A Macroeconomic Model of Equities and Real, Nominal, and Defaultable Debt

Eric T. Swanson

University of California, Irvine

eric.swanson@uci.edu

http://www.ericswanson.org

Abstract

Linkages between the real economy and financial markets can be extremely important, as demonstrated by the the recent global financial crisis and European sovereign debt crisis. In this paper, I develop a simple, structural macroeconomic model that is consistent with a wide variety of asset pricing facts, such as the size and variability of risk premia on equities, real and nominal government bonds, and corporate bonds—the equity premium puzzle, bond premium puzzle, and credit spread puzzle, respectively. I thus show how to unify a variety of asset pricing puzzles from finance into a simple, structural framework. Conversely, I show how to bring standard macroeconomic models into agreement with a wide range of asset pricing facts.

1. Introduction

Traditional macroeconomic models, such as Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007), ignore asset prices and risk premia and, in fact, do a notoriously poor job of matching financial market variables (e.g., Mehra and Prescott, 1985; Backus, Gregory, and Zin, 1989; Rudebusch and Swanson, 2008). At the same time, traditional finance models, such as Dai and Singleton (2003) and Fama and French (2015), ignore the real economy; even when these models use a stochastic discount factor or consumption rather than a latent factor framework, those economic variables are taken to be exogenous, reduced-form processes.

Despite this traditional separation, the linkages between the real economy and financial markets can be extremely important. During the 2007–09 global financial crisis and the ongoing European sovereign debt crisis, concerns about asset values caused lending and the real economy to plummet, while the deteriorating economy caused private-sector risk premia to increase and asset prices to spiral further downward (e.g., Mishkin, 2011; Gorton and Metrick, 2012; Lane, 2012). These crises also led to dramatic fiscal and monetary policy interventions that were well beyond the range of past experience.[1] Reduced-form finance models that perform well based on past empirical correlations may perform very poorly when those past correlations no longer hold, such as when there is a structural break or unprecedented policy intervention as observed during these crises. A structural macroeconomic model is more robust to these types of breaks and can immediately provide answers and insights into their possible effects on risk premia, financial markets, and the real economy. Macroeconomic models can also provide useful intuition about *why* output, inflation, and asset prices co-move in certain ways and how that comovement may change in response to policy interventions and structural breaks.

In the present paper, I develop a simple, structural macroeconomic model that is consistent with a wide range of asset pricing facts, such as the size and variability of risk premia on equities and real, nominal, and defaultable debt. Thus, unlike traditional macroeconomic models, the model I present here is able to match asset prices and risk premia remarkably well. Unlike traditional finance models, the model I develop here can give us insight into the effects of novel policy

---

[1] For example, the U.S. Treasury bought large equity stakes in automakers and financial institutions, and insured money market mutual funds to prevent them from "breaking the buck." The Federal Reserve purchased very large quantities of longer-term Treasury and mortgage-backed securities and gave explicit forward guidance about the likely path of the federal funds rate for years into the future. See, e.g., Mishkin (2011) and Gorton and Metrick (2012). For Europe, the European Union established the European Stability Mechanism to provide quick financial backing to member countries in need, while the European Central Bank provided large and unprecedented three-year loans to banks and announced that it would purchase large quantities of euro area members' bonds if the yields on those bonds became excessively stressed (e.g., *Wall Street Journal*, 2012; European Central Bank, 2012).

interventions and structural breaks on asset prices, and provides a unified structural explanation for the behavior of risk premia on a variety of assets.

The model has two essential ingredients: generalized recursive preferences (as in Epstein and Zin, 1989, Weil, 1989, and Tallarini, 2000) and nominal rigidities (as in the textbook New Keynesian models of Woodford, 2003, and Galí, 2008). Generalized recursive preferences allow the model to generate substantial risk premia without greatly distorting the behavior of macroeconomic aggregates. Nominal rigidities allow the model to match the behavior of inflation, nominal interest rates, and nominal assets such as Treasuries and corporate bonds.

My results have important implications for both macroeconomics and finance. For macroeconomics, I show how standard dynamic structural general equilibrium (DSGE) models can be brought into agreement with a wide variety of asset pricing facts. I thus address Cochrane's (2008) critique that a failure of macroeconomic models to match even basic asset pricing facts is a sign of fundamental flaws in those models.[2] Moreover, bringing those models into better agreement with asset prices makes it possible to use those models to study the linkages between risk premia, financial markets, and the real economy.

For finance, I unify a variety of asset pricing puzzles into a simple, structural framework. This framework can then be used to study the relationships between the different puzzles with each other and with the economy. For example, Backus, Gregory, and Zin (1989), Donaldson, Johnsen, and Mehra (1990), and Den Haan (1995) argue that the yield curve ought to slope downward on average because interest rates tend to be low during recessions, implying that bond prices are high when consumption is low (which would lead to an insurance-like, negative risk premium). According to the model here, the nominal yield curve can slope upward even if the real yield curve slopes downward if technology shocks (or other "supply" shocks) are an important source of economic fluctuations. Technology shocks cause inflation to rise when consumption falls, so that long-term nominal bonds *lose* rather than gain value in recessions, implying a positive risk premium. Similarly, the model developed here can be used to study the changes in correlations between stock and bond returns documented by Baele, Bekaert, and Inghelbrecht (2010), Campbell, Sundaram, and Viceira (2013), and others.

Previous macroeconomic models of asset prices have tended to focus exclusively on a single

---

[2] As Cochrane (2008) points out, asset markets are the mechansim by which marginal rates of substitution are equated to marginal rates of transformation in a macroeconomic model. If the model is wildly inconsistent with basic asset pricing facts, then by what mechanism does the model equate these marginal rates of substitution and transformation?

type of asset, such as equities (e.g., Boldrin, Christiano, and Fisher, 2001; Tallarini, 2000; Guvenen, 2009; Barillas, Hansen, and Sargent, 2009) or debt (e.g., Rudebusch and Swanson, 2008, 2012; Van Binsbergen et al., 2012; Andreasen, 2012). A disadvantage of this approach is that it's unclear whether the results in each case generalize to other asset classes. For example, Boldrin, Christiano, and Fisher (2001) show that capital immobility in a two-sector DSGE model can fit the equity premium by making the price of capital (and equity) more volatile, but this mechanism does not explain subtantial risk premia on long-term government bonds, which involve the valuation of a fixed stream of coupon payments. By focusing on multiple asset classes, I impose additional discipline on the model and ensure that its results apply more generally. Matching the behavior of a variety of assets also helps to identify model parameters, since different types of assets are relatively more informative about different aspects of the model. For example, nominal assets are helpful for identifying parameters related to inflation, and long-lived equities provide information about the longer-run features of the model.

A number of recent papers have begun to study stock and bond prices jointly in a traditional affine framework (e.g., Eraker, 2008; Bekaert, Engstrom, and Grenadier, 2010; Lettau and Wachter, 2011; Ang and Ulrich, 2013; Koijen, Lustig, and Van Nieuwerburgh, 2013).[3] Some of these studies work with latent factors, ignoring the real economy, while others relate asset prices to the reduced-form behavior of consumption. However, none of them uses a structural macroeconomic model, which has the advantages described above. Although reduced-form models often fit the data better than structural macroeconomic models, this can simply be a tautological implication of Roll's (1977) critique (that any mean-variance efficient portfolio perfectly fits the mean returns of all assets), as noted by Cochrane (2008). It is only the correspondence of financial risk factors to plausible economic risks that makes reduced-form financial factors interesting.

Chen, Collin-Dufresne, and Goldstein (2009), Bhamra, Kuehn, and Strebulaev (2010), and Chen (2010) model equity and corporate bond prices jointly in an endowment economy. Those authors undertake a more detailed, structural analysis of the corporate financing decision than I consider here, but at a cost of working in a much simpler, reduced-form macroeconomic environment. In other words, I use a simple, reduced-form model of the firm in order to better focus on the structural behavior of the economy, while Chen et al. (2010), Bhamra et al. (2010), and Chen (2010) use a simple, reduced-form model of the macroeconomy to better focus on the structural

---

[3] See also Campbell, Sundaram, and Viceira (2012), who price stocks and bonds jointly in a quadratic latent-factor framework.

finance behavior of the firm. The advantages of the structural macroeconomic approach I take here are discussed above.

The three papers most closely related to the present paper are Tallarini (2001), Rudebusch and Swanson (2012), and Campbell, Pflueger, and Viceira (2013). Tallarini (2001) incorporates Epstin-Zin-Weil preferences into a real business cycle model to match the equity premium. Relative to Tallarini (2001), the model here matches nominal as well as real features of the economy; explains the behavior of multiple assets, such as equities and real, nominal, and defaultable debt; and works within the now-standard New Keynesian DSGE framework rather than a real business cycle framework, which allows monetary policy to play a potentially important role in business cycles. Following Tallarini (2001), Rudebusch and Swanson (2012) incorporate Epstein-Zin-Weil preferences into a standard New Keynesian DSGE model to match the behavior of nominal long-term bonds.[4] In contrast to Rudebusch and Swanson (2012), the model here is much simpler (to clarify its essential features) and is extended to match the behavior of multiple asset classes. Campbell, Pflueger, and Viceira (2013, henceforth CPV) study stock and bond prices in a reduced-form three-equation New Keynesian model. In contrast to the present paper, CPV use a stochastic discount factor that is related to their New Keynesian model only in an ad hoc, reduced-form manner.[5] In fact, the ad hoc connection between the stochastic discount factor and the economy is crucial for CPV's results: as shown by Lettau and Uhlig (2000) and Rudebusch and Swanson (2008), CPV's Campbell-Cochrane (1999) habit specification is typically unable to produce significant risk premia when households can endogenously smooth consumption, because households endogenously choose a path for consumption that is so smooth the stochastic discount factor is stabilized.[6] In the present paper, I undertake a more structural approach, specifying a complete—but simple—macroeconomic model in which the stochastic discount factor is internally consistent with the other equations of the model.

Throughout the present paper, a recurring theme is the simplicity of the model, in the

---

[4] See also Van Binsbergen et al. (2012) and Andreasen (2012) for variations on the analysis in Rudebusch and Swanson (2012).

[5] In this respect, their analysis is similar to the term-structure studies of Rudebusch and Wu (2007) and Bekaert, Cho, and Moreno (2010), both of whom also use a reduced-form three-equation New Keynesian framework and an ad hoc stochastic discount factor.

[6] Households with Campbell-Cochrane (1999) habits are extremely averse to high-frequency fluctuations in consumption. In a DSGE model (as opposed to an endowment economy), households can self-insure themselves from these fluctuations by varying their hours of work or savings. In fact, for plausible parameterizations of DSGE models, households endogenously choose a path for consumption that is so smooth the stochastic discount factor does not vary much more than in the model without habits, leading risk premia to be about the same as without habits. See Rudebusch and Swanson (2008) and Lettau and Uhlig (2000).

interest of clarity and to help provide intuition for the underlying mechanisms. Thus, the model here is not designed to match very detailed features of the economy or asset prices; indeed, if one pushes the model far enough, it is certain to fail at matching some features of financial markets or the macroeconomy. (In a way, that failure is by design, due to the model's emphasis on simplicity.) Thus, the model here should be viewed as a "proof of concept" that the standard New Keynesian DSGE framework can be adapted to match asset prices quite well and shows a great deal of promise for future development in this direction. The approach I take here is thus analogous to Kydland and Prescott (1982), who showed that the stochastic growth model could be extended to match key features of business cycle fluctuations. Their stylized model failed to match many details of business cycles (e.g., unemployment, inflation), but opened the door to the equilibrium modeling of these phenomena.

The remainder of the paper proceeds as follows. In Section 2, I develop a simple New Keynesian DSGE model with nominal rigidities and Epstein-Zin preferences, show how to solve the model, and discuss the calibration of the model and its implications for macroeconomic quantities. In Section 3, I derive the prices of stocks and real, nominal, and defaultable bonds within the framework of the model, and compare the behavior of those asset prices to the data. In Section 4, I show how the model generates endogenous conditional heteroskedasticity, which is crucial for producing time-varying risk premia. Section 5 provides additional analysis and discussion related to issues raised in Sections 2 and 3. Section 6 concludes. Three Appendices present all the equations of the model, discuss the numerical solution method in more detail, and provide additional figures and analysis of the basic results.

2. A Simple Macroeconomic Model

In this section, I develop a simple dynamic macroeconomic model with generalized recursive preferences and nominal rigidities. Generalized recursive preferences, as in Epstein and Zin (1989) and Weil (1989), are required for the model to match the size of risk premia in the data.[7] Nominal rigidities are necessary for the model to match the basic behavior of inflation, nominal interest rates, and nominal assets such as Treasuries and corporate bonds.

Throughout this section, I strive to keep the model as simple as possible while still matching the essential behavior of macroeconomic variables and asset prices. The goal is to maximize

---

[7] See the previous footnote and Rudebusch and Swanson (2008) for a discussion of why habits in household preferences, such as Campbell and Cochrane (1999), are unable to match the size of risk premia in DSGE models.

intuition and insight into the relationships between the macroeconomy and asset prices, and avoid tangential complications. For this reason, I deliberately follow the very simple, "textbook" New Keynesian models of Woodford (2003) and Galí (2008), extended to the case of generalized recursive preferences. In principle, more realistic, medium-scale New Keynesian models such as Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007) could also be extended to the case of Epstein-Zin preferences to achieve an even better empirical fit to the data, but at the cost of being much more complicated.

2.1 Households

Time is discrete and continues forever. There is a unit continuum of representative households, each with generalized recursive preferences as in Epstein and Zin (1989) and Weil (1989). In each period t, the representative household receives the utility flow

$$u(c_t, l_t) \equiv \log c_t - \eta \frac{l_t^{1+\chi}}{1+\chi}, \tag{1}$$

where c_t and l_t denote household consumption and labor in period t, respectively, and $\eta > 0$ and $\chi > 0$ are parameters. Note that equation (1) differs from Epstein and Zin (1989) and Weil (1989) in that period utility depends on labor as well as consumption.

The assumption of additive separability in (1) follows Woodford (2003) and Galí (2008) and simplifies many aspects of the model. For example, the household's intertemporal elasticity of substitution is unity, its Frisch elasticity of labor supply is $1/\chi$, and its stochastic discount factor (defined below) is related to c_{t+1}/c_t; without additive separability, the expressions for these quantities would all be much more complicated. The similarity of the stochastic discount factor to versions of the model without labor also facilitates comparison to the finance literature. In addition, assuming logarithmic preferences over consumption ensures that the model is consistent with balanced growth (King, Plosser, and Rebelo, 1988, 2002) and is a standard benchmark in macroeconomics (e.g., King and Rebelo, 1999).

Households can borrow and lend in a default-free one-period nominal bond market at the continuously-compounded interest rate i_t. The use of continuous compounding simplifies the bond-pricing equations below and enhances comparability to the finance literature. Each period, the household faces a flow budget constraint

$$a_{t+1} = e^{i_t} a_t + w_t l_t + d_t - c_t, \tag{2}$$

where a_t denotes beginning-of-period nominal assets and w_t and d_t denote the nominal wage and exogenous transfers to the household, respectively. The household faces a standard no-Ponzi-scheme constraint,

$$\lim_{T \to \infty} E_t \prod_{\tau=t}^{T} e^{-i_{\tau+1}} a_{T+1} \geq 0. \tag{3}$$

Let $(c^t, l^t) \equiv \{(c_\tau, l_\tau)\}_{\tau=t}^{\infty}$ denote a state-contingent plan for household consumption and labor from time t onward, where the explicit state-dependence of the plan is suppressed to reduce notation. Following Epstein and Zin (1989), Weil (1989), and Tallarini (2000), I assume that the household has preferences over state-contingent plans ordered by the recursive functional

$$\widetilde{V}(c^t, l^t) \;=\; (1-\beta)\, u(c_t, l_t) \;-\; \beta\alpha^{-1}\log\left[E_t\,\exp\left(-\alpha\widetilde{V}(c^{t+1}, l^{t+1})\right)\right], \tag{4}$$

where $\beta \in (0,1)$ and $\alpha \in \mathbb{R}$ are parameters, E_t denotes the mathematical expectation conditional on the state of the economy at time t, and (c^{t+1}, l^{t+1}) denotes the state-contingent plan (c^t, l^t) from date $t+1$ onward. The case $\alpha = 0$ in (4) is defined by letting $\alpha \to 0$ and corresponds to the special case of expected utility preferences. When $\alpha \neq 0$ in (4), the expectation operator is effectively "twisted" and "untwisted" by the exponential function with coefficient $-\alpha$. This leaves the household's intertemporal elasticity of substitution in (4) the same as for expected utility, but amplifies (or attenuates) the household's risk aversion with respect to gambles over future utility flows by the additional curvature parameter α, with larger values of α corresponding to greater risk aversion. Thus, generalized recursive preferences allow the household's intertemporal elasticity of substitution and coefficient of relative risk aversion to be parameterized independently. Following Hansen and Sargent (2001), the specific form of generalized recursive preferences in (4) is often referred to as "multiplier preferences".

In each period, the household maximizes (4) subject to the budget constraint (2)–(3). The state variables of the household's optimization problem are a_t and Θ_t, where the latter is a vector denoting the state of the aggregate economy at time t. The household's "generalized value function" $V(a_t; \Theta_t)$ satisfies the generalized Bellman equation

$$V(a_t; \Theta_t) \;=\; \max_{(c_t, l_t)}\, (1-\beta)\, u(c_t, l_t) \;-\; \beta\alpha^{-1}\log\left[E_t\,\exp\left(-\alpha V(a_{t+1}; \Theta_{t+1})\right)\right], \tag{5}$$

where a_{t+1} is given by (2).

It's straightforward to show (e.g., Rudebusch and Swanson, 2012), that the household's

stochastic discount factor is given by

$$m_{t+1} \equiv \beta \frac{c_t}{c_{t+1}} \frac{\exp\left(-\alpha V(a_{t+1}; \Theta_{t+1})\right)}{E_t \exp\left(-\alpha V(a_{t+1}; \Theta_{t+1})\right)}. \tag{6}$$

Let r_t denote the one-period continuosly-compounded risk-free real interest rate. Then

$$e^{-r_t} = E_t m_{t+1}. \tag{7}$$

2.2 Firms

The economy also contains a continuum of infintely-lived monopolistically competitive firms indexed by $f \in [0,1]$, each producing a single differentiated good. Firms hire labor from households in a competitive market and have identical Cobb-Douglas production functions,

$$y_t(f) = A_t k^{1-\theta} l_t(f)^\theta, \tag{8}$$

where $y_t(f)$ denotes firm f's output, A_t is aggregate productivity affecting all firms, k and $l_t(f)$ denote the firm's capital and labor inputs at time t, respectively, and $\theta \in (0,1)$ is a parameter. For simplicity, and following Woodford (2003) and Galí (2008), I assume that firms' capital stocks are fixed, so that labor is the only variable input to production. Intuitively, movements in the capital stock are small at business-cycle frequencies and are dominated by fluctuations in labor.[8]

Technology, A_t, follows an exogenous AR(1) process,

$$\log A_t = \rho_A \log A_{t-1} + \varepsilon_t^A, \tag{9}$$

where $\rho_A \in (-1, 1]$, and ε_t^A denotes an i.i.d. white noise process with mean zero and variance σ_A^2. For simplicity and comparability to the finance literature, I set $\rho_A = 1$ in the baseline calibration of the model, below, but consider alternative values of ρ_A as well. For simplicity and ease of exposition, I abstract from technology growth in the baseline calibration of (9) as well (i.e., the mean of $\log(A_t/A_{t-1})$ is 0).[9]

---

[8] Woodford (2003, p. 167) compares a model with fixed firm-specific capital to a model with endogenous capital and investment adjustment costs and finds that the basic business-cycle features of the two models are very similar. In models with endogenous capital (e.g., Christiano et al., 2005; Smets and Wouters, 2007; Altig et al., 2011), investment adjustment costs are typically included to keep the capital stock stable at higher frequencies. Thus, one can think of the fixed-capital assumption as a simple way of achieving the same result. Woodford (2003) and Altig et al. (2011) also show that firm-specific capital stocks help generate inflation persistence that is consistent with the data (see particularly Woodford, 2003, pp. 163-173).

[9] If the mean rate of technology growth is μ_A, then the firm-specific capital stocks k must also grow at rate μ_A/θ in order for the model to have balanced growth.

Firms set prices optimally subject to nominal rigidities in the form of Calvo (1983) price contracts, which expire with exogenous probability $1 - \xi$ each period, $\xi \in [0, 1)$. Each time a Calvo contract expires, the firm sets a new contract price $p_t^*(f)$ freely, which then remains in effect for the life of the new contract, with indexation to the (continuously-compounded) steady-state inflation rate $\overline{\pi}$ each period.[10] In each period $\tau \geq t$ that the contract remains in force, the firm must supply whatever output is demanded at the contract price $p_t^*(f)e^{(\tau - t)\overline{\pi}}$, hiring labor $l_\tau(f)$ from households at the market wage w_τ.

Firms are jointly owned by households and distribute all profits and losses back to households each period in an aliquot, lump-sum manner. When a firm's price contract expires, the firm chooses the new contract price $p_t^*(f)$ to maximize the value to shareholders of the firm's cash flows over the lifetime of the contract,[11]

$$\sum_{j=0}^{\infty} \xi^j E_t \left\{ m_{t,t+j}(P_t/P_{t+j}) \left[p_t^*(f)e^{j\overline{\pi}} y_{t+j}(f) - w_{t+j} l_{t+j}(f) \right] \; \middle| \; p_{t+j}(f) = p_t^*(f) \right\}, \qquad (10)$$

where $m_{t,t+j} \equiv \prod_{i=1}^{j} m_{t+i}$ denotes shareholders' stochastic discount factor from period $t+j$ back to t, P_t the aggregate price level (defined below), w_t the nominal wage at time t, and $y_{t+j}(f)$ and $l_{t+j}(f)$ denote the firm's output and labor in period $t + j$, respectively, conditional on the contract price $p_t^*(f)$ still being in effect.

The output of each firm f is purchased by a perfectly competitive final goods sector, which aggregates the differentiated goods into a single final good using a CES production technology,

$$Y_t = \left[\int_0^1 y_t(f)^{1/\lambda} df \right]^\lambda, \qquad (11)$$

where Y_t denotes the quantity of the final good and $\lambda > 1$ is a parameter. Each intermediate firm f thus faces a downward-sloping demand curve for its product with elasticity $\lambda/(1 - \lambda)$,

$$y_t(f) = \left(\frac{p_t(f)}{P_t} \right)^{-\lambda/(\lambda-1)} Y_t, \qquad (12)$$

where $p_t(f)$ denotes the price in effect for firm f at time t (so $p_t(f) = p_\tau^*(f)e^{(t-\tau)\overline{\pi}}$, letting $\tau \leq t$ denote the most recent period in which firm f reset its contract price), and P_t is the CES

---

[10] The assumption of indexation keeps the model well-behaved with respect to changes in steady-state inflation. The continuous compounding is notationally simpler for some of the equations below.

[11] Equivalently, the firm can be viewed as choosing a state-contingent plan for prices that maximizes the value of the firm to shareholders.

aggregate price of the final good,

$$P_t \equiv \left[\int_0^1 p_t(f)^{1/(1-\lambda)} df \right]^{1-\lambda}. \tag{13}$$

Differentiating (10) with respect to $p_t^*(f)$ and setting the derivative equal to zero yields the standard New Keynesian price optimality condition,

$$p_t^*(f) = \lambda \frac{\sum_{j=0}^{\infty} \xi^j E_t \{ m_{t,t+j}(P_t/P_{t+j}) y_{t+j}(f) \mu_{t+j}(f) \mid p_{t+j}(f) = p_t^*(f) \}}{\sum_{j=0}^{\infty} \xi^j E_t \{ m_{t,t+j}(P_t/P_{t+j}) \xi^j y_{t+j}(f) e^{j\bar{\pi}} \mid p_{t+j}(f) = p_t^*(f) \}}, \tag{14}$$

where $\mu_t(f)$ denotes the (nominal) marginal cost for firm f at time t,

$$\mu_t(f) \equiv \frac{w_t l_t(f)}{\theta y_t(f)}. \tag{15}$$

That is, the firm's optimal contract price $p_t^*(f)$ is a monopolistic markup λ over a discounted weighted average of expected future marginal costs over the lifetime of the contract.[12]

2.3 Aggregate Resource Constraints and Government

Let L_t denote the aggregate quantity of labor demanded by firms,

$$L_t = \int_0^1 l_t(f) df. \tag{16}$$

Then L_t satisfies

$$Y_t = \Delta_t^{-1} A_t K^{1-\theta} L_t^{\theta}, \tag{17}$$

where $K = k$ denotes the aggregate capital stock and

$$\Delta_t \equiv \left[\int_0^1 \left(\frac{p_t(f)}{P_t} \right)^{\lambda/((1-\lambda)\theta)} df \right]^{\theta} \tag{18}$$

measures the cross-sectional dispersion of prices across firms. Δ_t has a minimum value of unity when $p_t(f) = P_t$ for all firms f; a greater degree of cross-sectional price dispersion increases Δ_t and reduces the economy's efficiency at producing final output.[13]

Labor market equilibrium requires that $L_t = l_t$, firms' labor demand equals the aggregate labor supplied by households. Equilibrium in the final goods market requires $Y_t = C_t$, where

---

[12] To be more precise, $p_t^*(f)$ is a weighted average of marginal costs deflated by the inflation index rate, $\mu_{t+j}(f)/e^{j\bar{\pi}}$. In addition, the weights in (14) depend on $y_{t+j}(f)$, which depend on the left-hand-side variable $p_t^*(f)$, so (14) is not a closed-form solution for $p_t^*(f)$. However, the closed-form solution for $p_t^*(f)$, reported in the Appendix, has the same form as (14).

[13] See Appendix C for additional discussion of price dispersion in the model.

$C_t = c_t$ denotes aggregate consumption demanded by households. For simplicity, there are no government purchases or investment in the baseline version of the model.

Finally, there is a monetary authority that sets the one-period nominal interest rate i_t according to a Taylor (1993)-type policy rule,

$$i_t = r + \pi_t + \phi_\pi(\pi_t - \overline{\pi}) + \frac{\phi_y}{4}(y_t - \overline{y}_t), \tag{19}$$

where $r = -\log\beta$ denotes the continuously-compounded steady-state real interest rate, $\pi_t \equiv \log(P_t/P_{t-1})$ denotes the inflation rate, $\overline{\pi}$ the monetary authority's inflation target, $y_t \equiv \log Y_t$,

$$\overline{y}_t \equiv \rho_{\overline{y}}\overline{y}_{t-1} + (1 - \rho_{\overline{y}})y_t \tag{20}$$

denotes a trailing moving average of log output, and $\phi_\pi, \phi_y \in \mathbb{R}$ and $\rho_{\overline{y}} \in [0, 1)$ are parameters.[14] The term $(\pi_t - \overline{\pi})$ in (19) represents the deviation of inflation from policymakers' target and $(y_t - \overline{y}_t)$ is a measure of the "output gap" in the model.[15]

2.4 Solution Method

I solve the model above by writing each equation in recursive form, dividing nonstationary variables (Y_t, C_t, w_t, etc.) by A_t so that the resulting ratios have a stable nonstochastic steady state. I then use the method of local approximation around the nonstochastic steady state, or perturbation methods, to compute a numerical solution to the model.[16] For the complete set of recursive equations that define the model and additional discussion, see Appendix A.

Macroeconomic models similar to the one developed above are typically solved using a first-order approximation (a linearization or log-linearization), but this solution method reduces all

---

[14] Note that interest rates and inflation in (19) are at quarterly rather than annual rates, so ϕ_y corresponds to the sensitivity of the annualized short-term interest rate to the output gap, as in Taylor (1993). I also exclude a lagged interest rate "smoothing" term on the right-hand side of (19) for simplicity and to keep the number of state variables in the model to a minimum. Rudebusch (2002) argues that the degree of federal funds rate smoothing from one quarter to the next is essentially zero, and that instead the Federal Reserve's deviations from the Taylor rule (19) are serially correlated—i.e., that the residuals ε_t^i in the empirical version of (19) are serially correlated.

[15] This is an empirically motivated definition of the output gap: it implies that the central bank will raise short-term nominal interest rates when output rises above its recent history and lower rates when output falls below that history, all else equal. The behavior of monetary policy is very important for the sign and size of risk premia on nominal and real bonds in the model. In order for the model to match these risk premia in the data, it's important that monetary policy act in a way that is consistent with the data. Defining the output gap to be the deviation of output from flexible-price output implies that interest rates would behave in an opposite manner, and generally would not allow the model to match empirical risk premia on nominal and real bonds.

[16] The equity price p_t^e is normalized by A_t^ν rather than A_t, where ν denotes the degree of leverage (see below). The value function V_t is normalized by defining $\tilde{V}_t \equiv V_t - \log A_t$.

risk premia in the model to zero.[17] A second-order approximation to the model produces risk premia that are nonzero but constant over time (a constant function of the variance σ_A^2). In order for risk premia in the model to vary with the state of the economy, the model must be solved to at least third order around the steady state. Note that second- and third-order terms in the model solution can be non-negligible as long as the model is sufficiently "curved", which is the case when risk aversion (related to the Epstein-Zin parameter α) is sufficiently large.

I compute third- and higher-order solutions of the model using the Perturbation AIM algorithm of Swanson, Anderson, and Levin (2006), which can compute general nth-order Taylor series approximate solutions to discrete-time recursive rational expectations models. The model above has two state variables $(\Delta_t, \overline{y}_t)$ and a single shock (ε_{t+1}^A) and thus can be solved to third order very quickly, in just a few seconds on a laptop computer. To obtain greater accuracy over a wider range of values for the state variables, the model can be solved to higher order; the results reported below are for the fifth-order solution unless stated otherwise. (Results for fourth- and sixth-order solutions are very similar, suggesting that the Taylor series has essentially converged over the relevant range for the state variables.) Aruoba et al. (2006) compare a variety of numerical solution techniques for standard macroeconomic models and find that higher-order perturbation solutions are among the most accurate globally as well as being the fastest to compute. Swanson, Anderson, and Levin (2006) provide details of the algorithm and discuss the global convergence properties of nth-order Taylor series approximations.

A noteworthy feature of the nonlinear solution algorithm I use here, relative to the loglinear-lognormal approximation typically used in finance, is that second- and higher-order terms of the Taylor series display endogenous conditional heteroskedasticity. Letting x_t denote a generic state variable and ε_{t+1} a generic shock, the second-order Taylor series solution has terms of the form $x_t \varepsilon_{t+1}$, which have a one-period-ahead conditional variance that depends on the economic state x_t (that is, $\text{Var}_t(x_t \varepsilon_{t+1})$ depends on x_t). Thus, even though the model's exogenous driving shocks ε_{t+1}^A are homoskedastic, the nonlinear solution algorithm I use here preserves the endogenous conditional heteroskedasticity that is naturally generated by the nonlinearities in the model.

---

[17] In the finance literature, it is standard to log-linearize the model and then take expectations of all variables assuming joint lognormality. This approximate solution method produces nonzero (but constant) risk premia, but effectively treats higher-order moments of the lognormal distribution on par with first-order economic terms. Standard perturbation methods (e.g., Judd, 1998; Swanson, Anderson, and Levin, 2006) explicitly relate higher-order moments of the shock distribution to the corresponding order of the state variables (so variance is a second-order term, skewness a third-order term, etc.), because their magnitudes are the same in theory.

<div align="center">Table 1: Parameter Values, Baseline Calibration</div>

β	0.992	θ	0.6	ϕ_π	0.5
χ	3	ξ	0.8	ϕ_y	0.75
η	0.545	λ	1.1	$\bar{\pi}$	0.008
RRA (R^c)	60	ρ_A	1	$\rho_{\bar{y}}$	0.9
		σ_A	0.007		
		$K/(4Y)$	2.5		

2.5 Calibration

The model described above is meant to be illustrative rather than provide a comprehensive empirical fit to the data, so I calibrate rather than estimate its key parameters. The baseline calibration is reported in Table 1, and is meant to be standard, following along the lines of parameter values estimated by Smets and Wouters (2007), Altig et al. (2011), and Del Negro, Giannoni, and Schorfheide (2015) using quarterly U.S. data.

I set the household's discount factor, β, to .992, implying a nonstochastic steady-state real interest rate of a little more than 3 percent per year. Although this might seem a bit high, households' risk aversion drives the unconditional mean of the risk-free real rate close to 2 percent in the stochastic case.

The household's logarithmic preferences over consumption imply an intertemporal elasticity of substitution of unity, which is higher than estimates based on aggregate data (e.g., Hall, 1988), but is similar to estimates using household-level data (e.g., Vissing-Jorgensen, 2002). Bansal and Yaron (2004) and Dew-Becker (2012) argue that estimates based on aggregate data are biased downward, providing further support for the value of unity used here. In addition, logarithmic preferences over consumption are a standard benchmark in macroeconomics (e.g., King and Rebelo, 1999).[18]

The calibrated value of $\chi = 3$ implies a Frisch elasticity of labor supply of 1/3, consistent with estimates in Del Negro et al. (2015) and estimates from household data (e.g., MaCurdy, 1980; Altonji, 1986). I set the parameter η so as to normalize $L = 1$ in steady state.

The parameter α is calibrated to imply a coefficient of relative risk aversion $R^c = 60$ in steady state, using the closed-form expressions derived in Swanson (2013) for models with

---

[18] My results are not sensitive to setting the IES equal to unity. For example, specifications with $u(c_t, l_t) = c_t^{1-\gamma}/(1-\gamma) - \eta l^{1+\chi}/(1+\chi)$ produce very similar results when γ is set to 0.9 or 1.1. Of course, these specifications do not satisfy balanced growth and are nonstationary in response to permanent technology shocks.

labor.[19] Although this value is high, it is a well-known byproduct of the model's simplicity:[20] for example, households in the model have perfect knowledge of all the model equations, parameter values, shock dynamics, and shock distributions, so the quantity of risk in the model is very small relative to the actual U.S. economy. As a result, the household's aversion to risk in the model must be correspondingly larger to fit the risk premia seen in the data. Barillas, Hansen, and Sargent (2009) formalize this intuition by showing that high risk aversion in an Epstein-Zin specification is isomorphic to a model in which households have low risk aversion but a moderate degree of uncertainty about the economic environment.[21] As an alternative to high risk aversion, one could increase the quantity of risk in the model instead, such as by introducing long-run risk as in Bansal and Yaron (2004), or disaster risk as in Rietz (1988) and Barro (2006).[22]

Turning to the production side of the economy, I set the elasticity of output with respect to labor $\theta = 0.6$. I calibrate the Calvo contract parameter $\xi = 0.8$, implying an average contract duration of five quarters, consistent with the estimates in Altig et al. (2010) and Del Negro et al. (2015). I calibrate the monopolistic markup λ for intermediate goods to 1.1, consistent with the estimates in Smets and Wouters (2007) and Altig et al. (2010). The technology process A_t is calibrated to be a random walk in the baseline calibration, $\rho_A = 1$. The standard deviation of technology shocks, σ_A, is set to .007, following estimates in King and Rebelo (1999). The steady-state ratio of the capital stock to annualized output is calibrated to 2.5.

The response of monetary policy to inflation, ϕ_π, is set to 0.5, as in Taylor (1993, 1999). I set $\phi_y = 0.75$, between the values of 0.5 and 1 used by Taylor (1993) and Taylor (1999). I set the monetary authority's inflation target $\bar{\pi}$ to 0.8 percent per quarter, implying a nonstochastic

---

[19] Swanson (2013) derives the coefficient of relative risk aversion for generalized recursive preferences with flexible labor and arbitrary period utility function $u(c_t, l_t)$. For multiplier preferences with period utility function (1) and $l = 1$ in steady state, risk aversion is given by $R^c = \alpha + (1+(\eta/\chi))^{-1}$. See Swanson (2013) for the derivation and details. In general, risk aversion is lower when labor supply can vary because the household is better able to insure itself from shocks.

[20] For example, Piazzesi and Schneider (2006) estimate a value of 57, Rudebusch and Swanson (2012) a value of 110, Van Binsbergen et al. (2012), Andreasen (2012), and Campbell and Cochrane (1999) a value of about 80, and Tallarini (2000) a value of about 50. The nonstationarity of technology implied by $\rho_A = 1$ in the present paper increases the quantity of risk in the model here relative to Rudebusch and Swanson (2012), which allows me to use a lower coefficient of relative risk aversion here.

[21] See also Campanale, Castro, and Clementi (2010), who emphasize that the quantity of consumption risk in a standard DSGE model is very small, and thus the risk aversion required to match asset prices must be correspondingly larger.

[22] The simplifying representative-household assumption could also be dropped. Mankiw and Zeldes (1991), Parker (2001), and Malloy, Moskowitz, and Vissing-Jorgensen (2009) show that the consumption of stockholders is more volatile (and more correlated with the stock market) than the consumption of nonstockholders, so the required level of risk aversion in a representative-agent model is higher than it would be in a model that recognized that stockholders have more volatile consumption (Guvenen, 2009).

steady-state inflation rate of about 3.2 percent per year. Although this is higher than the value of about 2 percent used by many central banks as their current official inflation target, there are two reasons why a higher number is appropriate here: First, a steady state inflation rate of 2 percent is too low to explain the historical average level of nominal yields in the U.S. and U.K. (and many other countries), even over relatively recent samples such as 1990–2007, as I will show below. Second, households' risk aversion drives the unconditional mean of inflation in the stochastic version of the model somewhat below the nonstochastic steady-state value. Finally, I calibrate $\rho_{\bar{y}} = 0.9$, implying that the monetary authority uses the deviation of current output from its average level over the past roughly 2.5 years to approximate the output gap.

2.6 Impulse Response Functions

Figure 1 plots impulse response functions for the model to a one-standard-deviation (0.7 percent) positive technology shock, under the baseline calibration described above. Recall that, because $\rho_A = 1$, the effect of the shock on productivity is permanent. The dashed red lines in each panel report standard impulse response functions for the first-order (log-linear) solution to the model, while the solid blue lines report impulse response functions for the nonlinear, fifth-order Taylor series solution to the model.[23] I start by describing the linear impulse response functions (dashed red lines), and then describe how the fifth-order impulse response functions (solid blue lines) differ from their linear counterparts.

The top left panel of Figure 1 reports the impulse response function for consumption, C_t, to the shock. Consumption jumps upward on impact, as higher productivity increases the supply of output and makes households wealthier in present-value terms, increasing consumption demand. The first-order impulse response function for C_t does not jump all the way to its new long-run level on impact, however, because of the increase in the real interest rate (described shortly). Instead, consumption continues to increase gradually over time to approach its new steady state.

---

[23] The impulse response functions for the fifth-order solution to the model are computed as follows: The state variables of the model are initialized to their nonstochastic steady-state values. The impulse response function is computed as the period-by-period difference between a "one-shock" and a "no-shock" (baseline) scenario. In the one-shock scenario, ε_t^A is set equal to .007 in period 1, and equal to 0 from period 2 onward. In the no-shock scenario, ε_t^A is set equal to 0 in every period. Agents in the model do not have perfect foresight, so they still act in a precautionary manner even though the realized shocks turn out to be deterministically equal to 0 from period 1 onward. In principle, this nonlinear impulse response function can vary as one varies the initial point of the simulation, or may scale nonlinearly with the size of the shock ε_t^A. In practice, however, the impulse responses in Figure 2 do not vary much with the initial point and do not display much nonlinearity in the size of the shock. For example, the fifth-order impulse response functions to a *negative* 0.7 percent technology shock, which are reported in Figure B1 in Appendix B, look very similar to the negative of the blue lines in Figure 1 (although they are a bit larger in magnitude).
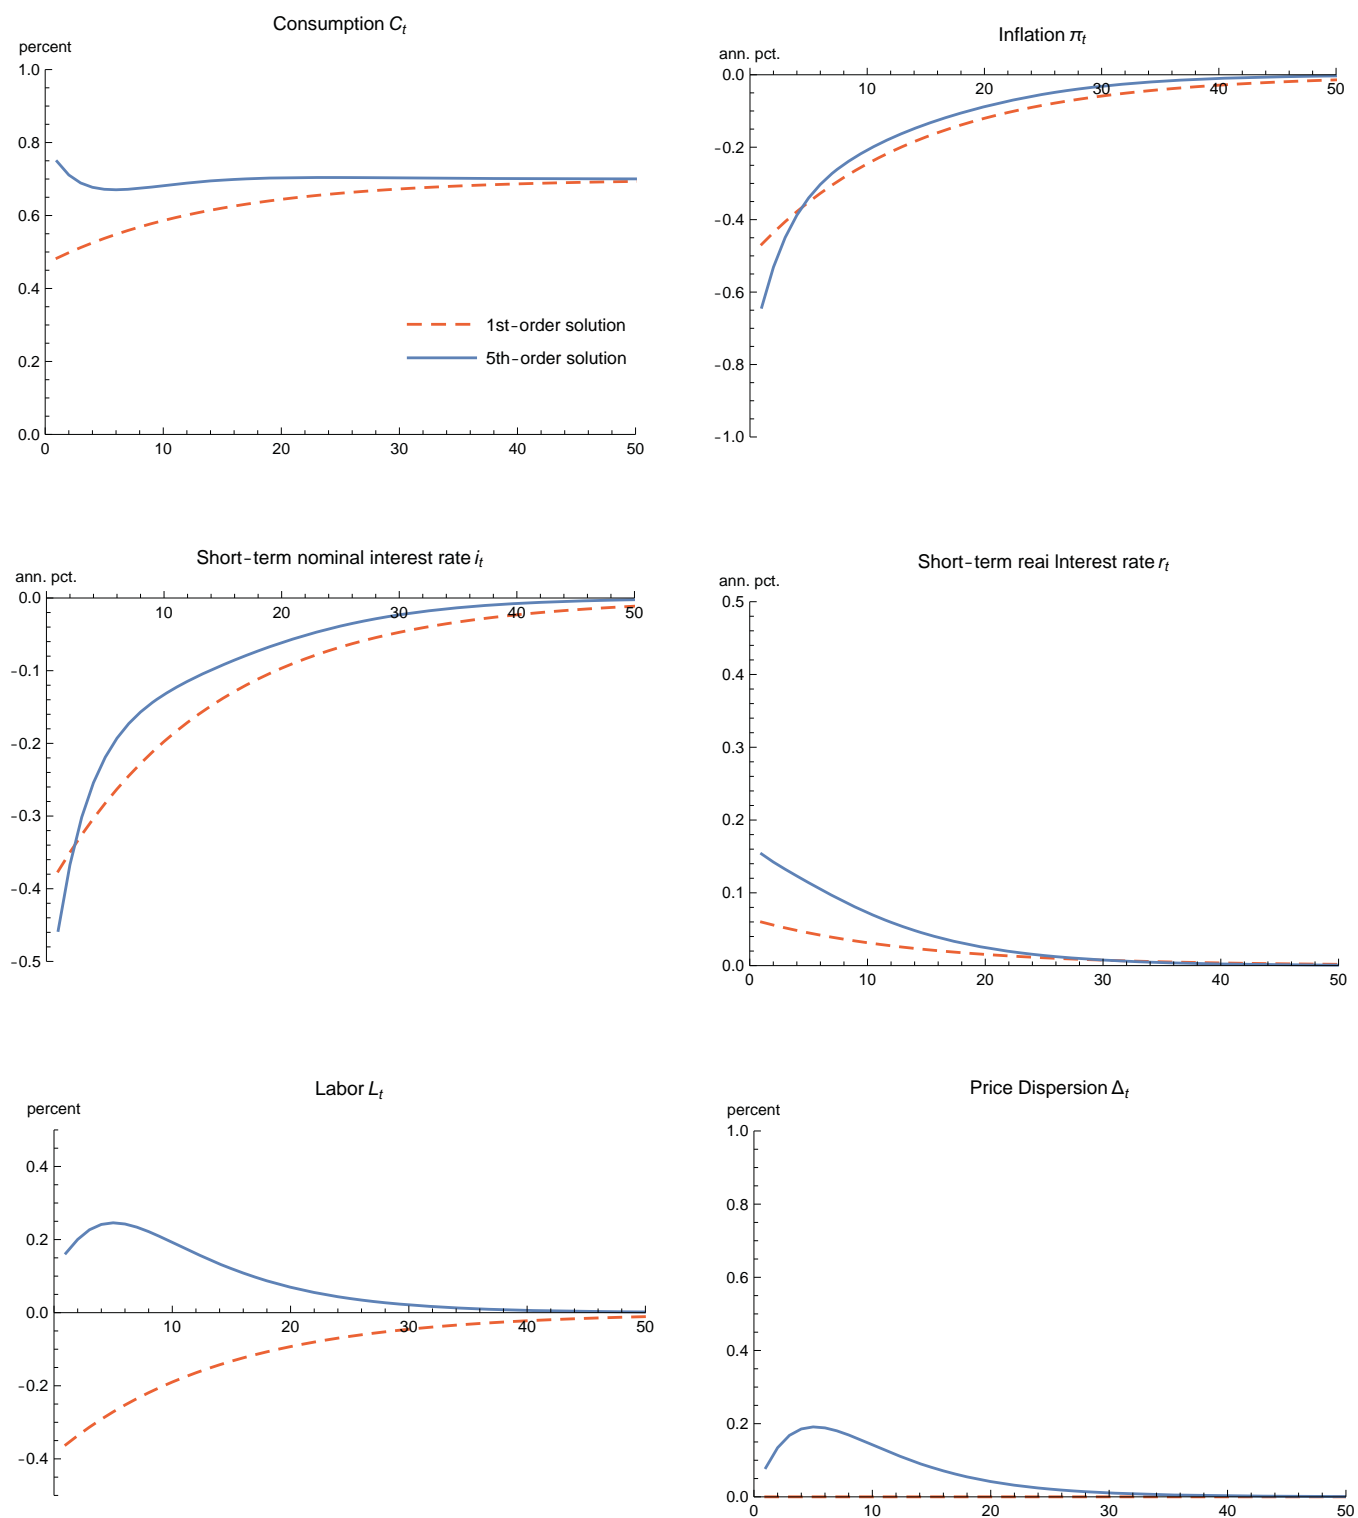
Figure 1. First-order (dashed red lines) and fifth-order (solid blue lines) impulse response functions for consumption C_t, inflation π_t, short-term nominal interest rate i_t, short-term real interest rate r_t, labor L_t, and price dispersion Δ_t to a one-standard-deviation (0.7 percent) positive technology shock in the model. See text for details.

The top right panel reports the impulse response for inflation, π_t. The higher level of technology reduces firms' marginal costs of production, and monopolistic firms set their price equal to a constant markup λ over expected future marginal costs, whenever they are able to reset their price. Thus, inflation falls on impact (by about 0.5 percent at an annualized rate) as those firms who are able to reset their prices do so. The response of inflation is persistent, however, as firms' price contracts expire only gradually.

The nominal interest rate i_t, in the middle left panel, is set by the monetary authority as a function of output and inflation according to the policy rule (19). Interest rates respond more strongly to inflation than to output, causing the nominal interest rate to decline moderately, on net, in response to the shock, about 40 basis points (at an annual rate) on impact before gradually returning to steady state. However, the nominal interest rate falls by less than inflation in response to the shock, so the real interest r_t *rises* about 5 basis points (at an annual rate) on impact, as can be seen in the middle right panel.[24] The real rate then gradually falls back to steady state.

The response of labor, L_t, is plotted in the bottom left panel. After the technology shock, households are wealthier in present value terms and want to consume more leisure; this tends to push labor downward. However, because prices are sticky and firms are monopolistic, firms hire whatever labor is necessary to satisfy output demand, which tends to push labor upward. For the simple model here, solved to first order, the former effect dominates, causing labor to decline slightly on net; indeed, this result is common in simple New Keynesian models, as pointed out by Galí (1999).[25]

However, this is no longer true for the fifth-order solution of the model, as can be seen by comparing the solid blue and red dashed lines in the bottom left panel. There are two main reasons for this difference: First, price dispersion Δ_t increases in response to the shock—as can be seen in the bottom right panel of Figure 1—but only for the nonlinear solution, because the linearized version of equation (18) implies shocks have no effect on Δ_t.[26] The increase in price dispersion reduces the economy's ability to produce final output efficiently, and increases the amount of labor required to produce any given level of output (see equation 17). Indeed, the hump shape

---

[24] Recall that r_t is the ex ante real interest rate, so $r_t = i_t - E_t\pi_{t+1}$ to first order.

[25] In more complicated and realistic models, such as Altig et al. (2011), increased demand for investment following the technology shock is typically enough to make the increase in firms' labor demand dominate. Alternatively, a stronger monetary policy response that drives the short-term real interest rate down in response to the shock would cause consumption to jump above 0.7 percent on impact and lead to an increase in labor.

[26] The linearized version of equation (18) is $\Delta_t = \xi\Delta_{t-1}$, which implies Δ_t is invariant to shocks.

in dispersion is clearly visible in the nonlinear impulse response function for labor (and to a lesser extent, consumption).[27] Second, the positive technology shock reduces the volatility of households' stochastic discount factor, for reasons discussed in detail in Section 4, below. The lower volatility of the SDF makes households effectively less risk averse and reduces their demand for precautionary savings, leading to an increase in consumption C_t relative to the linear case (as can be seen in the top-left panel). Households' greater demand for consumption requires firms to hire more labor, putting further upward pressure on L_t in the bottom-left panel.[28]

3. Asset Prices and Risk Premia

The stochastic discount factor implied by the simple macroeconomic model above can now be used to price any asset in the model. In particular, we can derive the implications of the model for the prices of equity and real, nominal, and defaultable debt.

3.1 Equity

I define an equity security in the model to be a levered claim on the aggregate consumption stream. The definition of equity as a consumption claim maximizes comparability to the finance literature and simplifies the intuition in the model; the results are very similar if equity is instead defined to be a claim on the profits of the monopolistic intermediate firm sector.[29] Each period, equity pays a dividend equal to C_t^ν, where ν is the degree of leverage. Consistent with Abel (1999), Bansal and Yaron (2004), and Campbell et al. (2013), I calibrate $\nu = 3$. Note that ν can be interpreted as the sum of operational and financial leverage in the economy, where operational leverage results from fixed costs of production for firms (Gourio, 2012; Campbell et al., 2013).

Let p_t^e denote the ex-dividend price of an equity security at time t. In equilibrium,

$$p_t^e = E_t m_{t+1}(C_{t+1}^\nu + p_{t+1}^e). \tag{21}$$

---

[27] The effect is essentially symmetric for a negative technology shock—i.e., Δ_t *decreases* in response to a negative technology shock (see Figure B1 in Appendix B). In the stochastic version of the model, inflation is often below the nonstochastic steady state value (due to precautionary behavior by firms, discussed below), so even if there are no shocks, $\Delta_t > 1$ will hold. Negative technology shocks can thus decrease Δ_t.

[28] The higher level of consumption in the nonlinear case also causes the real interest rate to rise by more in the middle-right panel. The higher level of labor in the nonlinear case increases firms' marginal costs, which puts upward pressure on inflation. Inflation nevertheless falls a bit more on impact in the nonlinear case for reasons discussed in Section 4, below. The response of the nominal interest rate i_t in the nonlinear case follows in a straightforward manner from C_t and π_t, given the policy rule (19).

[29] This is because consumption, output, and monopolistic firm profits are very highly correlated in the model: firms' profits are essentially a levered claim on the output stream, and $Y_t = C_t$. Note that adding fixed costs of production to the model would increase the degree of leverage.

Let R_{t+1}^e denote the realized gross return on equity,

$$R_{t+1}^e \equiv \frac{C_{t+1}^\nu + p_{t+1}^e}{p_t^e} \,. \tag{22}$$

I define the equity premium at time t, ψ_t^e, to be the expected excess return to holding equity for one period,

$$\psi_t^e \equiv E_t R_{t+1}^e - e^{r_t} \,. \tag{23}$$

Note that

$$\begin{aligned}
\psi_t^e &= \frac{E_t m_{t+1} E_t (C_{t+1}^\nu + p_{t+1}^e) - E_t m_{t+1}(C_{t+1}^\nu + p_{t+1}^e)}{p_t^e E_t m_{t+1}} \\
&= \frac{-\mathrm{Cov}_t(m_{t+1}, R_{t+1}^e)}{E_t m_{t+1}} \\
&= -\mathrm{Cov}_t\Big(\frac{m_{t+1}}{E_t m_{t+1}}, R_{t+1}^e\Big) \,, \tag{24}
\end{aligned}$$

where Cov_t denotes the covariance conditional on information at time t.[30]

The recursive equity pricing and equity premium equations (21)–(23) can be appended to the equations of the macroeconomic model in the previous section, allowing the equity premium (23) to be solved numerically along with the rest of the model. For the baseline calibration of the model reported in Table 1, solved to fifth order, the expected excess return on equity is about 1.05 percent per quarter (4.19 percent at an annualized rate), evaluating the model's state variables at their nonstochastic steady-state values. Empirical estimates of the equity premium typically range from about 3 to 6.5 percent for quarterly excess returns at an annual rate (e.g., Campbell, 1999, Fama and French, 2002), so the equity premium implied by the model is consistent with the data.

The model-implied equity premium is very sensitive to both the level of risk aversion R^c and the persistence of the technology shock ρ_A. Table 2 reports values for the equity premium for several different values of R^c and ρ_A, holding the other parameters of the model fixed at their baseline values from Table 1. The equity premium increases about linearly along with the household's coefficient of relative risk aversion, R^c, consistent with the analysis in Swanson (2013).[31] Perhaps more surprising is the substantial drop in the equity premium for values of

---

[30] If m_{t+1} and R_{t+1}^e are jointly lognormally-distributed, as is typically assumed in the finance literature, then the equation $E_t m_{t+1} R_{t+1}^e = 1$ implies $E_t r_{t+1}^e - r_t^f = -\mathrm{Cov}_t(\log m_{t+1}, r_{t+1}^e) - \frac{1}{2}\mathrm{Var}_t r_{t+1}^e$, where $r_{t+1}^e \equiv \log R_{t+1}^e$. Equation (24) says essentially the same thing without assuming joint lognormality.

[31] The equity premium increases linearly with risk aversion to second order around the nonstochastic steady state. The equity premium in Table 2 is computed to fifth order and thus is not strictly linear in risk aversion, but the intuition from the analysis in Swanson (2013) still holds.

TABLE 2: EQUITY PREMIUM AS A FUNCTION OF RISK AVERSION AND SHOCK PERSISTENCE

Risk aversion R^c	Shock persistence ρ_A	Equity premium ψ^e
10	1	0.62
30	1	1.96
60	1	4.19
90	1	6.70
60	.995	1.86
60	.99	1.08
60	.98	0.53
60	.95	0.17

Model-implied equity premium ψ^e, in annualized percentage points, for different values of relative risk aversion R^c and technology shock persistence ρ_A, holding the other parameters of the model fixed at their baseline values from Table 1. State variables of the model are evaluated at the nonstochastic steady state. See text for details.

ρ_A that are only slightly less than unity—for example, reducing ρ_A from 1 to .995 reduces the equity premium by more than half, and reducing ρ_A from .995 to .99 cuts the equity premium almost in half again. There are two reasons why ψ^e is so sensitive to ρ_A: First, equity is very long-lived, so it is sensitive to changes in the consumption dividend even at distant horizons. Second, the household's value function V_t, which enters into the stochastic discount factor (6), is also sensitive to consumption at long horizons. Reductions in ρ_A below unity have a very large effect on consumption at distant horizons, and thus significantly attenuate the response of both the equity price and the stochastic discount factor to a technology shock. The subsantially lower covariance between these two variables reduces the equity premium (equation (24)). (Note that the long-run risks literature, beginning with Bansal and Yaron, 2004, makes long-run consumption *more* volatile than my baseline calibration of $\rho_A = 1$; as a result, their models imply a larger equity premium than my model here, or a similar-sized equity premium with a lower degree of risk aversion.)

The equity premium in the model also varies substantially over time. Figure 2 plots the nonlinear (fifth-order) impulse response functions for the equity price p_t^e and the equity premium ψ_t^e to the technology shock, computed the same way as the nonlinear impulse response functions in Figure 1. The left-hand panel of Figure 2 depicts the response of the equity price, which jumps about 2.5 percent in response to the technology shock on impact. The risk-neutral increase in the equity price would be about 2.1 percent (the leverage ratio times the increase in dividends of about 0.7 percent every period); the additional 0.4 percent increase in the price is due to the decline in the risk premium that investors require to hold the risky asset. This can be seen in the
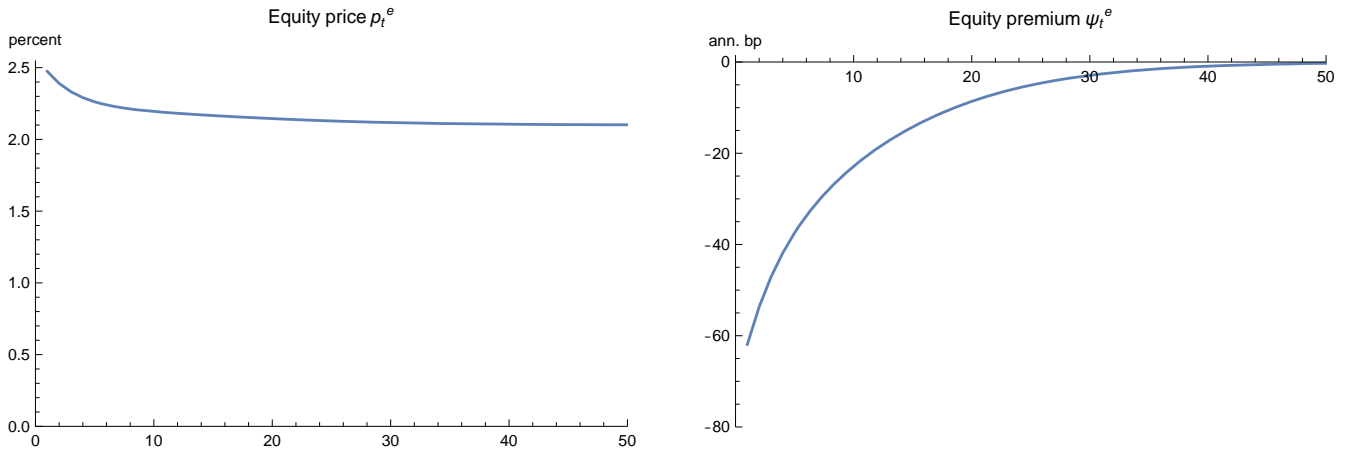
Figure 2. Nonlinear impulse response functions for the equity price p_t^e and equity premium ψ_t^e to a one-standard-deviation (0.7 percent) positive technology shock in the model, with state variables initialized to their nonstochastic steady state values. See text for details.

right-hand panel of Figure 2, where the equity premium drops about 60 basis points (bp) at an annual rate on impact before rising slowly back toward its initial level. Thus, the model produces an equity premium that is countercyclical (at least, in response to supply shocks), consistent with conventional wisdom in the literature (e.g., Fama and French, 1989; Campbell and Cochrane, 1999; Cooper and Priestley, 2008). The reason for this countercyclicality is that the volatility of the households' stochastic discount factor falls after a positive technology shock (and increases after a negative shock), for reasons I discuss below. Over the course of a year, the standard deviation of the equity premium in the model is about 103 bp, obtained by summing the squares of the first four quarters of the impulse response and taking the square root. (Note that this is the standard deviation of the *expected* excess return on equity; I discuss the standard deviation of the ex post excess return shortly.)

To compare the model-implied time variation in the equity premium to the data, it is useful to compute the model-implied Sharpe ratio, $\psi_t^e / \sqrt{\mathrm{Var}_t r_{t+1}^e}$, which is the standard measure used in the empirical finance literature. The average quarterly (non-annualized) Sharpe ratio in the model is $1.05/2.5 = 0.42$, which is about in line with the typical estimates of 0.2 to 0.4 in the literature (e.g., Campbell and Cochrane, 1999; Lettau and Ludvigson, 2010). The fact that the model's Sharpe ratio is at the high end of this range is not surprising since the model here is driven by a single shock and thus understates the overall volatility of equity prices; adding a monetary policy shock to the model, for example, would increase the volatility of equity without much altering its excess return (because monetary policy shocks are much less persistent than

technology shocks and have only a small effect on the equity premium), and lead to a lower Sharpe ratio more in line with the data.

The quarterly standard deviation of the (non-annualized) Sharpe ratio in the model is about $0.62/2.5 = 0.25$. Again, this is in line with estimates of the quarterly standard deviation of the Sharpe ratio in the literature, which range between 0.09 and 0.47 (e.g., Campbell and Cochrane, 1999; Lettau and Ludvigson, 2010, Table 11.7).

The *ex post* excess return on equity has a quarterly standard deviation of about 2.5 percent per quarter, or 5 percent per year. Empirical estimates in the literature are typically in the range of 6 to 12 percent per quarter, or 12 to 24 percent per year (e.g., Campbell, 1999; Lettau and Ludvigson, 2010), so the model-implied volatility for equity returns is substantially lower than the data. However, this is again not surprising, given that the stylized model here is driven by a single shock. Adding additional shocks to the model, such as fiscal or monetary policy shocks as in the New Keynesian DSGE literature (e.g., Smets and Wouters, 2007), would bring equity price volatility closer to the data.

From equation (24), we know that the decline in the model-implied equity premium in Figure 2 must be due to a drop in the conditional covariance of the equity price with the stochastic discount factor. In other words, the model generates *endogenous* conditional heteroskedasticity in response to shocks, even though the exogenous technology shock that drives the model is homoskedastic. This is a striking and very important feature of the model. I discuss how the model generates heteroskedasticity in detail in Section 4, below (and see also Section 2.4, above), but the key factor is the behavior of price dispersion Δ_t. In response to a technology shock, price dispersion moves in the same direction as the shock (see the bottom-right panels of Figures 1 and B1). Because greater price dispersion reduces output and aggregate productivity (see equation 17), the response of price dispersion to a technology shock tends to dampen the responses of the model's other variables to the shock. In addition, the sensitivity of price dispersion to a technology shock is greater if there has been a positive technology shock in the recent past—see Section 4, below. Thus, a positive technology shock today leads to a greater sensitivity of price dispersion Δ_t to future shocks, which reduces the volatility of the other variables of the model (such as consumption and the stochastic discount factor) to future shocks.

The result is that the stochastic discount factor displays substantial endogenous conditional heteroskedasticity. In a perfectly homogeneous, homoskedastic model—such as the ones typically used in finance that have no labor and no nominal rigidities—there is no endogenous conditional

heteroskedasticity at all. The only way to generate a time-varying equity premium in those models is to assume that the exogenous driving shock itself is conditionally heteroskedastic (see, e.g., Bansal and Yaron, 2004).

3.2 Real and Nominal Default-Free Bonds

A default-free zero-coupon real bond in the model pays one unit of consumption at maturity. Let $\tilde{p}_t^{(n)}$ denote the nominal price of an n-period zero-coupon real bond, and $p_t^{(n)} \equiv \tilde{p}_t^{(n)}/P_t$ its real price, with $p_t^{(0)} \equiv 1$. Then for $n \geq 1$,

$$p_t^{(n)} = E_t \, m_{t+1} p_{t+1}^{(n-1)} \tag{25}$$

in equilibrium in each period t. In particular, $p_t^{(1)} = e^{-r_t}$.

A default-free zero-coupon nominal bond pays one nominal dollar at maturity. Let $p_t^{\$(n)}$ denote the nominal price of an n-period zero-coupon nominal bond, with $p_t^{\$(0)} \equiv 1$. Then for $n \geq 1$,

$$p_t^{\$(n)} = E_t \, m_{t+1} e^{-\pi_{t+1}} p_{t+1}^{\$(n-1)} \tag{26}$$

in each period t. In particular, $p_t^{\$(1)} = e^{-i_t}$.

Let $r_t^{(n)}$ denote the n-period continuously-compounded yield to maturity on a real zero-coupon bond, and $i_t^{(n)}$ the corresponding yield on an n-period nominal bond. Then

$$r_t^{(n)} = -\frac{1}{n} \log p_t^{(n)} \tag{27}$$

and

$$i_t^{(n)} = -\frac{1}{n} \log p_t^{\$(n)}. \tag{28}$$

Note that even though these bonds are free from default, they are risky in the sense that their prices can fluctuate in response to shocks, for $n > 1$.

The risk premium on a bond is typically written as a *term premium,* the difference between the yield to maturity on the bond and the hypothetical, risk-neutral yield to maturity on the same bond. The risk-neutral real price $\hat{p}_t^{(n)}$ of an n-period zero-coupon real bond is given by

$$\hat{p}_t^{(n)} = e^{-r_t} E_t \, \hat{p}_{t+1}^{(n-1)}, \tag{29}$$

where $\hat{p}_t^{(0)} \equiv 1$. The n-period real term premium $\psi_t^{(n)}$ is then given by

$$\psi_t^{(n)} \equiv \frac{1}{n} \left(\log \hat{p}_t^{(n)} - \log p_t^{(n)} \right)$$

$$\approx \frac{1}{n \overline{p}^{(n)}} \left(\hat{p}_t^{(n)} - p_t^{(n)} \right), \tag{30}$$

where $\overline{p}^{(n)}$ denotes the steady-state real bond price.[32] The formula for the term premium on a nominal n-period bond, $\psi_t^{\$(n)}$, is analogous. Note that

$$
\begin{aligned}
\hat{p}_t^{(n)} - p_t^{(n)} &= E_t m_{t+1} E_t \hat{p}_{t+1}^{(n-1)} - E_t m_{t+1} p_{t+1}^{(n-1)} \\
&= -\mathrm{Cov}_t\big(m_{t+1}, p_{t+1}^{(n-1)}\big) + e^{-r_t} E_t\big(\hat{p}_{t+1}^{(n-1)} - p_{t+1}^{(n-1)}\big) \\
&= -E_t \sum_{j=0}^{n-1} e^{-r_{t,t+j}} \mathrm{Cov}_{t+j}\big(m_{t+j+1}, p_{t+j+1}^{(n-j-1)}\big),
\end{aligned}
\tag{31}
$$

where $r_{t,t+j} \equiv \sum_{\tau=t+1}^{t+j} r_\tau$ and the last line of (31) follows from forward recursion. Equation (31) shows that, even though the bond price depends only on the one-period-ahead covariance between the stochastic discount factor and next period's bond price, the risk premium on the bond depends on this covariance over the entire lifetime of the bond. Substituting (31) into (30) gives

$$
\psi_t^{(n)} = -\frac{1}{n\overline{p}^{(n)}} E_t \sum_{j=0}^{n-1} e^{-r_{t,t+j}} \mathrm{Cov}_{t+j}\big(m_{t+j+1}, p_{t+j+1}^{(n-j-1)}\big).
\tag{32}
$$

Intuitively, the term premium is larger the more negative the covariance between the stochastic discount factor and the price of the bond over the lifetime of the bond.

The bond pricing and bond yield equations (25)–(29) are recursive and can be appended to the macroeconomic model described above and solved numerically along with the macroeconomic variables, equity price, and equity premium. (Note that, to consider a bond with n periods to maturity, $n-1$ bond pricing equations must be appended to the model, one for each maturity from 2 to n.)

Table 3 reports the real yield curve implied by the model, along with the corresponding average real yields estimated from inflation-indexed government bonds in the U.S. and U.K. over different sample periods. Data for U.S. inflation-indexed Treasuries (TIPS) are taken from the updated Gürkaynak, Sack, and Wright (2010) online dataset. The first TIPS were issued in 1998, and a yield curve for maturities of 5 years or more can be estimated beginning in 1999. The first row of Table 3 thus reports average TIPS yields from 1999 to 2015. Real yields over this sample averaged about 1.3 to 1.8 percent per year. Zero-coupon yields for shorter-maturity TIPS (down to 2 years; neither Gürkaynak et al., 2010, nor the Bank of England report zero-coupon real yields with a maturity less than 2 years) can be estimated beginning in 2004, and are reported in the

---

[32] The first-order approximation on the first line of (30) is useful for gaining intuition. However, when I solve for bond prices and risk premia in the model numerically below, the solution will always include second-, third-, and higher-order terms as well as first-order terms.

TABLE 3: REAL ZERO-COUPON BOND YIELDS, DATA VS. MODEL

	2-yr.	3-yr.	5-yr.	7-yr.	10-yr.	(10y)−(3y)
US TIPS, 1999–2015[a]			1.29	1.55	1.82	
US TIPS, 2004–2015[a]	0.14	0.28	0.60	0.89	1.22	0.94
US TIPS, 2004–2007[a]	1.39	1.52	1.74	1.91	2.09	0.57
UK indexed gilts, 1983–1995[b]	6.12	5.29	4.34		4.12	−1.17
UK indexed gilts, 1985–2015[c]		1.91	2.05	2.16	2.25	0.34
UK indexed gilts, 1990–2007[c]		2.79	2.78	2.79	2.80	0.01
macroeconomic model	1.94	1.93	1.93	1.93	1.93	0.00

[a]Gürkaynak, Sack, and Wright (2010) online dataset.
[b]Evans (1999).
[c]Bank of England web site.
Estimated zero-coupon real yields from inflation-indexed bonds in the U.S. and U.K., and zero-coupon real yields implied by the macroeconomic model presented above. The last column reports the difference between the 10-year and 3-year yields in each row. See text for details.

second row of Table 3, along with the average yields for longer maturities over the same sample. This sample also excludes the period of lower TIPS liquidity in the first few years after they were issued. Over this sample, average real yields are lower, between about 0.15 and 1.2 percent. However, the period from 2008–15 is unusual in that the financial crisis and severe recession led the Federal Reserve to reduce short-term interest rates to record lows, and to some extent we might expect this to show up in shorter-term real yields as well, both as a lower level of yields and as a steeper yield curve slope. Thus, the third row of Table 3 reports results from 2004–07, a short sample, but one that avoids both the low liquidity of TIPS in its first few years and the financial crisis and recession. Over this sample, real yields average between about 1.4 and 2.1 percent.

However, this is a short sample and the period from 2004 to 2005 was also characterized by very easy monetary policy and a very low level of short-term U.S. yields as the Federal Reserve worked to facilitate recovery from the 2001 recession. Thus, the next three rows of Table 3 report average real yields on inflation-indexed gilts in the U.K., for which we have a longer sample (indexed gilts have traded since at least the early 1980s in the U.K.). Evans (1999) estimates real zero-coupon U.K. yields from 1983 to 1995, reported in the fourth row of Table 3, which average between about 4 and 6 percent over that sample. Interestingly, the real U.K. gilt yield curve slopes *downward* rather than upward over this period, by about 100–200 basis points. However, as in the U.S., the early years of the U.K. indexed gilt market may have suffered from low liquidity and correspondingly higher yields. Thus, the fifth row of Table 3 reports estimated real yields

from 1985 to the present, from the Bank of England's web site. Over this longer sample, real U.K. yields average about 1.9 to 2.25 percent, and the yield curve sloped upward by about 34 bp. The sixth row of Table 3 reports results for the U.K. excluding both the early years of the sample and the financial crisis and recession period, for the same reasons as for the U.S. Over this sample, 1990–2007, real yields in the U.K. are a bit higher, averaging about 2.8 percent, and the yield curve is about flat, sloping upward by 1 bp.

While the exact level and slope of the real yield curve depend on the sample period and country considered (U.S. vs. U.K.), the overall pattern suggests an average real yield of approximately 2 percent per year, with a slope that is relatively flat—neither strongly upward-sloping nor downward-sloping on average. The macroeconomic model presented above fits these features of real yields in the data quite well. Real yields in the model average a bit less than 2 percent under the baseline calibration, evaluating the model's state variables at the nonstochastic steady state.[33] The model also implies that the real yield curve is about flat on average, with essentially no spread between the 10-year and 3-year yields, and a -1 bp spread between the 10-year and 2-year yields.

A downward-sloping real yield curve is a standard feature of traditional real-business-cycle models—see Backus, Gregory, and Zin (1989), Donaldson, Johnsen, and Mehra (1990) and Den Haan (1995). Intuitively, if short-term real interest rates fall in recessions, then the price of a long-term real bond will tend to rise in recessions, which is when households value consumption the most. Thus, long-term real bonds act like recession insurance and should carry a negative risk premium. In the macroeconomic model I develop here, the response of the short-term real interest rate to the shock is fairly small (see Figure 1), implying a relatively small change in the real long-term bond price (see Figure 3, below). Moreover, the fall in the real term premium after the shock (Figure 3) attenuates the response of the real bond price even more. As a result, the price of a real long-term bond is not very countercyclical and the insurance properties of the bond are minor, resulting in only a small risk premium, consistent with the data.

Figure 3 reports nonlinear impulse response functions for the 10-year real bond price and term premium, computed in the same way as in Figures 1 and 2. The bond price falls only about 0.2 percent on impact, due to the small increase in short-term real rates in Figure 1, and the

---

[33] Note that the baseline value of β from Table 1 implies a real yield of a little more than 3 percent in the nonstochastic steady state. However, the real yield $r_t = 1/E_t m_{t+1}$, and $E_t m_{t+1}$ is substantially greater than $1/\beta$ in the stochastic case due to Jensen's inequality terms. Intuitively, households' aversion to risk drives up their demand for precautionary savings in the riskless asset, lowering the risk-free rate below its nonstochastic steady-state value.
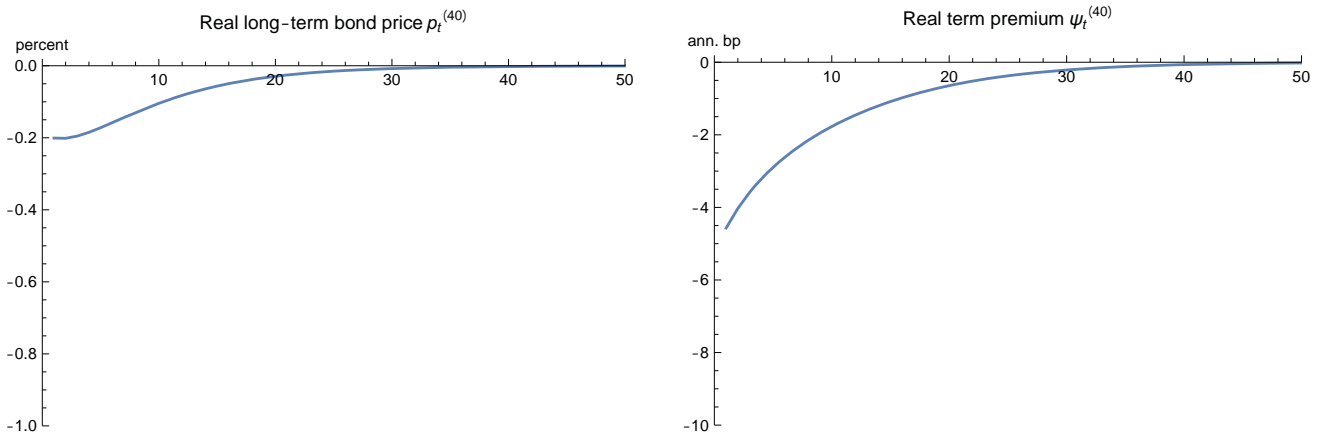
Figure 3. Nonlinear impulse response functions for real long-term bond price $p_t^{(40)}$ and real term premium $\psi_t^{(40)}$ to a one-standard-deviation (0.7 percent) positive technology shock in the model, with state variables initialized to their nonstochastic steady state values. See text for details.

offsetting fall in the real term premium (see the right-hand panel of Figure 3). The real term premium declines in response to the shock, but by much less than the equity premium, only about 4 bp.

Table 4 compares the nominal yield curves implied by the model to the data. Gürkaynak, Sack, and Wright (2007) estimate zero-coupon nominal Treasury yields for the U.S. going back to 1961 for maturities out to 7 years, and 1971 for maturities out to 10 years (data through the present are available from the online version of their dataset). From 1961 to 2015, nominal yields averaged about 5.3 to 6.2 percent. From 1971 to 2015, the average is a bit higher, about 5.4 to 6.7 percent, with an average yield curve slope of about 130 bp. Just as for real yields, though, the period from 2008–15 may be atypical in that short-term interest rates hit record lows in response to the financial crisis and recession, and were constrained by the zero lower bound on nominal rates. The "Great Inflation" period of the 1970s and early 1980s may also be problematic in that monetary policy may have experienced a structural break since that period and is now conducted in a more aggressive anti-inflationary manner (e.g., Clarida, Galí, and Gertler, 1999). Thus, the third row of Table 4 reports average yields from 1990 to 2007, a period that excludes both the Great Inflation and recent Great Recession periods. Over this sample, nominal Treasury yields averaged about 4.5 to 6 percent, with a yield curve slope of about 140 bp.

The Bank of England also reports estimated zero-coupon yield curves for the U.K. going back to 1970. From 1970 to 2015, nominal gilt yields in the U.K. averaged between about 6.9 and 7.9 percent, with a yield curve slope of about 95 bp, as reported in Table 4. Restricting attention

Table 4: Nominal Zero-Coupon Bond Yields, Data vs. Model

	1-yr.	2-yr.	3-yr.	5-yr.	7-yr.	10-yr.	(10y)−(1y)
US Treasuries, 1961–2015[a]	5.27	5.50	5.68	5.97	6.18		
US Treasuries, 1971–2015[a]	5.42	5.66	5.86	6.18	6.43	6.71	1.29
US Treasuries, 1990–2007[a]	4.56	4.84	5.06	5.41	5.68	5.98	1.42
UK gilts, 1970–2015[b]	6.92	7.10	7.26	7.51	7.70	7.89	0.96
UK gilts, 1990–2007[b]	6.20	6.29	6.38	6.47	6.50	6.48	0.28
macroeconomic model	5.35	5.59	5.80	6.09	6.27	6.44	1.09

[a]Gürkaynak, Sack, and Wright (2007) online dataset.
[b]Bank of England web site.
Empirical estimates of zero-coupon nominal yields from government bonds in the U.S. and U.K., and zero-coupon nominal yields implied by the macroeconomic model presented above. The last column reports the difference between the 10-year and 1-year yield in each row. See text for details.

to the period from 1990 to 2007, for the same reasons as above, average U.K. nominal yields are a bit lower, about 6.2 to 6.5 percent, with a slope of just 28 bp.

Again, the exact level and slope of the nominal yield curve depends on the sample period and country considered, but nominal yields appear to average about 5 or 6 percent and have an upward slope of about 100 bp. The model is able to reproduce these features of the data quite well: the average level of nominal yields in the model (evaluating the state variables at the nonstochastic steady state) is between about 5.4 and 6.4 percent, with an upward slope of 109 bp. Thus, although the model-implied real yield curve is flat, the implied nominal yield curve slopes upward substantially. As discussed by Rudebusch and Swanson (2012), this is because technology shocks in the model make nominal bonds risky: A negative technology shock causes inflation to rise persistently at the same time that consumption falls; as a result, long-term nominal bonds in the model *lose* value in recessions. This implies that long-term nominal bonds should carry a substantial risk premium, about 110 bp over the corresponding risk-neutral yield. Thus, the simple model presented here provides a straightforward answer to the puzzle posed by Backus, Gregory, and Zin (1989), Donaldson, Johnsen, and Mehra (1990), and Den Haan (1995): namely, why does the nominal yield curve slope upward? The answer is technology shocks, or more generally, any type of "supply shock" that causes inflation and output to move in opposite directions, such as an oil price shock or markup shock.

Of course, the larger and more important are technology or supply shocks in the model, the larger the term premium on nominal bonds. Thus, if supply shocks were relatively larger in the 1970s and early 1980s than in the 1960s or more recently, the model predicts that we should
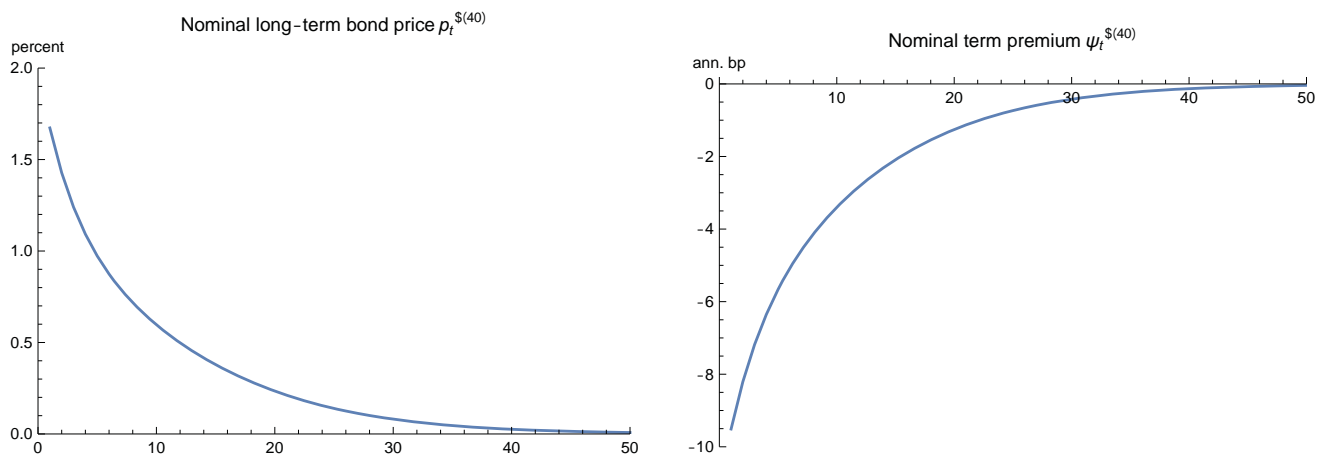
Figure 4. Nonlinear impulse response functions for nominal long-term bond price $p_t^{\$(40)}$ and term premium $\psi_t^{\$(40)}$ to a one-standard-deviation (0.7 percent) technology shock in the model, with state variables initialized to their nonstochastic steady state values. See text for details.

see a larger term premium on nominal bonds in those periods when supply shocks were larger. And in fact, this prediction seems to be consistent with the data: Rudebusch, Sack, and Swanson (2007) graph several measures of the term premium—from a VAR, affine no-arbitrage models with latent or observable factors, and the Cochrane-Piazzesi (2005) "tent-shaped" predictor of excess returns—and for all of these measures, the estimated term premium on long-term nominal bonds in the U.S. is higher in the 1970s and early 1980s than in the 1960s or more recently.

Campbell, Sundaram, and Viceira (2013) also document changes in the correlation between stock and nominal bond returns over time. Although the baseline calibration of the model here has only a single shock, making it stochastically singular, extending the model to include fiscal and/or monetary policy shocks is straightforward and is standard in the medium-scale New Keynesian DSGE literature (e.g., Smets and Wouters, 2007). In these models, if the relative importance of technology or supply shocks is varied, then the size of the term premium and the correlation between stock and bond returns will vary as well. Thus, changing correlations between stock and bond returns can be mapped back to more structural features of the model.

Figure 4 plots the nonlinear impulse response functions for the 10-year nominal bond price and term premium to a one-standard-deviation positive technology shock, computed in the same way as in Figures 1–3. As discussed above, a positive technology shock causes inflation and the short-term nominal interest rate to fall (Figure 1) and the nominal long-term bond price to rise substantially (Figure 4), about 1.7 percent on impact before gradually returning back to steady state. The nominal term premium falls about 10 bp on impact, so part of the strong

price response of the long-term nominal bond is due to the fall in the term premium. The reason for that fall is essentially the same as for the equity premium: the decline in the volatility of the households' stochastic discount factor after the positive technology shock. Importantly, the model's prediction of a countercyclical term premium is consistent with the evidence in Fama and French (1989), Cooper and Priestley (2008), Piazzesi and Swanson (2008), and conventional wisdom in the literature (e.g., Campbell and Cochrane, 1999). Over the course of a year, the standard deviation of the term premium is about 16 bp.

Estimates of the quarterly standard deviation of the term premium in the data range between about 8 to 40 bp: standard affine term structure models with latent factors, such as Kim and Wright (2005), imply a quarterly standard deviation of about 30–35 bp, but Rudebusch and Wu (2007) argue that these highly-parameterized models tend to overfit the high-frequency fluctuations in long-term yields, and that fluctuations in the term premium are smaller, only about 8 bp from quarter to quarter (see also the survey of empirical estimates in Rudebusch, Sack, and Swanson, 2007). The term premium implied by the model of the present paper is consistent with this range of estimates, but lies toward the lower end, consistent with the less highly-parameterized models of Rudebusch and Wu (2007) and others.

3.3 Defaultable Bonds

The simple macroeconomic model developed above also matches many features of defaultable bond prices. For simplicity, I follow the theoretical finance literature and model a defaultable bond as a depreciating consol that has some probability of defaulting each period (see, e.g., Leland, 1994, 1998; Duffie and Lando, 2001; and Chen, 2010). The *credit spread* in the model is the difference in yield between the defaultable consol and an otherwise identical consol that is free from default. I consider two cases in the analysis below: first, where the probability of default and the recovery rate given default are constant over time, and second, where those quantities vary cyclically in line with the data.

A default-free depreciating nominal consol is an infinitely-lived bond that pays a geometrically declining coupon of δ^n dollars in each period $n = 1, 2, \ldots$ after issuance. The equilibrium ex-coupon price p_t^c of the consol in period t is given by

$$p_t^c = E_t\, m_{t+1} e^{-\pi_{t+1}} (1 + \delta p_{t+1}^c), \tag{33}$$

where the size of the next coupon payment is normalized to one dollar. The very simple recursive structure of (33) makes this type of long-term bond extremely convenient to work with and

generalizes naturally to the case where the bond may default.[34] When $\delta = 0$, the consol reduces to a one-period zero-coupon bond, and when $\delta = 1$, it behaves like a traditional nondepreciating consol. By choosing δ appropriately, the consol can be given any desired Macauley duration and made to behave very similarly to the corresponding zero-coupon bond.

The continuously-compounded yield to maturity, i_t^c, for the consol satisfies

$$p_t^c = \frac{1}{e^{i_t^c}} + \frac{\delta}{e^{2i_t^c}} + \frac{\delta^2}{e^{3i_t^c}} + \cdots , \tag{34}$$

implying

$$i_t^c = \log\left(\frac{1}{p_t^c} + \delta\right). \tag{35}$$

The Macauley duration of the consol is

$$-\frac{d\log p_t^c}{di_t^c} = 1 + \delta p_t^c. \tag{36}$$

When calibrating the model below, I set δ so that the consol has a Macauley duration of 10 years, corresponding to the approximate duration of the longer-term coupon bonds in Moody's indexes.

A defaultable consol pays a nominal coupon each period in the same way as a default-free consol, but in addition there is a chance each period that the bond will default and cease paying interest forever. In the event of default, bondholders receive a recovery rate times the previous value of the bond, which we can calibrate to the data. Thus, the defaultable consol price p_t^d satisfies

$$p_t^d = E_t\, m_{t+1} e^{-\pi_{t+1}}\Big[(1 - \mathbf{1}_{t+1}^d)(1 + \delta p_{t+1}^d) + \mathbf{1}_{t+1}^d\, \omega_{t+1}\, p_t^d\Big], \tag{37}$$

where $\mathbf{1}_t^d$ is an indicator variable equal to 1 if the bond defaults in period t and 0 otherwise, and ω_t denotes the recovery rate on the bond in the event of default. The yield to maturity i_t^d and duration of the defaultable bond are defined by equations (35)–(36), with p_t^d in place of p_t^c. The credit spread is the yield differential, $i_t^d - i_t^c$.

It remains to calibrate $\Pr_t\{\mathbf{1}_{t+1}^d = 1\}$ and ω_t in (37). The average rate of default for bonds initially rated Baa or BBB is about 0.6 percent per year (e.g., Moody's, 2006; Standard & Poor's, 2014), and the average recovery rate on defaulted bonds is about 42 percent (Chen, Collin-

---

[34] Leland (1994), Duffie and Lando (2001), and Chen (2010) use a nondepreciating consol to model corporate bonds, while Leland (1998) uses a depreciating consol. Rudebusch and Swanson (2008) use a (default-free) depreciating consol to study the long-term bond premium puzzle. The behavior of the depreciating consol in the simple model above and in Rudebusch and Swanson (2008) is very similar to that of a zero-coupon bond with the same Macauley duration.

TABLE 5: MODEL-IMPLIED CREDIT SPREAD ON DEFAULTABLE BONDS

average ann. default prob.	cyclicality of default prob.	average recovery rate	cyclicality of recovery rate	credit spread (bp)
.006	0	.42	0	34.0
.006	−0.3	.42	0	130.9
.006	−0.3	.42	2.5	143.1
.006	−0.15	.42	2.5	78.9
.006	−0.6	.42	2.5	367.4
.006	−0.3	.42	1.25	137.0
.006	−0.3	.42	5	155.2

Model-implied credit spread $i_t^d - i_t^c$ for defaultable vs. default-free depreciating consols with Macauley duration of 10 years. Average annualized default probability is calibrated to bonds initially rated Baa. Cyclicality of default probability and recovery rate are the loadings on the output gap, $y_t - \overline{y}_t$. See text for details.

Dufresne, and Goldstein, 2009; Chen, 2010).[35] As a first calibration, then, I set $\Pr_t\{\mathbf{1}_{t+1}^d = 1\}$ to an exogenous, constant rate of 0.15 percent per quarter, and ω_t to an exogenous constant of 42 percent.

The credit spread implied by the model for this calibration is reported in the first row of Table 5. With a constant average annual default probability of 0.6 percent, the model-implied credit spread is about 34 bp. This is essentially the risk-neutral expected loss each period from default, $(.006)(.58) = 34.8$ bp, and is far less than the historical average credit spread on Baa-rated bonds of about 120 bp (e.g., Chen, Collin-Dufresne, and Goldstein, 2009; Chen, 2010).[36] Intuitively, if the risk of default in the model is uncorrelated with the stochastic discount factor, there is no additional risk premium attached to losses from default.

Empirically, however, corporate bond defaults are highly countercyclical and recovery rates highly procyclical (see, e.g., Chen, 2010; Giesecke, Longstaff, Schaefer, and Strebulaev, 2011; Standard & Poor's, 2011). For example, in Figure 1 of Chen (2010), the default rate averages about 0.9 percent for all bonds over the postwar period, but spikes to about 3.7 percent in 1990, 4 percent in 2001, and 5.5 percent in 2009, with smaller increases in earlier recessions (and a spike

---

[35] The default rate on bonds *currently* rated Baa/BBB is much lower, about 0.15 percent per year on average. However, these bonds also lose value when they are downgraded, which happens with much higher probability than default. Rather than keep track of credit ratings, the probability of downgrades, and capital losses in the event of downgrade, I simply keep track of the default rate for bonds *initially* rated Baa/BBB.

[36] This is the average difference between the yield on Moody's Baa and Aaa seasoned corporate bond indexes from 1921–2013. The average spread over alternative sample periods is similar. The spread between Baa-rated corporate bonds and U.S. Treasuries is even larger, about 185 bp. However, U.S. Treasuries carry an additional premium for their extreme liquidity and beneficial tax treatment, so the Baa-Aaa spread is often used in the literature to measure the credit spread (Aaa corporate bonds are similar in liquidity and tax treatment to Baa-rated bonds and the probability of default on Aaa-rated bonds is still extremely low; see, e.g., Chen et al., 2009).

to 8.5 percent in 1933). In boom years, the default rate falls to essentially zero. Recovery rates average about 42 percent, but drop to about 20–25 percent in 1990, 2001, and 2009, while they rise to 50–60 percent in boom years.

Thus, the next rows of Table 5 consider cases where the default rate, recovery rate, or both are correlated with the output gap in the model, $y_t - \overline{y}_t$. I calibrate the cyclicality of the model's annualized default rate to a value of -0.3, which implies a drop in output of 5 percent below trend is associated with an increase in the default rate of about 1.5 percentage points. While this cyclicality is lower than in Chen (2010), my focus here is on bonds initially rated Baa/BBB, while the data in Chen (2010) is for all bonds, which includes many that were issued at ratings below investment grade.[37] The second row of Table 5 reports the model-implied credit spread when the default rate is countercyclical, holding the recovery rate constant over time. This greatly increases the model-implied credit spread, to about 131 bp, consistent with the observed spread in the data.

The third row of Table 5 considers the case where the recovery rate is also cyclical. I calibrate the cyclicality of the recovery rate in the model to 2.5, so that a fall in output of 5 percent below trend is associated with a roughly 12.5-percentage-point decrease in the recovery rate on defaulted corporate bonds, in line with the variation reported in Chen (2010). Given this degree of cyclicality, the credit spread in the model increases a bit further, to 143 bp, still close to (and even a bit above) the value of 120 bp in the data.

In the last four rows of Table 5, I vary these cyclicality parameters to check how sensitive the credit spread is to their variation. Changes in the cyclicality of default have a large effect on the spread, while changes in the cyclicality of the recovery rate have a much smaller effect, moving the spread by only a few basis points. Intuitively, a marginal increase in the probability of default is much more costly to households because it implies an increase in the chance of a large loss; in contrast, a marginal fall in the recovery rate implies only a small chance (0.15 percent per quarter) of a modest increase in the loss. Thus, the cyclicality of recovery rates is much less important in the model and can largely be ignored.

Figure 5 reports nonlinear impulse response functions for the defaultable bond price and credit spread to a positive one-standard-deviation technology shock, computed the same way as in previous figures. The default probability and recovery rate in the model are assumed to

---

[37] See the discussion in footnote 35. Also, to prevent the default rate in the model from becoming negative, I model it in logarithms rather than in levels. That is, the cyclicality of the log default rate is set to -50, which, when multiplied by the average default rate of .006 per year, produces -0.3.
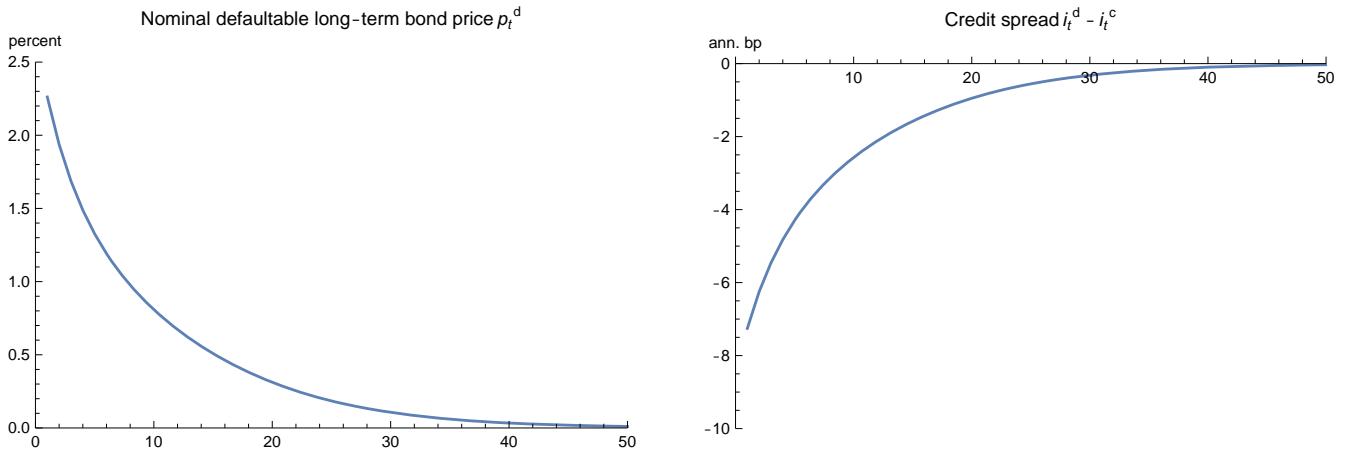
Figure 5. Nonlinear impulse response functions for defaultable long-term bond price p_t^d and credit spread $i_t^d - i_t^c$ to a one-standard-deviation (0.7 percent) technology shock in the model, with state variables initialized to their nonstochastic steady state values. See text for details.

have the same cyclicality as in the third row of Table 5, consistent with the data. On impact, the defaultable bond price jumps about 2.3 percent, in between the responses of the default-free nominal bond and equity prices in Figures 4 and 2. This is intuitive, since defaultable bonds are riskier than default-free bonds but less risky than equity in the model.

The credit spread, depicted in the right-hand panel of Figure 5, drops about 7 bp on impact. This is somewhat less than in the data; the standard deviation of the post-war quarterly change in the Baa-Aaa spread is about 20 bp. However, as discussed above, the stylized model here has only one driving shock; extending the model to include additional shocks would increase the overall volatility of the credit spread and bring it closer to the data.

To some extent, the model's ability to jointly fit equity returns and corporate bond yields is not surprising, since Chen, Collin-Dufresne, and Goldstein (2009), Bhamra, Kuehn, and Strebulaev (2010), and Chen (2010) achieve similar simultaneous fits in an endowment economy. Nevertheless, the present paper is the first to jointly match these data in a fully-specified macroeconomic model. The distinction is important because results in an endowment economy do not necessarily carry over to the case where households can choose their consumption stream endogenously (see, e.g., the discussion of Campbell-Cochrane (1999) habits in the Introduction above and in Rudebusch and Swanson, 2008). Like the present paper, Bhamra et al. (2010) and Chen (2010) use Epstein-Zin preferences, albeit with consumption and inflation taken to be exogenous, reduced-form processes. The advantage of the structural macroeconomic approach I take here is its greater robustness to structural breaks and ability to consider novel policy interventions,

which cannot be studied in a reduced-form macroeconomic environment. On the other hand, the much simpler macroeconomic structure in Bhamra et al. (2010) and Chen (2010) allows them to perform a more structural analysis of firms' corporate financing and default decisions. In other words, I have adopted a simplistic, reduced-form model of the firm in order to better focus on the structural behavior of the macroeconomy, while Bhamra et al. (2010) and Chen (2010) have adopted a very simplistic, reduced-form model of the macroeconomy to better focus on the structural finance behavior of the firm.

4. Endogenous Conditional Heteroskedasticity

The macroeconomic model developed above naturally generates endogenous conditional heteroskedasticity in response to shocks, as mentioned previously in Sections 2.4 and 3.1. Indeed, this feature of the model is crucial for generating time-varying risk premia on assets, as discussed in Section 2.4: If the stochastic discount factor and the asset return are both homoskedastic, then the risk premium on the asset must be constant over time, as can be seen in equation (24). The fact that the model produces time-varying risk premia in Figures 2 through 5 is therefore (indirect) evidence that the model-implied stochastic discount factor is heteroskedastic.

In Figure 6, I provide direct evidence of this heteroskedasticity. Given that $m_{t+1} = \beta(C_{t+1}/C_t)^{-1}\big[\exp(-\alpha V_{t+1})/E_t\exp(-\alpha V_{t+1})\big]$, the conditional volatility of the stochastic discount factor, $\mathrm{Var}_t\log m_{t+1}$, can be decomposed into two parts: the variance of consumption growth, $\mathrm{Var}_t\log(C_{t+1}/C_t)^{-1}$, and the variance of the Epstein-Zin term, $\mathrm{Var}_t\log\big[\exp(-\alpha V_{t+1})/E_t\exp(-\alpha V_{t+1})\big]$. I report nonlinear impulse response functions for each of these two components in the left- and right-hand panels of Figure 6.[38]

In response to a positive, one-standard-deviation technology shock, the conditional variance of the SDF falls about 50 percent in total. The conditional variance of consumption growth in the left-hand panel of Figure 6 falls about 40 percent, while the conditional variance of the Epstein-Zin term (the right-hand panel) falls about 10 percent, so the fall in the variance of current-period

---

[38] The nonlinear impulse response functions are computed in the same way as in previous figures. I compute the conditional variance of a variable X_{t+1} in the model by defining $\mu_t^X \equiv E_t X_{t+1}$ and then $V_t^X \equiv E_t(X_{t+1} - \mu_t^X)^2$. I append these recursive equations to the rest of the model and solve them nonlinearly along with the other model variables as described earlier. Note that the conditional variance V_t^X is linearized and not log-linearized around the nonstochastic steady state because it equals zero at that point (when the variance of the technology shock σ_A^2 is set to zero). I report the change in variances V_t^X in Figure 6 in percentage terms by dividing the impulse responses (in levels) through by a constant; namely, the variance V_t^X solved to fifth order and evaluated with each state variable at the nonstochastic steady state, but with the variance of the technology shock σ_A^2 set to $.007^2$.
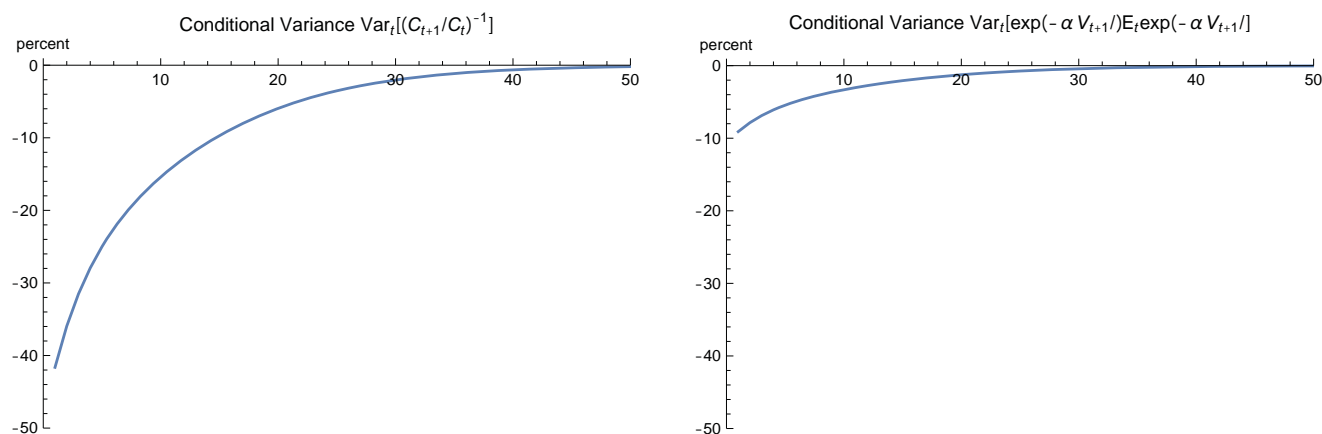
Figure 6. Nonlinear impulse response functions for conditional variances $\mathrm{Var}_t(C_{t+1}/C_t)^{-1}$ and $\mathrm{Var}_t\big[\exp(-\alpha V_{t+1})/E_t\exp(-\alpha V_{t+1})\big]$ to a one-standard-deviation (0.7 percent) positive technology shock in the model, with state variables initialized to their nonstochastic steady-state values. Impulse responses are in percentage deviation from steady state. The model-implied stochastic discount factor is $m_{t+1} = \beta(C_{t+1}/C_t)^{-1}\big[\exp(-\alpha V_{t+1})/E_t\exp(-\alpha V_{t+1})\big]$, so the two panels decompose the response of the conditional variance of the SDF to the shock. See text for details.

consumption growth is the primary driver of the decline in volatility of the SDF.[39] Nevertheless, the Epstein-Zin term is still important because its average volatility is so much higher than that of consumption growth—about 0.47^2 vs. 0.0065^2. The high average volatility of the Epstein-Zin term is what makes the relative fall in consumption growth volatility quantitatively important for risk premia in the model.

In Figure 7, I investigate what drives the decline in consumption volatility in Figure 6. The left column reports nonlinear impulse response functions for a one-standard-deviation (0.7 percent) positive technology shock, while the right column reports the analogous impulse response functions for a *negative* (-0.7 percent) technology shock. In each panel of Figure 7, two lines are plotted: The dashed blue line is the standard nonlinear impulse response function computed in the same way as in previous figures (see footnote 23); that is, the period-by-period difference between a "one shock" and a "no shock" (baseline) scenario, with all state variables of the model initialized to their nonstochastic steady-state values. The solid green line in each panel of Figure 7 is the period-by-period difference starting from a *different* initial point: instead of the nonstochastic steady state, the impulse responses are computed starting from the point right after a positive 0.7 percent technology shock in the previous period. Thus, the solid green lines in Figure 7 depict

[39] As with previous figures, the impulse responses in Figure 6 are essentially symmetric for a negative technology shock. That is, a negative technology shock causes the conditional volatility of the SDF to increase, with a similar magnitude to (in fact, slightly larger magnitude than) Figure 6.
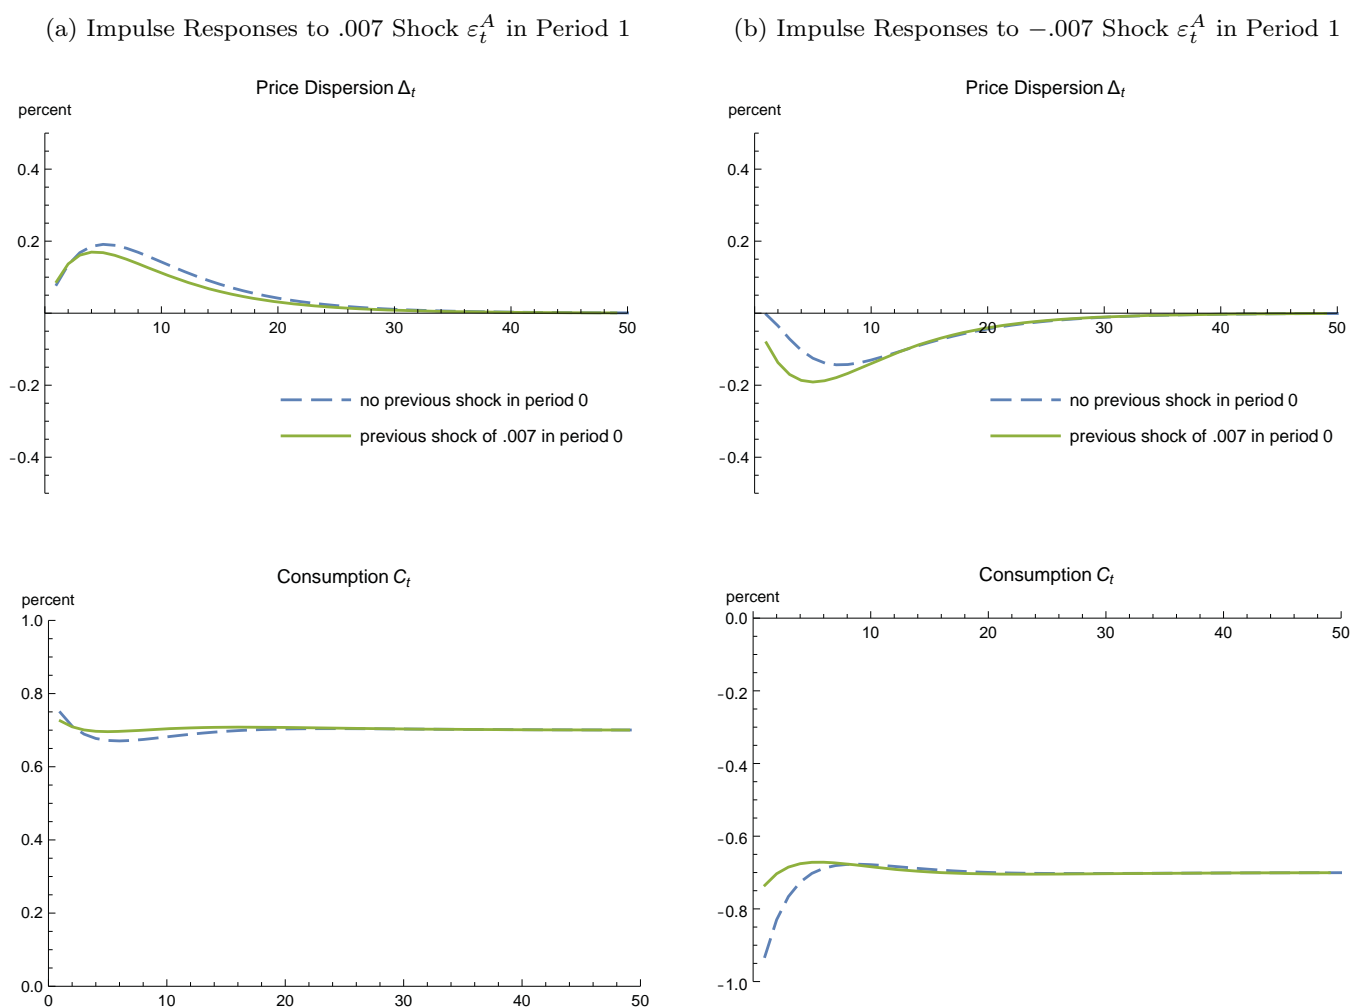
(a) Impulse Responses to .007 Shock ε_t^A in Period 1   (b) Impulse Responses to $-.007$ Shock ε_t^A in Period 1

Price Dispersion Δ_t

Price Dispersion Δ_t

Consumption C_t

Consumption C_t

Figure 7. Comparison of nonlinear impulse response functions for price dispersion Δ_t and consumption C_t to a one-standard-deviation (0.7 percent) (a) positive vs. (b) negative technology shock in period 1. Dashed blue lines depict impulse response functions relative to a baseline of no previous shocks, with state variables initialized to their nonstochastic steady-state values, as in previous figures. Solid green lines depict impulse response functions relative to a different baseline, after a positive 0.7 percent technology shock in period 0. The figure shows that the conditional volatility of consumption is lower after a positive 0.7 percent technology shock in period 0, particularly in response to negative shocks in period 1. See text for details.

the difference between a "two shock" and a "one shock" (baseline) scenario.[40]

The lower conditional volatility of consumption after a positive technology shock can be seen clearly in the bottom panels of Figure 7, particularly the bottom-right panel. There, consumption falls substantially less in response to a negative technology shock (in period 1) if that shock was

---

[40] This is computed as follows: in period -1, the state variables of the model are initialized to their nonstochastic steady-state values. In the baseline ("one shock") scenario, ε_t^A is set equal to .007 in period 0, and set equal to 0 from period 1 onward. In the "two shock" scenario, ε_t^A is set equal to .007 in period 0, to .007 in period 1 (for the left-hand column of Figure 7, or to $-.007$ in period 1 for the right-hand column of Figure 7), and then set equal to 0 from period 2 onward. The impulse response function is computed as the period-by-period difference between the "two-shock" scenario and the baseline ("one shock") scenario, beginning in period 1.

preceded by a positive technology shock the period before (in period 0)—that is, the solid green line does not fall by as much as the dashed blue line. In the bottom-left panel, the response of consumption to a positive technology shock (in period 1) is fairly similar whether or not there was a positive technology shock in the previous period (period 0). This lower volatility of consumption growth in the bottom panels of Figure 7 is typical of other sizes of shocks as well.[41]

I now turn to the causes of this lower conditional volatility. The top row of Figure 7 plots the impulse responses of the economy's price dispersion, Δ_t. Note first how price dispersion tends to offset or attenuate the effects of the technology shock: for example, after a negative technology shock in the right-hand column of Figure 7, price dispersion falls, which tends to increase output all else equal (see equation 17). Although the net effect of the shock on output and consumption is still negative (see the bottom-right panel of Figure 7), the change in price dispersion moderates the effect of the technology shock. This is true for a positive technology shock in the left-hand column of Figure 7 as well.

Importantly, the effect of the technology shock on price dispersion is larger when prices are more distorted to begin with. For example, in the top-right panel of Figure 7, the solid green line falls *more* than the dashed blue line in response to the -0.7 percent shock. In the solid green line simulation, prices are more distorted to begin with because of the positive technology shock that hit the economy the period before.

Thus, after a positive technology shock, the moderating effects of price dispersion in the model are *greater*, causing the volatility of output and consumption to decline. In other words, a positive technology shock leads to a lower conditional volatility of consumption.

In Appendix C, I discuss in detail *why* Δ_t offsets the effect of the technology shock and prove that the response of Δ_t to a technology shock is larger when Δ_t itself is larger. However, the intuition for these effects can be seen in Figure 8. As shown in Appendix C, price dispersion $\Delta_t \geq 1$ holds in general, and $\log \Delta_t$ is a convex function of $\log(p_t^*/P_t)$, as in the figure. If we start the economy at the nonstochastic steady state and there is no uncertainty in the model, then firms find it optimal to set $p_t^* = P_t$ in period t. This corresponds to point A in Figure 8. Even if there is uncertainty, firms still find point A optimal in the linearized version of the model—at this point, firms' expected profits over the lifetime of the price contract are maximized, as in equation (14).

---

[41] As in previous figures, the effects are also essentially symmetric: if the economy is hit by a *negative* technology shock in period 0, then the conditional volatility in Figure 6 *increases* rather than decreases, and consumption in the bottom-right panel of Figure 7 is relatively *lower* after the negative technology shock—i.e., the solid green line lies below the dashed blue line.
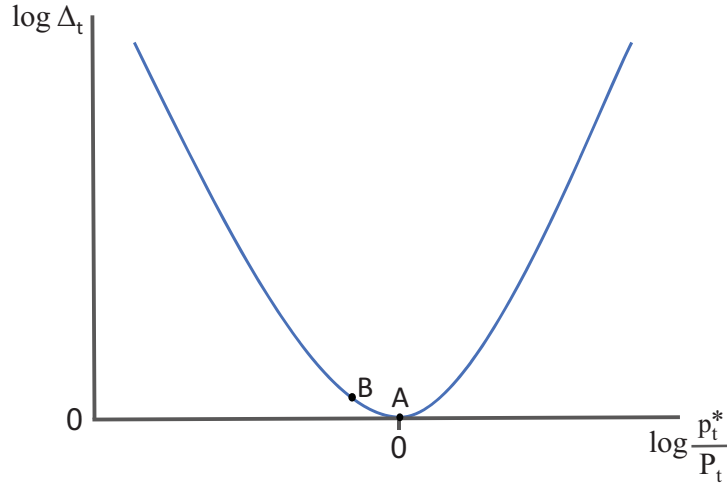
Figure 8. Illustrative graph of price dispersion Δ_t as a function of monopolistic firms' time-t relative price p_t^*/P_t. Point A denotes firms' risk-neutral optimal price (or the optimal price in the linearized model); point B denotes firms' risk-averse optimal price. See text for details.

In recessions, firms will lose a bit of profit because output is too low, and in expansions firms will lose a bit of profit because output (and hence marginal cost) is too high, but in expectation the firms' price strikes an optimal balance between these two.

Now consider the case where the firms' owners (here, the households) are risk averse, as in the nonlinear version of the model. In this case, recessions are more painful than expansions, so it is optimal for firms to put more weight on generating profit in recessions. This causes the firms' optimal price to be *lower* than in the risk-neutral case, so as to generate more output and profit in the event of recessions. Thus, the risk-averse firms' optimal price in Figure 8 lies at a point like B, to the left of A. In other words, firms have an incentive to set prices slightly lower than the trend rate of inflation, $\bar{\pi}$, would imply, in order to insure themselves in case of a recession.

The result is that price dispersion, Δ_t, in the stochastic economy will typically satisfy $\Delta_t > 1$, as at point B, rather than $\Delta_t = 1$. Note that at point B, the derivative $\partial \log \Delta_t / \partial \log(p_t^*/P_t) < 0$. Thus, positive technology shocks—which cause firms' marginal costs to fall and lower p_t^*/P_t— raise Δ_t. Firms are already setting prices too low from an aggregate efficiency standpoint (because they are risk-averse), and the positive technology shock exacerbates this inefficiency. Similarly, negative technology shocks lower Δ_t. This explains the impulse response functions for Δ_t in Figure 7, and why they move in such a way as to attenuate or offset the effects of the technology shock.

To see why this effect is larger when price dispersion is greater, notice that $\log \Delta_t$ in Figure 8 is a convex function of $\log(p_t^*/P_t)$. If Δ_t increases from point B, this implies moving further up

the curve to the left. At this point, the slope of the curve is more negative, implying a higher derivative. Thus, technology shocks lead to larger changes in Δ_t when Δ_t is larger to begin with, consistent with the results in Figure 7. It is because of this effect that the conditional volatility of consumption growth and the SDF fall after a positive technology shock, as in Figure 6. In turn, the decline in the volatility of the SDF is what causes risk premia to change after a shock. For more detailed derivations and discussion, see Appendix C.

5. Additional Discussion and Extensions

The macroeconomic model developed above is essentially a "proof of concept" that standard, structural macro models can be brought into agreement with a variety of asset pricing puzzles. Up to this point, I've deliberately kept the model as simple as possible in order to focus on the most crucial features for matching the broad behavior of macroeconomic variables and asset prices. Of course, some of those simplifying assumptions raise questions, and the model's overall simplicity invites extension in a number of directions, which I now discuss. First, I compare the model's assumption of a unitary intertemporal elasticity of substitution to the typical assumption that the IES > 1 in the long-run risks literature. Second, I discuss the relationship between the conditional heteroskedasticity in the model and the literature on "uncertainty shocks" (e.g., Bloom, 2009). Third, I extend the model to include additional shocks—in particular, a monetary policy shock and a fiscal policy shock—and discuss how this affects the results. Finally, I discuss the implications of the model's ability to price assets endogenously for financial frictions models, such as Bernanke, Gertler, and Gilchrist (1999), Kiyotaki and Moore (1997), and others.

5.1 The Intertemporal Elasticity of Substitution

In the long-run risks literature, such as Bansal and Yaron (2004), the intertemporal elasticity of substitution is typically assumed to be substantially greater than unity. There are two main reasons for that calibration: first, an IES > 1 ensures that a positive shock to consumption in their model causes stock prices to rise (rather than fall); and second, an IES > 1 ensures that an exogenous increase in volatility in their model causes stock prices to fall.

However, even in Bansal and Yaron (2004, henceforth BY), the assumption of an IES > 1 is not strictly necessary for these two criteria to be satisfied. For example, when equity represents a levered rather than unlevered consumption claim, then equity prices in BY rise in response to

a positive consumption shock if and only if the IES $> 1/\nu$, where ν is the degree of leverage. With leverage equal to 3, values of the IES down to $1/3$ still satisfy the criterion that a positive shock to consumption causes stock prices to rise. For a volatility shock, stock prices can respond negatively as long as the IES $> 1/\gamma$, where γ is the household's relative risk aversion.[42] If risk aversion is substantially greater than unity, then stock prices can respond negatively to volatility shocks even for values of the IES that are substantially less than unity.

Of course, the model in the present paper differs in many respects from that in BY and other reduced-form long-run risks models. In BY, consumption is an exogenous process with highly persistent shocks to consumption growth, while in the present paper, consumption is determined endogenously and is driven by exogenous technology growth shocks that are i.i.d. rather than persistent. Nevertheless, a positive shock to consumption (through technology) in the model here causes stock prices to rise, as can be seen clearly in Figure 2, even though the IES $= 1$. (In fact, the substitution effect dominates the wealth effect even for values of the IES that are somewhat below unity, consistent with the discussion above even though the model here differs from BY.) Thus, the model here satisfies the first criterion that positive output shocks cause stock prices to rise.

To investigate the second criterion—that an increase in volatility causes stock prices to fall—I extend the model to include exogenous stochastic volatility. In particular, let the standard deviation of the technology shock each period, $\sigma_{A,t}$, follow the autoregressive process

$$\log \sigma_{A,t} \;=\; (1 - \rho_\sigma) \log \bar{\sigma}_A \;+\; \rho_\sigma \log \sigma_{A,t-1} \;+\; \varepsilon_t^\sigma, \tag{38}$$

where $\bar{\sigma}_A = .007$, as in Table 1. Following Bansal and Yaron (2004), I calibrate $\rho_\sigma = 0.98$ and $\mathrm{Var}(\varepsilon_t^\sigma) = (0.1)^2$.[43]

Figure 9 plots the nonlinear impulse response functions of the model (computed the same way as in previous figures) to a positive one-standard-deviation shock to ε_t^σ. Volatility σ_t^A increases to about .0077 on impact and slowly declines back toward its initial level of .007. Consumption drops about 0.2 percent on impact, as households increase precautionary savings, and inflation falls gradually by about 0.1 percent in response to the decrease in demand. The increase in the

---

[42] In Bansal and Yaron (2004), the coefficient A_2 for the unlevered consumption claim requires $\theta < 0$, where $\theta = (1 - \gamma)/(1 - 1/\psi)$ and γ denotes risk aversion and ψ the IES in their paper. For the levered consumption claim, however, the coefficient $A_{2,m}$ requires $\theta/(1 - \theta) < 0$ (see their equation A20), which holds if either $\theta < 0$ or $\theta > 1$. Given $\gamma > 1$, then $\theta > 1$ if and only if $\psi > 1/\gamma$.

[43] Bansal and Yaron (2004) assume a more complicated (square-root rather than logarithmic) process for $\sigma_{A,t}$ than (38), but the magnitudes in (38) are essentially the same as theirs.
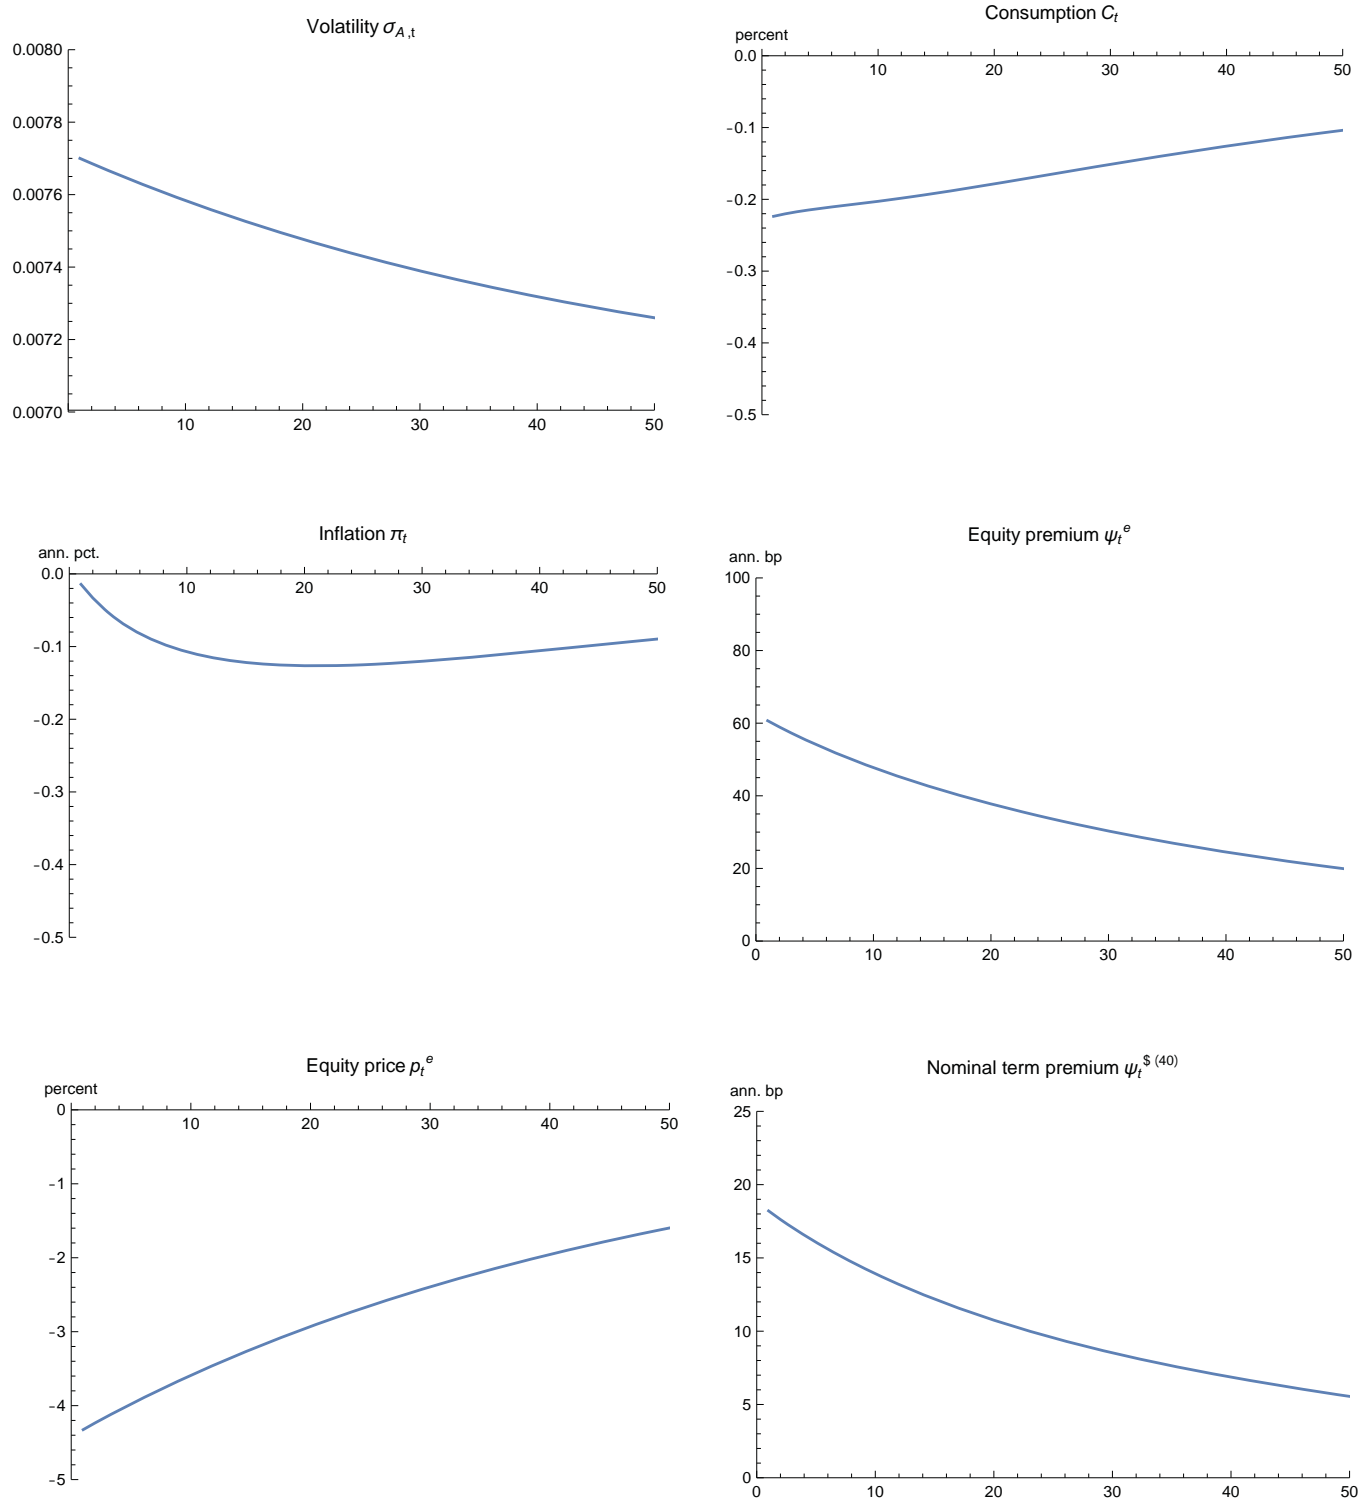
Figure 9. Nonlinear impulse response functions for volatility σ_t^A, consumption C_t, inflation π_t, the equity premium ψ_t^e, equity price p_t^e, and nominal term premium $\psi_t^{\$(40)}$ to a one-standard-deviation (0.1 percent, or .0007) positive volatility shock in the extended model. See text for details.

conditional volatility of consumption increases the volatility of the stochastic discount factor, which causes a large, 60 bp jump in the equity premium. (The nominal term premium also responds substantially to the volatility shock, increasing about 25 bp.) The large and persistent rise in the equity premium implies that the equity price must fall dramatically on impact, about 4.5 percent.[44] Thus, the model also satisfies the second criterion discussed above—that an exogenous increase in volatility causes stock prices to fall—without the need for an IES > 1, consistent with the discussion above even though the model differs from BY in many respects.

5.2 Uncertainty Shocks

Many authors use the volatility of stock prices as a measure of uncertainty and, as Bloom (2009) notes, recessions are typically associated with greater stock market uncertainty. Although one interpretation of this correlation is that exogenous increases in uncertainty lead to recessions, it is also possible that recessions lead to higher uncertainty in the stock market.

In fact, both of these channels can be seen clearly in the macroeconomic model developed above. The impulse response functions in Figure 9 depict the effects of an exogenous shock to volatility ε_t^σ (or uncertainty) in the model. As can be seen in that figure, consumption (and output) falls about 0.3 percent in response to the shock (and recovers quite slowly, because the shock in the model is so persistent). Thus, the model, extended to the case of exogenous stochastic volatility as in the previous section, matches the empirical findings in Bloom (2009) and others.

In addition, the model here implies that recessions cause uncertainty to increase *endogenously*, particularly as measured by volatility in the stock market. The intuition for this is essentially the same as for the model's endogenous conditional heteroskedasticity, discussed in Section 4. When the economy is weak, consumption is low and the household's stochastic discount factor becomes more sensitive to subsequent shocks. This drives up the equity premium, as shown in Figure 2, but it also increases the volatility of stock prices, as shown in Figure 10. In that latter figure, the one-step-ahead variance of the log equity price p_t^e, solved to fifth order using the same methods as above, is graphed as a function of technology A_t.[45] (Recall that A_t in the baseline parameterization of the model is a random walk, so the difference $\log A_t - \bar{y}_t$ essentially

---

[44] In order to generate an equity premium of 60 bp in the first period, stock prices must fall by about 0.6 percent below their second-period value. In order to generate an equity premium in each subsequent period, equity prices must continue to rise. This requires a large initial fall in the equity price so that in each subsequent period equity prices can rise in line with the implied equity premium.

[45] To compute the one-period-ahead variance in the model, I define two new variables, $\mathrm{expp}_t^e \equiv E_t p_{t+1}^e$ and $\mathrm{varp}_t^e \equiv E_t (p_{t+1}^e)^2 - (\mathrm{expp}_t^e)^2$. The latter variable, solved to fifth order, is plotted in Figure 10.
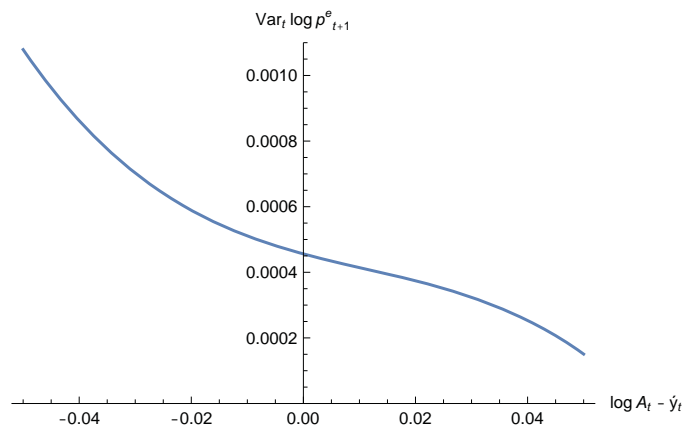
Figure 10. One-step-ahead variance of the log equity price p_t^e in the baseline model computed to fifth order. Low values of A_t relative to $\bar{y}_t$ correspond to a weaker economy, and to a more volatile equity price. See text for details.

"stochastically detrends" $\log A_t$.) When A_t is low (relative to $\bar{y}_t$), the economy is in a recession, with weakening output and consumption. And as can be seen in the figure, stock prices in the model are more volatile. When technology and output are about 5 percent below potential, the volatility of stock prices is about 2 to 2.5 times higher than when the economy is operating near potential.

Thus, the causality between recessions and uncertainty runs in both directions in the model developed above: exogenous increases in uncertainty cause the economy to weaken, but a weakening economy also causes uncertainty to rise. Changes in stock price volatility in the model are not a valid measure of "uncertainty shocks" without additional identifying assumptions that can plausibly isolate the *exogenous* changes in uncertainty from the *endogenous* changes in stock price volatility in response to the economy. (For empirical work along these lines, see Ludvigson, Ma, and Ng, 2016.)

5.3 The Financial Accelerator

Traditional models of the financial accelerator (e.g., Bernanke, Gertler, and Gilchrist, 1999; Kiyotaki and Moore, 1997; Gertler and Kiyotaki, 2011) allow for the possiblity of default, but ignore deviations from risk neutrality. The fact that borrowers might default introduces a wedge between borrowers and lenders that can act as an amplification and propagation mechanism for shocks: for example, in Kiyotaki and Moore (1997), a negative technology shock reduces the value of capital, which reduces firms' collateral; with less collateral, firms must scale back production and output falls by more than the effect of the technology shock alone. In Bernanke, Gertler, and Gilchrist

(1999, henceforth BGG), a weaker economy implies a higher probability of default, which raises costly state verification costs for financial intermediaries, which in turn leads to a greater spread between private borrowing rates and the risk-free rate.

The traditional financial accelerator mechanism captures many important features of a credit crunch and a financial crisis. At the same time, these models are essentially risk-neutral and abstract from risk premia—that is, in BGG and Kiyotaki and Moore (1997), the spread between private borrowing rates and the risk-free rate is just the risk-neutral expected loss from default. Yet an important part of the transmission mechanism in the 2007–08 financial crisis was the fall in the value of collateral beyond even the risk-neutral probability of default: for example, risk and liquidity premia rose dramatically even on securities that had little or no connection to subprime real estate lending (Gorton and Metrick, 2012), and the prices of many mortgage-backed securities and credit default swaps fell by much more than can be explained by any reasonable assumption for mortgage default and recovery rates (Stanton and Wallace, 2014). The dramatic increase in risk premia during the financial crisis caused huge drops in the value of collateral and historic increases in credit spreads. For the same reasons as in traditional financial accelerator models, we would expect these repercussions from rising risk premia to be an important part of the transmission mechanism from financial markets to the real economy.

Recently, some authors have begun to incorporate deviations from risk neutrality into financial accelerator models—see, e.g., He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014)—but those models are idiosyncratic. In contrast, the macroeconomic model I've developed here is canonical and would allow researchers to study the effects of risk premia on collateral values, credit spreads, and the economy within the well-understood and ubiquitous New Keynesian DSGE framework. Extending the model to include a financial intermediation sector is beyond the scope of the present paper, but would be a worthy topic for future research.

6. Conclusions

The simple macroeconomic model developed in this paper is consistent with a wide variety of asset pricing facts, such as the size and variability of the equity premium, real and nominal term premium, and the credit spread. Thus, I show that a wide variety of asset pricing puzzles are all consistent with the behavior of a very standard macroeconomic model—essentially the textbook New Keynesian model of Woodford (2003) and Galí (2008)—extended to include generalized recursive preferences, as in Epstein and Zin (1989) and Weil (1989).

Generalized recursive preferences allow the model to match the size and variability of risk premia in the data, essentially by increasing the model's risk aversion without greatly distorting its implications for macroeconomic variables. Thus, the variety of asset pricing puzzles that I consider in the present paper can all be thought of as different facets of a single, more fundamental puzzle: namely, why does risk aversion in macroeconomic models need to be so high to match the size and variability of risk premia in the data?

In the present paper, I do not directly address this last question, but good answers are provided by a number of recent studies in the macro-finance literature. For example, simple macroeconomic models substantially understate the true level of risk or uncertainty in the economy, such as general model uncertainty (e.g., Barillas, Hansen, and Sargent, 2009), parameter uncertainty (e.g., Weitzman, 2007), long-run risks (e.g., Bansal and Yaron, 2004), rare disasters (e.g., Rietz, 1988; Barro, 2006; Schmidt, 2015), and/or uninsurable idiosyncratic risk (e.g., Constantinides and Duffie, 1996; Schmidt, 2015). In addition, there is evidence that the consumption of stock- and bond-holders is more cyclical than that of non-asset-holders (e.g., Mankiw and Zeldes, 1991; Parker, 2001; Malloy, Moskowitz, and Vissing-Jorgensen, 2009), so the required level of risk aversion in a simple representative-agent model is higher than it would be in a model that recognized this heterogeneity (Guvenen, 2009). Similarly, Adrian, Etula, and Muir (2014) provide evidence that the marginal investor is a financial intermediary, whose principals' consumption is likely very highly correlated with market fluctuations. Extending the very simple, stylized model of the present paper to incorporate additional features along any of these lines would allow it to explain the asset pricing puzzles above with a substantially lower degree of risk aversion. The point of the present paper is not to incorporate all of these additional features and provide the best possible fit to the data, but rather to serve as a "proof of concept" that the asset pricing data can be matched within the standard New Keynesian modeling paradigm.

The simple, structural model I develop here provides a unified and intuitive framework for thinking about asset prices and asset pricing puzzles. Rather than studying each puzzle in isolation, the model here provides a reasonable description of the behavior of all the major asset classes. In addition, structural models have the well-known advantage of being more robust to structural breaks and novel policy interventions, such as those observed during the recent global financial crisis and European sovereign debt crisis. The model developed here can potentially provide insight in these situations, when more traditional, reduced-form models are largely uninformative.

Finally, by showing how a standard macroeconomic model can be made consistent with the behavior of risk premia in financial markets, the present paper opens the door to studying the feedback between those risk premia and the economy within the standard macreconomic modeling framework in use at central banks and other institutions around the world. As evidenced by the recent financial crises, this feedback from asset prices to the economy and back again can be very important. In the simple, stylized model of the present paper, asset prices have no feedback effects on the real economy, for simplicity. Nevertheless, it would be very interesting to combine the asset-pricing framework of the present paper with a macroeconomic model that includes a financial accelerator along the lines of Bernanke, Gertler, and Gilchrist (1999), Kiyotaki and Moore (1997), Gertler and Kiyotaki (2011), and others. In general, these models abstract from risk aversion and risk premia and focus instead on the effect of agency problems and collateral constraints on lending and investment. In a combined framework, shocks that cause the economy to deteriorate would lead to an increase in risk premia and a concomitant fall in asset prices, further amplifying the collateral constraint on firms and financial intermediaries. This channel appears to have been an important amplification and propagation mechanism in recent crises.

Appendix A: Model Equations

We can write the equations of the macroeconomic model in Section 2 in recursive form as follows. (Equations for equity and debt are essentially the same as in Section 3 and are not reproduced here.)

Value function:

$$V_t = (1-\beta)\left(\log C_t - \eta \frac{L_t^{1+\chi}}{(1+\chi)}\right) - \beta\alpha^{-1}\log \text{Vexp}_t, \tag{A1}$$

$$\text{Vexp}_t = E_t \exp(-\alpha V_{t+1}). \tag{A2}$$

Risk-free real rate and Euler equations:

$$e^{-r_t} = \beta E_t(C_{t+1}/C_t)^{-1}\big(\exp(-\alpha V_{t+1})/\text{Vexp}_t\big), \tag{A3}$$

$$C_t^{-1} = \beta E_t e^{i_t - \pi_{t+1}} C_{t+1}^{-1}\big(\exp(-\alpha V_{t+1})/\text{Vexp}_t\big). \tag{A4}$$

Optimal price setting by firms:

$$\left(\frac{p_t^*}{P_t}\right)^{1+\lambda(1-\theta)/((\lambda-1)\theta)} = \lambda \frac{z_t^n}{z_t^d}, \tag{A5}$$

$$z_t^n = \mu_t Y_t + \beta\xi E_t(C_{t+1}/C_t)^{-1}\big(\exp(-\alpha V_{t+1})/\text{Vexp}_t\big)(e^{\pi_{t+1}-\bar\pi})^{\lambda/((\lambda-1)\theta)} z_{t+1}^n, \tag{A6}$$

$$z_t^d = Y_t + \beta\xi E_t(C_{t+1}/C_t)^{-1}\big(\exp(-\alpha V_{t+1})/\text{Vexp}_t\big)(e^{\pi_{t+1}-\bar\pi})^{1/(\lambda-1)} z_{t+1}^d, \tag{A7}$$

$$(e^{\pi_t-\bar\pi})^{1/(1-\lambda)} = (1-\xi)\left(\frac{p_t^*}{P_t}\right)^{1/(1-\lambda)}(e^{\pi_t-\bar\pi})^{1/(1-\lambda)} + \xi. \tag{A8}$$

Marginal cost and real wage:

$$\mu_t = \frac{w_t}{P_t}\frac{Y_t^{(1-\theta)/\theta}}{\theta A_t^{1/\theta} K^{(1-\theta)/\theta}}, \tag{A9}$$

$$\eta L_t^\chi / C_t^{-1} = \frac{w_t}{P_t}. \tag{A10}$$

Production and resource constraint:

$$Y_t = A_t K^{1-\theta} L_t^\theta / \Delta_t, \tag{A11}$$

$$\Delta_t^{1/\theta} = (1-\xi)\left(\frac{p_t^*}{P_t}\right)^{-\lambda/((\lambda-1)\theta)} + \xi(e^{\pi_t-\bar\pi})^{\lambda/((\lambda-1)\theta)}\Delta_{t-1}^{1/\theta}, \tag{A12}$$

$$Y_t = C_t. \tag{A13}$$

Monetary policy rule:

$$i_t = \log(1/\beta) + \pi_t + \phi_\pi(\pi_t - \bar\pi) + \frac{\phi_y}{4}\log(Y_t/\bar Y_t), \tag{A14}$$

$$\log \bar Y_t = \rho_{\bar y}\log \bar Y_{t-1} + (1-\rho_{\bar y})\log Y_t. \tag{A15}$$

Technology shock:

$$\log A_t = \log A_{t-1} + \varepsilon_t^A. \tag{A16}$$

Equations (A1)–(A2) break the generalized value function into two equations to correspond to the syntax of Perturbation AIM and other rational expectations equation solvers, which typically require the model to be written as a system of equations in a form similar to $E_t F(X_{t-1}, X_t, X_{t+1}; \varepsilon_t) = 0$.

Equations (A5)–(A7) represent monopolistic firms' optimal price-setting equations. The exponent on (p_t^*/P_t) in (A5) follows from substituting out $y_{t+j}(f)$ in equation (14) in the main text, and is due to the presence of firm-specific capital stocks. Equations (A6)–(A7) are recursive versions of the infinite sums in the numerator and denominator of (14).

The other equations above follow in a straightforward manner from the equations in the main text.

Although capital stocks in the model above are fixed, the model nevertheless has a balanced growth path along which all variables are either constant or grow at constant rates if technology A_t itself grows at a constant rate. Along the balanced growth path, each of the variables Y_t, C_t, w_t, $\bar{Y}_t$, z_t^n, and z_t^d grow at the same rate as A_t. If we divide each of these variables through by A_t, the ratios have a nonstochastic steady state. Moreover, after a shock to A_t, these ratios converge back to their pre-shock levels. Thus, the nonstochastic steady state of these ratios constitutes a stable point around which we can approximate the model.

I thus transform the model by dividing each of the above variables by the level of technology A_t, and transform the value function V_t by defining $\tilde{V}_t \equiv V_t - \log A_t$. The transformed model then has a nonstochastic steady state around which I can compute an nth-order approximate solution as described in the text. These solutions are highly accurate in a neighborhood of the steady state, and become increasingly accurate over larger regions of the state space as the order of approximation n becomes large (see Swanson, Anderson, and Levin, 2006, for details and discussion).

Appendix B: Impulse Responses to a Negative Technology Shock

The nonlinear impulse response functions graphed in Figures 1 through 5 can be asymmetric for positive and negative shocks, because they are nonlinear. In practice, the New Keynesian model presented in the main text does not produce impulse response functions that are very asymmetric. This can be seen, for example, in Figure/ B1, which reproduces Figure 1 from the main text for the case of a negative one-standard-deviation ($-.007$) shock to technology A_t. The nonlinear impulse response functions in Figure B1 are computed in exactly the same was as in Figures 1 through 6, except that the shock has the opposite sign. Overall, the responses in Figure B1 are close to being symmetric counterparts to Figure 1.

Figure B2 presents nonlinear impulse response functions for the equity price p_t^e, equity premium ψ_t^e, real long-term bond price $p_t^{(40)}$, and term premium $\psi_t^{(40)}$ on the real long-term bond to the negative one-standard-deviation technology shock. The impulse responses are again close to being the symmetric counterparts to the nonlinear impulse responses functions in Figures 2 and 3, although the asset price and risk premium responses are slightly larger in magnitude for the negative shock than they are for the positive shock. For example, the equity premium increases by about 70 bp after the negative technology shock here, while it fell by about 62 bp after the positive shock in Figure 2. Similarly, the equity price falls by about 2.75 percent after the negative technology shock here, but rose by about 2.5 percent after the positive shock in Figure 2.

Figure B3 repeats the analysis for the nominal long-term bond price $p_t^{\$(40)}$, the term premium $\psi_t^{\$(40)}$ on the nominal long-term bond, defaultable bond price p_t^d, and credit spread $i_t^d - i_t^c$. As with Figure B2, the responses here are essentially symmetric to their counterparts in Figures 4 and 5, while also being somewhat larger in magnitude.
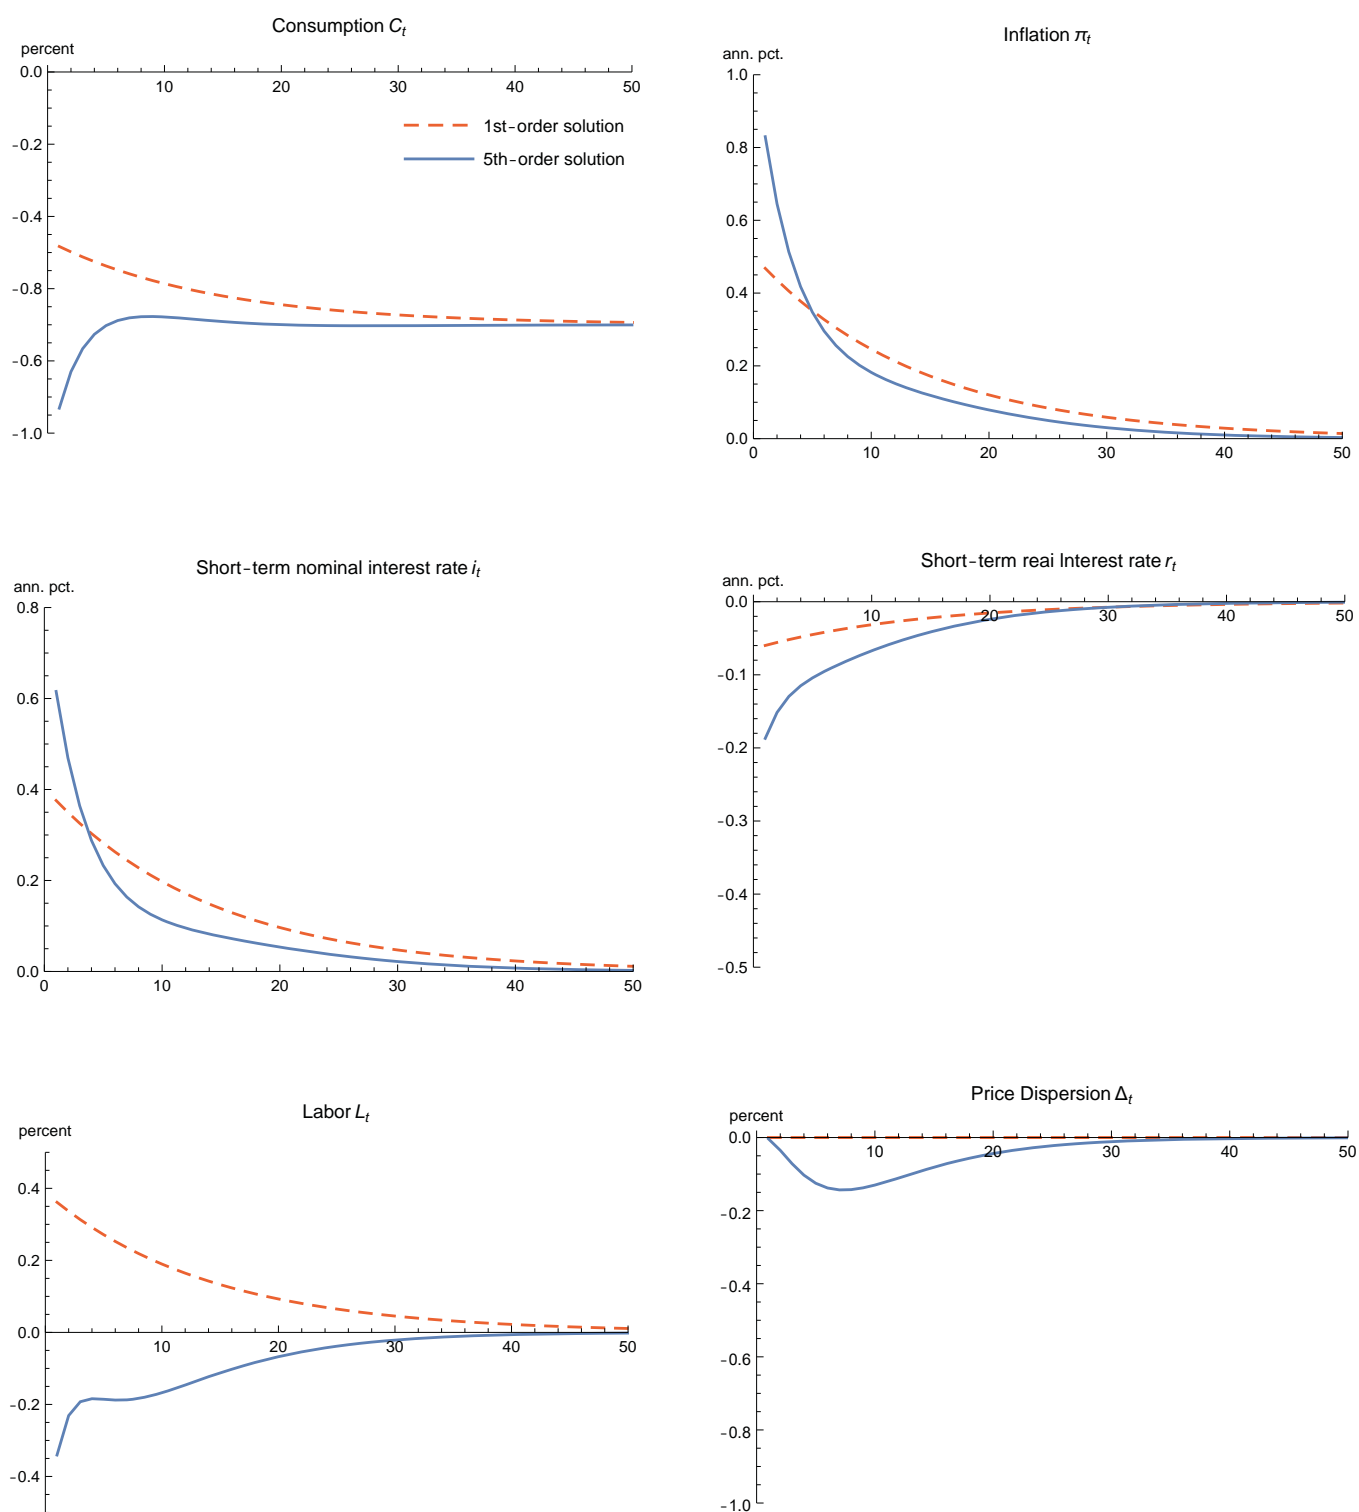
Figure B1. First-order (dashed red lines) and fifth-order (solid blue lines) impulse response functions for consumption C_t, inflation π_t, short-term nominal interest rate i_t, short-term real interest rate r_t, labor L_t, and price dispersion Δ_t to a one-standard-deviation *negative* (-0.7 percent) technology shock in the model. See Figure 1 for comparison and text for details.
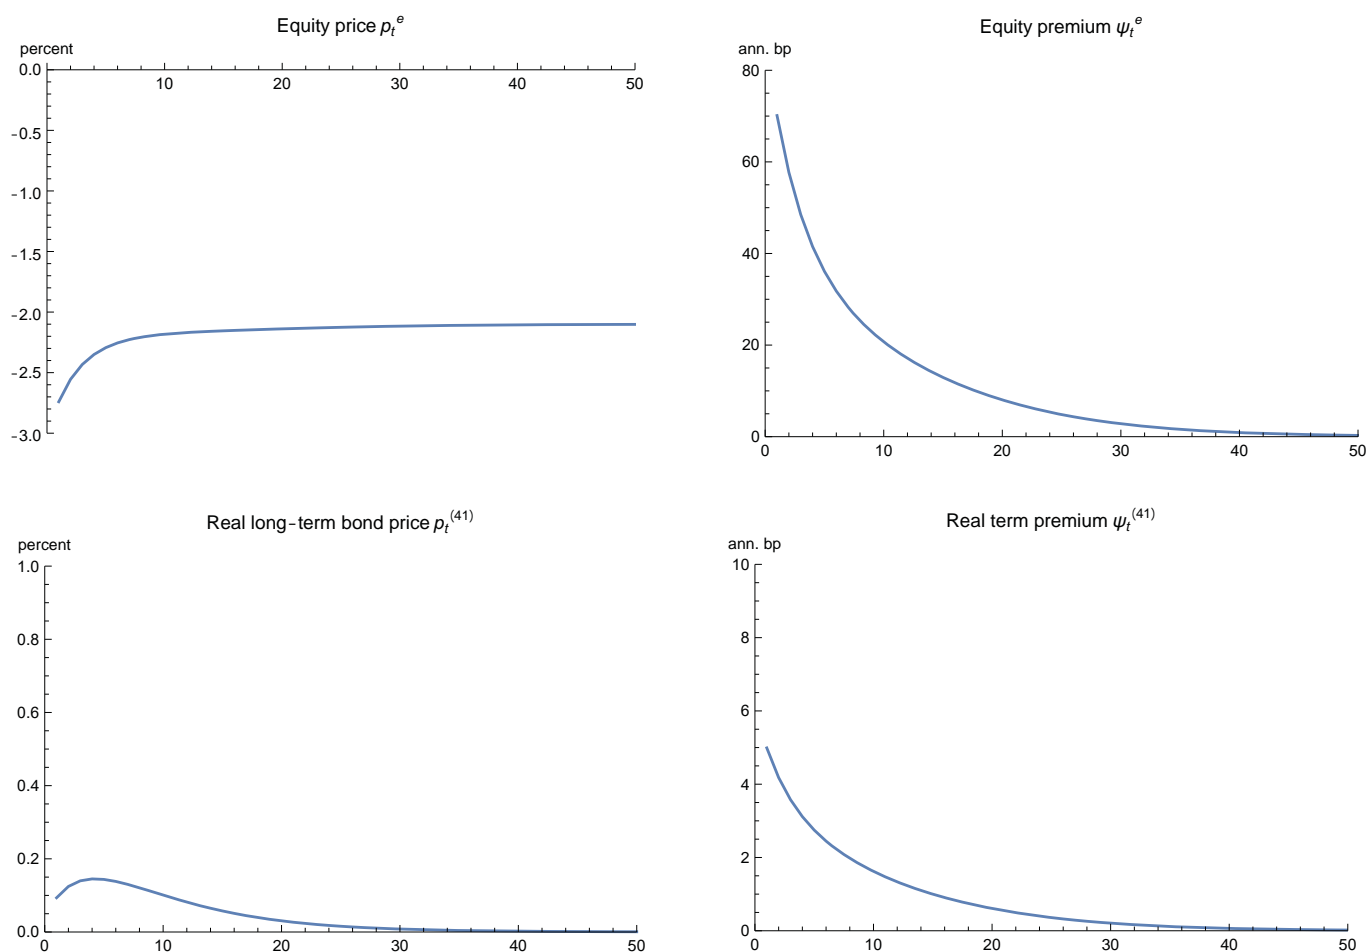
Figure B2. Nonlinear impulse response functions for the equity price p_t^e, equity premium ψ_t^e, real long-term bond price $p_t^{(40)}$ and real term premium $\psi_t^{(40)}$ to a one-standard-deviation *negative* (-0.7 percent) technology shock in the model, with state variables initialized to their nonstochastic steady state values. See Figures 2–3 for comparison and text for details.
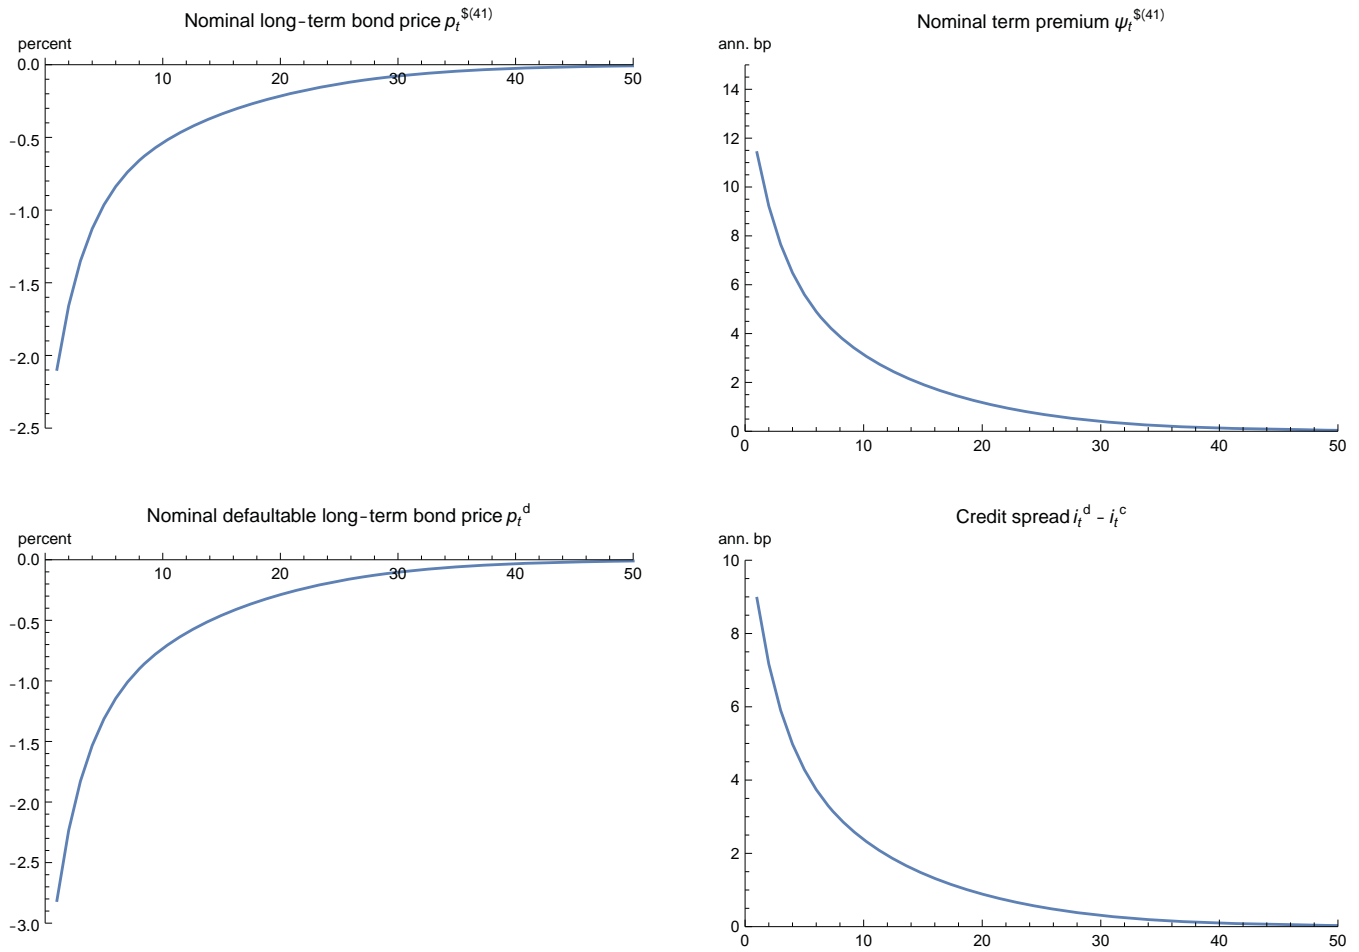
Figure B3. Nonlinear impulse response functions for the nominal long-term bond price $p_t^{\$(40)}$, nominal bond term premium $\psi_t^{\$(40)}$, defaultable long-term bond price p_t^d, and credit spread $i_t^d - i_t^c$ to a one-standard-deviation *negative* (-0.7 percent) technology shock in the model, with state variables initialized to their nonstochastic steady state values. See Figures 4–5 for comparison and text for details.

Appendix C: Price Dispersion and Endogenous Heteroskedasticity

As discussed in Section 4, price dispersion Δ_t plays a key role in generating endogenous conditional heteroskedasticity in the model. Here I explain the behavior of price dispersion and endogenous conditional heteroskedasticity in detail.

The first facts to note about price dispersion are the following (which are standard in the New Keynesian literature):

Proposition 1. *Price dispersion $\Delta_t \geq 1$. On the model's balanced growth path, $\Delta_t = 1$.*

PROOF: Recall that

$$\Delta_t \equiv \left[\int_0^1 \left(\frac{p_t(f)}{P_t} \right)^{\lambda/((1-\lambda)\theta)} df \right]^\theta. \tag{C1}$$

Because $\lambda > 1$ and $\theta \in (0,1)$, the function $x^{\lambda/\theta}$ is convex. Jensen's inequality then implies that

$$\left(\int_0^1 x(f)^{\lambda/\theta} df \right)^{\theta/\lambda} \geq \int_0^1 x(f) df. \tag{C2}$$

Letting $x(f) \equiv (p_t(f)/P_t)^{1/(1-\lambda)}$ implies

$$\Delta_t \geq \int_0^1 \left(\frac{p_t(f)}{P_t} \right)^{1/(1-\lambda)} df = 1, \tag{C3}$$

where the equality in (C3) follows from the definition of P_t. On the model's balanced growth path (the nonstochastic steady state of the transformed model), $p_t(f) = P_t$ for all f, implying $\Delta_t = 1$. $\square$

Writing equation (C1) in recursive form, we have

$$\Delta_t^{1/\theta} = (1-\xi) \left(\frac{p_t^*}{P_t} \right)^{-\lambda/((\lambda-1)\theta)} + \xi (e^{\pi_t - \bar{\pi}})^{\lambda/((\lambda-1)\theta)} \Delta_{t-1}^{1/\theta}. \tag{C4}$$

Writing equation (13) in recursive form, we have

$$P_t^{1/(1-\lambda)} = (1-\xi)(p_t^*)^{1/(1-\lambda)} + \xi e^{\bar{\pi}/(1-\lambda)} P_{t-1}^{1/(1-\lambda)}, \tag{C5}$$

which implies

$$e^{(\pi_t - \bar{\pi})/(\lambda-1)} = \frac{1}{\xi} - \frac{1-\xi}{\xi} \left(\frac{p_t^*}{P_t} \right)^{1/(1-\lambda)}. \tag{C6}$$

If we start the economy on its balanced growth path, then $\Delta_{t-1} = 1$, and Δ_t satisfies

$$\Delta_t^{1/\theta} = (1-\xi) \left(\frac{p_t^*}{P_t} \right)^{-\lambda/((\lambda-1)\theta)} + \xi^{1-\lambda/\theta} \left[1 - (1-\xi) \left(\frac{p_t^*}{P_t} \right)^{1/(1-\lambda)} \right]^{\lambda/\theta}. \tag{C7}$$

Price dispersion Δ_t is then a function of monopolistic firms' relative price at time t, p_t^*/P_t. The derivative of $\log \Delta_t$ with respect to the log relative price is

$$\frac{d \log \Delta_t}{d \log \frac{p_t^*}{P_t}} = \Delta_t^{-1/\theta} (1-\xi) \frac{\lambda}{\lambda-1} \left(\frac{p_t^*}{P_t} \right)^{-\lambda/((\lambda-1)\theta)} \left\{ \left[\frac{1}{\xi} \left(\frac{p_t^*}{P_t} \right)^{1/(\lambda-1)} - \frac{1-\xi}{\xi} \right]^{\lambda/\theta - 1} - 1 \right\}. \tag{C8}$$

The term before the curly brackets on the right-hand side of (C8) is always positive. The term inside the curly brackets is zero for $p_t^*/P_t = 1$, positive for $p_t^*/P_t > 1$, and negative for $p_t^*/P_t < 1$.

Proposition 2. *Given $\Delta_{t-1} = 1$, $\log \Delta_t$ is a convex function of $\log(p_t^*/P_t)$ in a neighborhood of 0.*

PROOF: Δ_t is minimized at the point $p_t^*/P_t = 1$, and (C8) is differentiable with respect to $\log(p_t^*/P_t)$. Thus, $d^2 \log \Delta_t / d \log(p_t^*/P_t)^2 > 0$ at $\log(p_t^*/P_t) = 0$. $\qquad\square$

Proposition 2 implies the graph of $\log \Delta_t$ as a function of the log relative price is as depicted in Figure 8, in a neighborhood of $p_t^*/P_t = 1$.

[To be completed.]

References

ABEL, ANDREW (1999). "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics* 43, 3–33.

ADRIAN, TOBIAS, ERKKO ETULA, AND TYLER MUIR (2014). "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance,* forthcoming.

ALTIG, DAVID, LAWRENCE CHRISTIANO, MARTIN EICHENBAUM, AND JESPER LINDÉ (2011). "Firm-Specific Capital, Nominal Rigidities, and the Business Cycle," *Review of Economic Dynamics* 14, 225–247.

ALTONJI, JOSEPH (1986). "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy* 94, S176–S215.

ANDREASEN, MARTIN (2012). "An Estimated DSGE Model: Explaining Variation in Nominal Term Premia, Real Term Premia, and Inflation Risk Premia," *European Economic Review* 56, 1656–1674.

ANG, ANDREW, AND MAXIM ULRICH (2012). "Nominal Bonds, Real Bonds, and Equity," unpublished manuscript, Columbia Business School.

ARUOBA, S. BORAGAN, JESÚS FERNÁNDEZ-VILLAVERDE, AND JUAN RUBIO-RAMÍREZ (2006). "Comparing Solution Methods for Dynamic Equilibrium Economies," *Journal of Economic Dynamics and Control* 30, 2477–2508.

BACKUS, DAVID, ALLAN GREGORY, AND STANLEY ZIN (1989). "Risk Premiums in the Term Structure," *Journal of Monetary Economics* 24, 371–399.

BAELE, LIEVEN, GEERT BEKAERT, AND KOEN INGHELBRECHT (2010). "The Determinants of Stock and Bond Comovements," *Review of Financial Studies* 23, 2374–2428.

BANSAL, RAVI, AND AMIR YARON (2004). "Risks for the Long Run: A Potential Resolution of Asset-Pricing Puzzles," *Journal of Finance* 59, 1481–1509.

BARILLAS, FRANCISCO, LARS HANSEN, AND THOMAS SARGENT (2009). "Doubts or Variability?" *Journal of Economic Theory* 144, 2388–2418.

BARRO, ROBERT (2006). "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics* 121, 823–866.

BEKAERT, GEERT, SEONGHOON CHO, AND ANTONIO MORENO (2010). "New Keynesian Macroeconomics and the Term Structure," *Journal of Money, Credit and Banking* 42, 33–62.

BEKAERT, GEERT, ERIC ENGSTROM, AND STEVEN GRENADIER (2010). "Stock and Bond Returns with Moody Investors," *Journal of Empirical Finance* 17, 867–894.

BERNANKE, BEN, MARK GERTLER, AND SIMON GILCHRIST (1999). "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of Macroeconomics* 1, 1341–1393.

BHAMRA, HARJOAT, LARS-ALEXANDER KUEHN, AND ILYA STREBULAEV (2010). "The Levered Equity Risk Premium and Credit Spreads: A Unified Framework," *Review of Financial Studies* 23, 645–703.

BLOOM, NICHOLAS (2009). "The Impact of Uncertainty Shocks," *Econometrica* 77, 623–685.

BOLDRIN, MICHELE, LAWRENCE CHRISTIANO, AND JONAS FISHER (2001). "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review* 91, 149–166.

BRUNNERMEIER, MARKUS, AND YULIY SANNIKOV (2014). "A Macroeconomic Model with a Financial Sector," *American Economic Review* 104, 379–421.

CALVO, GUILLERMO (1983). "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12, 383–398.

CAMPANALE, CLAUDIO, RUI CASTRO, AND GIAN LUCA CLEMENTI (2010). "Asset Pricing in a Production Economy with Chew-Dekel Preferences," *Review of Economic Dynamics* 13, 379–402.

CAMPBELL, JOHN (1999). "Asset Prices, Consumption, and the Business Cycle," *Handbook of Macroeconomics* 1, ed. John Taylor and Michael Woodford, Elsevier, 1231–1303.

CAMPBELL, JOHN, AND JOHN COCHRANE (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy* 107, 205–251.

CAMPBELL, JOHN, CAROLIN PFLUEGER, AND LUIS VICEIRA (2013). "Monetary Policy Drives of Bond and Equity Risks," unpublished manuscript, Harvard University.

CAMPBELL, JOHN, ADI SUNDARAM, AND LUIS VICEIRA (2013). "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds," unpublished manuscript, Harvard University.

CHEN, HUI (2010). "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure," *Journal of Finance* 65, 2171–2212.

CHEN, LONG, PIERRE COLLIN-DUFRESNE, AND ROBERT GOLDSTEIN (2009). "On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle," *Review of Financial Studies* 22, 3367–3409.

CHRISTIANO, LAWRENCE, MARTIN EICHENBAUM, AND CHARLES EVANS (2005). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* 113, 1–45.

CLARIDA, RICHARD, JORDI GALÍ, AND MARK GERTLER (1999). "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37, 1661–1707.

COCHRANE, JOHN (2008). "Financial Markets and the Real Economy," *Handbook of the Equity Risk Premium,* 237–326.

COCHRANE, JOHN, AND MONIKA PIAZZESI (2005). "Bond Risk Premia," *American Economic Review* 95, 138–160.

CONSTANTINIDES, GEORGE, AND DARRELL DUFFIE (1996). "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104, 219–240.

COOPER, ILAN, AND RICHARD PRIESTLEY (2008). "Time-Varying Risk Premiums and the Output Gap," *Review of Financial Studies* 22, 2801–2833.

DAI, QIANG, AND KENNETH SINGLETON (2003). "Fixed-Income Pricing," *Handbook of the Economics of Finance* 1B, 1207–1246.

DEL NEGRO, MARCO, FRANK SCHORFHEIDE, AND MARC GIANNONI (2015). "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics* 7, 168–196.

DEN HAAN, WOUTER (1995). "The Term Structure of Interest Rates in Real and Monetary Economies," *Journal of Economic Dynamics and Control* 19, 909–940.

DEW-BECKER, IAN (2012). "A Model of Time-Varying Risk Premia with Habits and Production," unpublished manuscript, Fuqua School of Business, Duke University.

DONALDSON, JOHN, THORE JOHNSEN, AND RAJNISH MEHRA (1990). "On the Term Structure of Interest Rates," *Journal of Economic Dynamics and Control* 14, 571–596.

DUFFIE, DARREL, AND DAVID LANDO (2001). "Term Structures of Credit Spreads with Incomplete Accounting Information," *Econometrica* 69, 633–664.

EPSTEIN, LARRY, AND STANLEY ZIN (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57, 937–969.

ERAKER, BJORN (2008). "Affine General Equilibrium Models," *Management Science* 54, 2068–2080.

EUROPEAN CENTRAL BANK (2012). "Technical Features of Outright Monetary Transactions," press release, 6 September, http://www.ecp.europa.eu/press/pr/date/2012/html/pr120906_1.en.html.

FAMA, EUGENE, AND KENNETH FRENCH (1989). "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics* 25, 23–49.

Fama, Eugene, and Kenneth French (2002). "The Equity Premium," *Journal of Finance* 57, 637–659.

Fama, Eugene, and Kenneth French (2015). "A Five-Factor Asset Pricing Model," *Journal of Financial Economics* 116, 1–22.

Galí, Jordi (2008). *Monetary Policy, Inflation, and the Business Cycle,* Princeton University Press.

Gertler, Mark, and Nobuhiro Kiyotaki (2011). "Financial Intermediation and Credit Policy in Business Cycle Analysis," *Handbook of Monetary Economics* 3A, 547–599.

Giesecke, Kay, Francis Longstaff, Stephen Schaefer, and Ilya Strebulaev (2011). "Corporate Bond Default Risk: A 150-Year Perspective," *Journal of Financial Economics* 102, 233–250.

Gorton, Gary, and Andrew Metrick (2012). "Getting Up to Speed on the Financial Crisis: A One-Weekend-Reader's Guide," *Journal of Economic Literature* 50, 128–150.

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

Gürkaynak, Refet, Brian Sack, and Jonathan Wright (2007). "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics* 54, 2291–2304.

Gürkaynak, Refet, Brian Sack, and Jonathan Wright (2010). "The TIPS Yield Curve and Inflation Compensation," *American Economic Journal: Macroeconomics* 2, 70–92.

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

Hall, Robert (1988). "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96, 339–357.

Hansen, Lars P. and Thomas Sargent (2001). "Robust Control and Model Uncertainty," *American Economic Review* 91, 60–66.

He, Zhiguo, and Arvind Krishnamurthy. "Intermediary Asset Pricing," *American Economic Review* 103 (2013), 732–770.

Judd, Kenneth (1988). *Numerical Methods in Economics,* MIT Press.

Kim, Don, and Jonathan Wright (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.

King, Robert, Charles Plosser, and Sergio Rebelo (1988). "Production, Growth, and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics* 21, 195–232.

King, Robert, Charles Plosser, and Sergio Rebelo (2002). "Production, Growth, and Business Cycles: Technical Appendix," *Computational Economics* 20, 87–116.

King, Robert, and Sergio Rebelo (1999). "Resuscitating Real Business Cycles," in *Handbook of Macroeconomics* 1, ed. John Taylor and Michael Woodford, Elsevier, 927–1007.

Kiyotaki, Nobuhiro, and John Moore (1997). "Credit Cycles," *Journal of Political Economy* 105, 211–248.

Koijen, Ralph, Hanno Lustig, and Stijn Van Nieuwerburgh (2013). "The Cross-Section and Time Series of Stock and Bond Returns," *NBER Working Paper* 15688.

Kydland, Finn, and Edward Prescott (1982). "Time to Build and Aggregate Fluctuations," *Econometrica* 50, 1345–1370.

Lane, Phillip (2012). "The European Sovereign Debt Crisis," *Journal of Economic Perspectives* 26, 49–68.

Leland, Hayne (1994). "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure," *Journal of Finance* 49, 1213–1252.

Leland, Hayne (1998). "Agency Costs, Risk Management, and Capital Structure," *Journal of Finance* 53, 1213–1243.

Lettau, Martin, and Sydney Ludvigson (2010). "Measuring and Modeling Variation in the Risk-Return Trade-off," *Handbook of Financial Econometrics* 1, 617–690.

Lettau, Martin, and Harald Uhlig (2000). "Can Habit Formation Be Reconciled with Business Cycle Facts?" *Review of Economic Dynamics* 3, 79–99.

Lettau, Martin, and Jessica Wachter (2011). "The Term Structure of Equity and Interest Rates," *Journal of Financial Economics* 101, 90–113.

Ludvigson, Sydney, Sai Ma, and Serena Ng (2016). "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" unpublished manuscript, New York University.

MaCurdy, Thomas (1981). "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy* 89, 1059–1085.

Malloy, Christopher, Tobias Moskowitz, and Annette Vissing-Jorgensen (2009). "Long-Run Stockholder Consumption, Risk and Asset Returns," *Journal of Finance* 44, 2427–2479.

Mankiw, N. Gregory, and Stephen Zeldes (1991). "The Consumption of Stockholders and Non-stockholders," *Journal of Financial Economics* 29, 97–112.

Mehra, Rajnish, and Edward Prescott (1985). "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145–161.

Mishkin, Frederic (2011). "Over the Cliff: From the Subprime to the Global Financial Crisis," *Journal of Economic Perspectives* 25, 49–70.

Parker, Jonathan (2001). "The Consumption Risk of the Stock Market," *Brookings Papers on Economic Activity* 2, 279–333.

Piazzesi, Monika, and Martin Schneider (2006). "Equilibrium Yield Curves," *NBER Macroeconomics Annual* 21, 389–442.

Piazzesi, Monika, and Eric Swanson. "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy," *Journal of Monetary Economics* 55 (2008), 677–691.

Rietz, Thomas (1988). "The Equity Risk Premium: A Solution," *Journal of Monetary Economics* 22, 117–131.

Roll, Richard (1977). "A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory," *Journal of Financial Economics* 4, 129–176.

Rudebusch, Glenn (2002). "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia," *Journal of Monetary Economics* 49, 1161–1187.

Rudebusch, Glenn, Brian Sack, and Eric Swanson (2007). "Macroeconomic Implications of Changes in the Term Premium," *Federal Reserve Bank of St. Louis Economic Review* 89, 241–269.

Rudebusch, Glenn, and Eric Swanson (2008). "Examining the Bond Premium Puzzle with a DSGE Model," *Journal of Monetary Economics* 55, 111–126.

Rudebusch, Glenn, and Eric Swanson (2012). "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks," *American Economic Journal: Macroeconomics* 4, 105–143.

Rudebusch, Glenn, and Tao Wu (2007). "Accounting for a Shift in Term-Structure Behavior with No-Arbitrage Macro-Finance Models," *Journal of Money, Credit and Banking* 39, 395–422.

Schmidt, Lawrence (2015). "Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk," unpublished manuscript, University of Chicago.

SMETS, FRANK, AND RAFAEL WOUTERS (2007). "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97, 586–606.

STANDARD & POOR'S (2011). *2011 Annual Global Corporate Default Study and Rating Transitions,* New York.

STANTON, RICHARD, AND NANCY WALLACE (2014). "The Bear's Lair: Index Credit Default Swaps and the Subprime Mortgage Crisis," *Review of Financial Studies* , forthcoming.

SWANSON, ERIC (2013). "Risk Aversion, Risk Premia, and the Labor Margin with Generalized Recursive Preferences," *Federal Reserve Bank of San Francisco Working Paper* 2012–17.

SWANSON, ERIC, GARY ANDERSON, AND ANDREW LEVIN (2006). "Higher-Order Perturbation Solutions to Dynamic, Discrete-Time Rational Expectations Models," *Federal Reserve Bank of San Francisco Working Paper* 2006–01.

TALLARINI, THOMAS (2001). "Risk-Sensitive Business Cycles," *Journal of Monetary Economics* 45, 507–532.

TAYLOR, JOHN (1993). "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39, 195–214.

TAYLOR, JOHN (1999). "A Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules,* ed. John Taylor (MIT Press: Cambridge MA), 319–341.

VAN BINSBERGEN, JULES, JESÚS FERNÁNDEZ-VILLAVERDE, RALPH KOIJEN, AND JUAN RUBIO-RAMÍREZ (2012). "The Term Structure of Interest Rates in a DSGE Model with Recursive Preferences," *Journal of Monetary Economics* 59, 634–648.

VISSING-JORGENSEN, ANNETTE (2002). "Limited Asset Market Participation and the Elasticity of Substitution," *Journal of Political Economy* 110, 825–853.

The Wall Street Journal (2012). "Secret Weapon: Europe's Loan Plan," Tom Lauricella and Matt Phillips, January 23.

WEIL, PHILLIPE (1989). "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics* 24, 401–421.

WEITZMAN, MARTIN (2007). "Subjective Expectations and Asset-Return Puzzles," *American Economic Review* 97, 1102–1130.

WOODFORD, MICHAEL (2003). *Interest and Prices,* Princeton University Press.