table 1 and lower adjustment costs, the volatility of the real wage is two orders of magnitude too large.

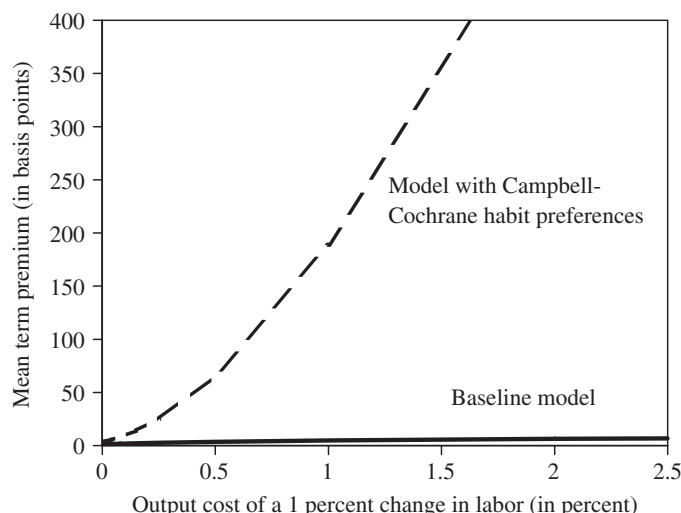


Fig. 1. Mean term premium in DSGE models with varying amounts of labor adjustment costs. The solid line shows the mean 10-year term premium in the baseline version of the benchmark model. The dashed line shows the mean term premium in the version of the model with Campbell–Cochrane habit preferences.

Table 3
Unconditional moments in three versions of the benchmark DSGE model

Variable	Campbell–Cochrane	C–C with quadratic adj. costs to labor	C–C with quadratic adj. costs to labor and real wage rigidity
sd[C]	1.11	0.89	4.54
sd[Y]	0.71	0.59	2.86
sd[L]	2.88	3.60	3.65
sd[w ^r]	2.14	220.9	1.90
sd[π]	2.25	19.7	5.98
sd[i]	2.05	7.66	5.94
sd[r]	1.96	21.4	2.81
sd[i ⁽¹⁰⁾]	0.57	1.19	3.84
Mean[ψ ⁽¹⁰⁾]	0.027	0.640	0.168
sd[ψ ⁽¹⁰⁾]	0.001	0.095	0.001
Mean[i ⁽¹⁰⁾ – i]	–0.035	0.004	0.004
sd[i ⁽¹⁰⁾ – i]	1.56	6.52	2.43
Mean[x ⁽¹⁰⁾]	0.033	0.729	1.46
sd[x ⁽¹⁰⁾]	7.23	23.03	23.52
β _{CS} ⁽¹⁰⁾	0.99	1.01	0.98

Except for the Campbell–Shiller regression coefficient, β_{CS}⁽¹⁰⁾, all variables are quarterly values expressed in percent. Inflation and interest rates, the term premium (ψ), and excess holding period returns (x) are expressed at an annual rate.

means that agents do not want to vary *either* labor *or* consumption in response to a shock. Yet when there is a shock, one or the other of these two quantities must give; as a result, the real wage must vary *tremendously* in order to achieve equilibrium. These large movements in the real wage in turn cause firms' marginal costs to be extremely volatile, which passes through to prices and inflation. The Taylor-type policy rule implies that the movements in inflation pass through to the short-term interest rate and the long-term bond yield. Both the marginal utility of consumption and the long-term bond price are much more volatile in this version of the model with adjustment costs, hence the term premium is much greater in magnitude.

5.2. Real wage rigidities

Instead of simple quadratic labor adjustment costs, other approaches to modeling labor market frictions might provide a better combination of macroeconomic and term premium fit. Real (and nominal) wage rigidities have been widely used in the macroeconomics literature and, following Blanchard and Galí (2005), we introduce a wage bargaining friction into our benchmark model. Specifically, we assume the real wage follows the process

$$\log w_t^r = \mu \log w_{t-1}^r + (1 - \mu)(\log w_t^{r*} + \omega), \tag{30}$$

where w^f denotes the real wage, w^{f*} denotes the frictionless real wage that would obtain in the absence of the wage rigidity, ω denotes a steady-state wedge between the real wage and households' marginal rate of substitution, and μ denotes the sluggishness of wages in adjusting toward the frictionless real wage. Although Eq. (30) does not explicitly model Nash bargaining between workers and firms, Blanchard and Galí motivate it as a simple friction that captures the essential features of real wage bargaining.

When we introduce Eq. (30) into our benchmark model, however, it turns out to have essentially no effect on the term premium, either in our baseline parameterization of the model or in our version with Campbell–Cochrane habits. Even in the version with C–C habits and $\mu = 0.99$ —an extremely rigid real wage—the mean term premium in the model increases from 2.7 to just 3.0 basis points. Setting $\mu = 0.999$ increases the term premium to only 3.4 basis points. Varying the parameters ω and μ over wide ranges has similarly small effects.

Intuitively, the real wage rigidity drives up the variability of w^{f*} and the household's marginal rate of substitution and stochastic discount factor. *Ceteris paribus*, increasing the variance of the stochastic discount factor should increase the magnitude of the term premium. However, the wage rigidity also makes firms' marginal costs much smoother than in the flexible-wage case; as a result, prices and inflation are much *less* volatile when there are wage rigidities than when there are not. The net effect of these two opposing forces is ambiguous, but for the wide range of parameterizations of the model we have considered, the net effect never amounted to more than a few basis points, far short of the magnitude we observe in the data.

Thus, real wage rigidities alone do not appear able to resolve the bond premium puzzle in our benchmark DSGE model, even when combined with Campbell–Cochrane habits. However, if quadratic adjustment costs to labor are also added to the mix, then perhaps the real wage rigidity would help to damp the excessive volatility of real wages of 221 percent that we saw previously. In the fourth column of Table 3, we report the mean term premium and unconditional standard deviations for the model with Campbell–Cochrane habits, quadratic adjustment costs to labor of $\kappa = 50\bar{Y}$, and with a real wage rigidity parameter of $\mu = 0.999$ (in the Campbell–Cochrane version of the model, the marginal rate of substitution is so volatile that only with extreme degrees of wage rigidity can the variation in real wages and inflation be brought back down to reasonable levels). While this extreme degree of wage rigidity does bring the standard deviations of the real wage and other macroeconomic variables back toward more reasonable levels, it also reduces the term premium, both in mean and standard deviation, to a point that is still about five times smaller than in the data. Thus, not only does the required degree of real wage rigidity appear to be implausibly large, but even assuming such wage rigidity, we are unable to fit both the term premium and macroeconomic variables.

5.3. Staggered nominal wage contracting

As an alternative to real wage frictions, we also consider staggered nominal wage contracts as in Erceg et al. (2000). Such contracts are prevalent in many medium- and large-scale DSGE models, such as Christiano et al. (2005) and Smets and Wouters (2003). Briefly, as in the Calvo price-contracting specification in the benchmark model, each household is now assumed to be a monopolistic supplier of a differentiated type of labor, which is bundled by a perfectly competitive labor market aggregator into the final labor input that is used by firms. A critical element that maintains tractability in the model is the assumption of complete financial markets, which allows households to trade state-contingent securities and ensures that—despite the heterogeneity across households in the wage charged and in hours worked—all households have identical wealth and consumption in every period in equilibrium. This assumption is standard in the literature because keeping track of a continuum of household-specific wealth holdings would be computationally intractable.

When we incorporate Calvo staggered nominal wage contracts into our benchmark model, either under our baseline parameterization or with Campbell–Cochrane habits, there is again no significant effect on the term premium. For example, in the Campbell–Cochrane version of the model, the term premium actually *decreases* from 2.7 to 1.1 basis points when we add Calvo wage contracts, and this result is robust when we vary the parameters of the model over wide ranges. Intuitively, the assumption of complete financial markets that is required for tractability in these models also has the side effect of allowing households to insure their consumption streams through financial markets. Thus, even though most households in the model cannot self-insure against negative shocks by working more hours, they can purchase state-contingent claims that pay off in the event of a negative income shock and in the event that the household is unable to reset its wage, which amounts to essentially the same thing. Because of this assumption, households are still able to insure their consumption streams from the consequences of negative shocks, and the term premium in the model remains very small.

6. Conclusions

All in all, our results cast a pessimistic light on the ability of habit-based DSGE models to fit the term premium. Even in versions of the model with large and persistent habits following Campbell and Cochrane (1999), the ability of households to vary their labor supply and thereby insure themselves against consumption fluctuations leads to term premiums that are far too small and far, far too stable relative to the data. Trying to reduce households' ability to self-insure by introducing standard labor market frictions into the model dramatically increases the volatility of households' marginal rate of substitution and the real wage—and hence marginal costs, inflation, and the short-term nominal interest rate—to a point that is far in excess of the data. Thus, the success that Wachter (2006) reports in fitting the term premium in an

endowment economy does not appear to generalize to the standard macroeconomic DSGE framework, where labor supply and production are endogenous.

While our results are somewhat discouraging from the point of view of habit-based DSGE models, there are other approaches still available that might allow one to fit the term premium in a DSGE framework. First, Piazzesi and Schneider (2006) have reported success in fitting the term premium in an endowment economy model using “generalized recursive” preferences, as in Epstein and Zin (1989). This approach may be more promising than habits in a DSGE setting because Epstein–Zin preferences separate the intertemporal elasticity of substitution from the coefficient of relative risk aversion, which should give the model a better chance of fitting both the macroeconomic quantities and asset prices. The methods employed in the present paper—such as second- and third-order approximations and a generalized consol to model the term premium—can likewise be applied to the case of recursive utility in a DSGE framework and make the solution of the term premium in such models computationally tractable, and we have begun to explore this variation in Rudebusch and Swanson (2008). Finally, models based on heterogeneous agents with incomplete insurance markets, as in Constantinides and Duffie (1996) and Krebs (2007), might also be able to explain the term premium, although merging heterogeneous-agent frameworks into standard macroeconomic DSGE models poses a significant computational challenge at present.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at [10.1016/j.jmoneco.2008.07.007](https://doi.org/10.1016/j.jmoneco.2008.07.007).

References

- Altig, D., Christiano, L., Eichenbaum, M., Lindé, J., 2004. Firm-specific capital, nominal rigidities, and the business cycle. Manuscript, Northwestern University.
- Alvarez, F., Jermann, U., 2001. Quantitative asset pricing implications of endogenous solvency constraints. *Review of Financial Studies* 14, 1117–1151.
- Ang, A., Piazzesi, M., 2003. No-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics* 50, 745–787.
- Backus, D., Gregory, A., Zin, S., 1989. Risk premiums in the term structure. *Journal of Monetary Economics* 24, 371–399.
- Bekaert, G., Cho, S., Moreno, A., 2005. New-Keynesian macroeconomics and the term structure. Manuscript, Columbia Business School.
- Blanchard, O., Gali, J., 2005. Real wage rigidities and the new Keynesian model. NBER Working Paper 11806.
- Boldrin, M., Christiano, L., Fisher, J., 2001. Habit persistence, asset returns, and the business cycle. *American Economic Review* 91, 149–166.
- Calvo, G., 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12, 383–398.
- Campbell, J., Cochrane, J., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205–251.
- Campbell, J.Y., Shiller, R.J., 1991. Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies* 58 (3), 495–514.
- Carrillo, J., Fève, P., Matheron, J., 2007. Monetary policy inertia or persistent shocks: a DSGE analysis. *International Journal of Central Banking* 3, 1–38.
- Christiano, L., Eichenbaum, M., Evans, C., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113, 1–45.
- Cochrane, J., 2007. Financial markets and the real economy. In: Mehra, R. (Ed.), *Handbook of the Equity Risk Premium*. Elsevier, Amsterdam, pp. 237–330.
- Cochrane, J.Y., Piazzesi, M., 2005. Bond risk premia. *American Economic Review* 95 (1), 138–160.
- Constantinides, G., Duffie, D., 1996. Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104, 219–240.
- Den Haan, W., 1995. The term structure of interest rates in real and monetary economies. *Journal of Economic Dynamics and Control* 19, 909–940.
- Diebold, F., Piazzesi, M., Rudebusch, G.D., 2005. Modeling bond yields in finance and macroeconomics. *American Economic Review, Papers and Proceedings* 95, 415–420.
- Doh, T., 2006. What moves the yield curve? Lessons from an estimated nonlinear macro model. Manuscript, University of Pennsylvania.
- Donaldson, J., Johnsen, T., Mehra, R., 1990. On the term structure of interest rates. *Journal of Economic Dynamics and Control* 14, 571–596.
- Epstein, L., Zin, S., 1989. Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57, 937–969.
- Erceg, C., Henderson, D., Levin, A., 2000. Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics* 46, 281–313.
- Gallmeyer, M.F., Hollifield, B., Zin, S.E., 2005. Taylor rules, McCallum rules and the term structure of interest rates. *Journal of Monetary Economics* 52, 921–950.
- Gurkaynak, R.S., Sack, B.P., Wright, J.H., 2007. The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54 (8), 2291–2304.
- Hördahl, P., Tristani, O., Vestin, D., 2006. A joint econometric model of macroeconomic and term structure dynamics. *Journal of Econometrics* 131, 405–444.
- Hördahl, P., Tristani, O., Vestin, D., 2007. The yield curve and macroeconomic dynamics. European Central Bank Working Paper 832.
- Jaccard, I., 2007. Asset pricing, habit memory, and the labor market. Unpublished manuscript, The Wharton School.
- Jermann, U., 1998. Asset pricing in production economies. *Journal of Monetary Economics* 41, 257–275.
- Kim, D.H., Wright, J.H., 2005. An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. Federal Reserve Board, Finance and Economics Discussion Series 2005-33.
- Krebs, T., 2007. Job displacement risk and the cost of business cycles. *American Economic Review* 97, 664–686.
- Lettau, M., Uhlig, H., 2000. Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics* 3, 79–99.
- Levin, A., Onatski, A., Williams, J.C., Williams, N., 2005. Monetary policy under uncertainty in micro-founded macroeconomic models. NBER Macro Annual 229–287.
- Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. *Journal of Monetary Economics* 15, 145–161.
- Piazzesi, M., Schneider, M., 2006. Equilibrium yield curves. NBER Macro Annual 389–442.
- Ravenna, F., Seppälä, J., 2006. Monetary policy and the term structure of interest rates. Manuscript, University of California, Santa Cruz.
- Rudebusch, G.D., 2002. Term structure evidence on interest rate smoothing and monetary policy inertia. *Journal of Monetary Economics* 49, 1161–1187.
- Rudebusch, G.D., 2006. Monetary policy inertia: fact or fiction? *International Journal of Central Banking* 2, 85–135.
- Rudebusch, G.D., Swanson, E., 2007. Examining the bond premium puzzle with a DSGE model. Federal Reserve Bank of San Francisco Working Paper 2007-25.
- Rudebusch, G.D., Swanson, E., 2008. The bond premium in a DSGE model with long-run risk. Manuscript, Federal Reserve Bank of San Francisco.

- Rudebusch, G.D., Wu, T., 2007. Accounting for a shift in term structure behavior with no-arbitrage and macro-finance models. *Journal of Money, Credit, and Banking* 39, 395–422.
- Rudebusch, G.D., Swanson, E., Wu, T., 2006. The bond yield “Conundrum” from a macro-finance perspective. *Monetary and Economic Studies* 24 (S-1), 83–128.
- Rudebusch, G.D., Sack, B., Swanson, E., 2007. Macroeconomic implications of changes in the term premium. *Federal Reserve Bank of St. Louis, Review* 89, 241–269.
- Smets, F., Wouters, R., 2003. An estimated stochastic dynamic general equilibrium model of the Euro area. *Journal of European Economic Association* 1, 1123–1175.
- Smith, J., Taylor, J.B., 2007. The long and the short end of the term structure of policy rules. Manuscript, Stanford University.
- Swanson, E., Anderson, G., Levin, A., 2006. Higher-order perturbation solutions to dynamic, discrete-time rational expectations models. Manuscript, Federal Reserve Bank of San Francisco.
- Taylor, J., 1980. Aggregate dynamics and staggered contracts. *Journal of Political Economy* 88, 1–23.
- Uhlig, H., 2007. Explaining asset prices with external habits and wage rigidities in a DSGE model. *American Economic Review, Papers and Proceedings* 97, 239–243.
- Wachter, J., 2006. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79, 365–399.
- Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics* 24, 401–421.
- Woodford, M., 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.
- Wu, T., 2006. Macro factors and the affine term structure of interest rates. *Journal of Money, Credit, and Banking* 38, 1847–1875.