

King, Thomas B.

Working Paper

A portfolio-balance approach to the nominal term structure

Working Paper, No. 2013-18

Provided in Cooperation with:

Federal Reserve Bank of Chicago

Suggested Citation: King, Thomas B. (2013) : A portfolio-balance approach to the nominal term structure, Working Paper, No. 2013-18, Federal Reserve Bank of Chicago, Chicago, IL

This Version is available at:

<https://hdl.handle.net/10419/96632>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Federal Reserve Bank of Chicago

A Portfolio-Balance Approach to the Nominal Term Structure

Thomas B. King

November 2013

WP 2013-18

A Portfolio-Balance Approach to the Nominal Term Structure

Thomas B. King*
Department of Economic Research
Federal Reserve Bank of Chicago

November 20, 2013

Abstract

Explanations of why changes in the relative quantities of safe debt seem to affect asset prices often appeal informally to a “portfolio balance” mechanism. I show how this type of effect can be incorporated in a general class of structural, arbitrage-free asset-pricing models using a numerical solution method that allows for a wide range of nonlinearities. I consider some applications in which the Treasury market is isolated, investors have mean-variance preferences, and the short-rate process is truncated at zero. Despite its simplicity, a version of this model incorporating inflation can fit longer-term yields well, and it suggests that fluctuations in Treasury supply explain a sizeable fraction of the historical time-series variation in term premia. Nonetheless, under plausible parameterizations central-bank asset purchases have a fairly small impact on the yield curve by removing duration from the market, and these effects are particularly weak when interest rates are close to their zero lower bound.

JEL Codes: C63, E43, E44, E52, E58, G11, G12

Keywords: Yield curve, LSAP, quantitative easing, preferred habitat, forward guidance

* Contact: Thomas.King@chi.frb.org. I thank Stefania D’Amico, Jonathan Goldberg, Francisco Vazquez-Grande, and seminar participants at the Federal Reserve Bank of Chicago and the Federal Reserve Board for helpful comments and discussions. Roger Fan provided excellent research assistance. The views expressed here do not reflect official positions of the Federal Reserve.

1 Introduction

Whether and how the mix of government liabilities affects financial conditions is relevant for optimal debt-management policy, the potential role of safe assets in the economy, and the transmission of monetary and fiscal policy through financial markets. These questions also have implications for theoretical and empirical modeling of interest rates, asset prices, and macrofinancial dynamics. They have received particular attention following the recent efforts of several central banks to reduce long-term interest rates by purchasing large quantities of government debt and other securities. In the United States, this policy has taken the form of the Federal Reserve’s Large-Scale Asset Purchase (LSAP) programs, which have, at the time of this writing, removed over \$3 trillion of long-term government debt from the market.

Recent empirical work has been nearly universal in concluding that fluctuations in the distribution of government liabilities do have effects on the term structure of interest rates and, most likely, other asset prices.¹ The story that is most often told to explain these findings is that changes in the amounts of long-term debt alter the quantity of interest-rate or “duration” risk held by the public and thereby affect the equilibrium price of that risk.² Yet, while one can find narrative versions of this “portfolio balance” mechanism dating at least to Culbertson (1957), only a few academic papers have tried to formalize the idea. In particular, Vayanos and Vila (2009) is to date the only study to incorporate portfolio balance—or, indeed, supply effects of any kind—into a no-arbitrage model of the term structure under rational expectations. However, as discussed further below, certain aspects of that model may complicate its application to the data and use in policy analysis.

This paper considers a broad class of arbitrage-free, rational-expectations portfolio-balance models in which the relative amounts of default-free assets held by the public can matter for the term structure of interest rates and other asset prices. I solve the models computationally in a way that is robust to functional forms and other modeling choices. In particular, the approach allows for departures from linearity, normality, and homoscedasticity, and the distribution of securities held by investors can shift in an essentially unrestricted way—features that may be particularly important when considering an

¹ See Bernanke et al. (2004), Kuttner (2006), Gagnon et al. (2010), Greenwood and Vayanos (2010), Krishnamurthy and Vissing-Jorgensen (2011), Meaning and Zhu (2011), D’Amico et al., (2012), Hamilton and Wu (2012), Ihrig et al. (2012), Joyce et al. (2012), Li and Wei (2012), Rosa (2012), D’Amico and King (2013), Cahill et al. (2013), and Chabot and Herman (forthcoming).

² For example, Gagnon et al. (2010) appeal strongly to this idea. Federal Reserve officials who have advocated a similar view include Sack (2009) and Bernanke (2010). With respect to LSAPs, Bauer and Rudebusch (forthcoming), argue that some of the effect likely results from the signals those programs convey to investors about future short-term rates. However, several of the studies listed above find that asset purchases also have large effects on term premia or that they alter the configuration of longer-term rates in ways that cannot easily be explained through expectations for the short-rate path. Moreover, the signaling channel cannot account for the evidence from long time series or other historical episodes, which is largely driven by debt fluctuations that are independent of monetary policy.

environment in which the short-term rate is near its zero lower bound (ZLB) or situations in which there are large, infrequent shocks to outstanding supply. A discrete-time version of the Vayanos-Vila model is a special case. It is also straightforward to evaluate portfolio-balance models that incorporate attributes such as stochastic volatility and time-varying risk aversion and models in which investors derive convenience yields from certain securities, as in Krishnamurthy and Vissing-Jorgenson (2012).

The requirement that arbitrage opportunities cannot exist in equilibrium severely restricts the types of effects that supply fluctuations may have on the yield curve. Broadly speaking, “local supply” effects, in which the yields on particular securities move in direct response to fluctuations in the amounts of those specific securities outstanding, are ruled out. Instead, asset quantities affect asset prices by changing the equilibrium return on the market portfolio, since the latter is just the sum of the returns on individual assets weighted by the quantities held. So long as the stochastic discount factor depends on the market return, changing asset quantities changes the vector of asset returns itself. Of course, this in turn changes the market return—which changes the discount factor, which changes the market return, and so on. The algorithm for solving the models exploits this recursion. Adding embellishments such as imperfect substitutability between different securities, preferred-habitat demand, or additional risk factors like inflation, can alter the shape and magnitude of the supply effects but does not affect the basic transmission channel. Indeed, the key mechanism potentially applies to any model—for example, those involving recursive preferences over consumption—in which the pricing kernel involves the return on wealth.³

For exposition, I apply the approach to some simple models in which investors have myopic preferences over asset returns and only hold government liabilities. These restrictions are useful for isolating the effects of duration risk cleanly and for maintaining comparability to the previous literature. But another key advantage is that, by setting aside Wallace neutrality, the Roll critique, intertemporal substitution, and long investment horizons, these assumptions attribute the maximum possible puissance to the duration channel. The results thus likely represent an upper bound on the size of supply effects that could plausibly be consistent with rational expectations and an equilibrium absence of arbitrage opportunities. Despite the simplicity and starkness of their assumptions, these models match the yield-curve data surprisingly well. However, it is important to note that, although accounting for the quantities held by the market is a crucial element of the approach, variation in quantities themselves contributes a relatively small portion of the variation in yields in practice; as in most term-structure models, expectations for short rates are the primary drivers of yields, at least for short to medium maturities.

³ As discussed further below, this mechanism does not conflict with Eggertsson and Woodford’s (2003) well known irrelevance proposition for portfolio-balance effects, because their proof applies to cases in which the pricing kernel is independent of the return on wealth.

Nonetheless, changes in Treasury supply are estimated to have contributed significantly to variation in term premia. In particular, the increase in the duration of government debt held by the public in the late 1970s and early 1980s appears to have boosted the ten-year term premium by over 100 basis points. On the other hand, most term-premium movements since that time—notably, their large secular decline—are estimated to have been due to fluctuations in inflation risk, consistent with previous empirical evidence. (See, for example, the literature surveyed by Gurkaynak and Wright, 2012.)

Under the estimated model parameters, supply shocks on the scale of the 2008 – 2012 LSAP programs result in a decline in the ten-year term premium of only about 20 basis points through the duration channel, compared to empirical findings in the literature of as much as 250 basis points for the overall effects of these programs. Since the models considered are rigged to generate duration effects that are, if anything, unrealistically *large*, this finding implies that the impact of LSAPs is probably not primarily due to “removing duration from the market,” as many policymakers and academics have suggested. That conclusion is broadly consistent with the results of research that has tried to disentangle the various channels of LSAPs empirically (e.g., Krishnamurthy and Vissing-Jorgensen, 2011). The primary reason is that, during the LSAP period, the short-term interest rate has been close to the ZLB. The truncation of the lower tail of the rate distribution reduces volatility along the entire yield curve, and—since duration effects depend on the product of duration and volatility—render this channel less effective, a point also made by Doh (2010). The reduction in interest-rate volatility is even greater when the central bank broadcasts its medium-term intentions for short-term interest rates through so-called “forward guidance.” Thus, forward guidance and LSAPs, two tools of monetary policy that are often viewed as complements at the ZLB, actually offset each other somewhat in terms of their effects on longer-term yields. The nonlinear solution method is particularly useful in exploring this result, as this type of situation—a short-rate process that is both nonlinear and heteroscedastic—is not amenable to analysis in an affine term-structure framework.

The methodological approach developed here can be viewed as a generalization of early work by Friedman (1980), Roley (1982), and Frankel (1985). These papers showed how to invert a portfolio-choice problem to solve for expected returns as a function of asset shares. (More recently, Piazzesi and Schneider, 2007, extended this basic approach to a consumption-based framework.) However, by assuming that the distribution of returns is unaffected by changes in portfolio weights and by placing no restrictions on returns across assets, these models violated rational-expectations and no-arbitrage conditions. Piazzesi and Schneider (2008) embed a similar portfolio-choice problem in an arbitrage-free framework, using survey data and vector autoregressions to generate estimates of the conditional distribution of future yields, and derive bond-demand relationships based on those estimates. While this method captures the state dependence of the return distribution and is clearly preferable to the assumption

that that distribution is constant, it still does not impose that investors know that future prices will be determined by the no-arbitrage relationships of the model and it may thus still lead to internal inconsistency under rational expectations.⁴

As noted above, Vayanos and Vila (2009) features both no-arbitrage and rational expectations. In that model, mean-variance arbitrageurs interact with “preferred-habitat” agents who have downward-sloping demand curves for assets of particular maturities. Assuming a linear-Gaussian factor structure, Vayanos and Vila obtain an affine solution in which demand shocks that remove long-term debt from the hands of arbitrageurs push down long-term interest rates. As the only model to incorporate supply effects into an otherwise standard representation of the term structure, Vayanos-Vila has been highly influential in the way that economists have designed and interpreted the recent empirical studies.⁵ But the presence of preferred-habitat agents clouds the issues somewhat when it comes to studying the effects of overall debt supply, since, as Greenwood and Vayanos (forthcoming) illustrate, prices depend on quantities in the model even if those agents are absent. The relevance of the preferred-habitat mechanism thus becomes an empirical question. Furthermore, Vayanos and Vila make several strong assumptions about functional forms. Although these assumptions succeed in delivering a familiar-looking affine term structure, it is also important to entertain more-general cases, such as those allowing for the ZLB. Since the Vayanos-Vila model cannot generally be solved in closed form anyway, little is lost in moving away from specifications whose chief appeal is analytical convenience.

The paper proceeds as follows. Section 2 lays out the portfolio balance framework, derives the basic asset-pricing equations, and discusses the solution algorithm. Section 3 considers the simplest possible version of these models, in which investors have myopic, mean-variance preferences and the short rate is the only source of stochastic variation, to illustrate the basic workings of portfolio balance and the duration channel. Section 4 shows how the duration effects in this model can be extended in various directions, and Section 5 explores ways of introducing imperfect substitutability among Treasuries. Section 6 estimates a three-factor model with time-varying inflation risk and stochastic Treasury supply, and Section 7 applies this model to examine the effects of the LSAP programs. Section 8 concludes.

2 The Portfolio-Balance Framework

I use the term “Portfolio Balance” to encompass a broad set of models in which the equilibrium prices of financial assets are related to the quantities of those assets that investors must hold and in which the quantities can be considered to have an exogenous or policy-dependent component. In particular, suppose

⁴ As Piazzesi and Schneider (2008) point out, if expectations are not rational investors may not believe that future asset prices will be determined by the same relationships that determine present prices. In this case, their use of survey evidence for these expectations may be more appropriate.

⁵ See, for example, Doh (2010), Hamilton and Wu (2011), and Li and Wei (2012).

there exists a finite number of assets N , with time- t prices $\mathbf{p}_t = (p_{1t} \dots, p_{Nt})$, and par values $\mathbf{X}_t \equiv (X_{1t} \dots, X_{Nt})$. The idea behind portfolio balance is that fluctuations in the state of the economy that change the value of \mathbf{X}_t will change \mathbf{p}_t because expected returns—and therefore current prices—must adjust to make investors willing to hold the outstanding supply of securities at each point in time. In the most straightforward cases, the supply of each security is treated as an exogenous quantity issued by a price-insensitive entity, such as the government. More generally, \mathbf{X}_t may itself depend on prices.

This section lays out the general properties of asset-pricing models in which portfolio-balance effects may arise and describes how to solve such models numerically. Specific examples of these models will be developed in the remainder of the paper.

2.1 General asset-pricing relations

Suppose the state of the economy in time t can be described by a vector of state variables \mathbf{s}_t . The state may include observable or unobservable factors, and it is assumed to have a transition density of Markov form $\tau(\mathbf{s}_t | \mathbf{s}_{t-1})$. I assume that τ is continuous almost everywhere and that the support of \mathbf{s}_t , \mathbf{S} , is bounded. Following Krishnamurthy and Vissing-Jorgensen (2012) and others, I allow asset returns to include a convenience yield to investors beyond their price appreciation. Specifically, the return on the n^{th} asset between periods t and $t+1$ is given by

$$R_{nt+1} \equiv \frac{q_{nt+1} + b_n(\mathbf{s}_t, \mathbf{p}_t)}{p_{nt}} \quad (1)$$

where q_{nt+1} is the asset's nominal payoff in period $t+1$ and b_n represents liquidity, safety, or other benefits derived from holding security n , denominated in dollars. It is assumed that these benefits accrue to the holders of the assets with perfect certainty. Collect the convenience yield in the vector $\mathbf{b}_t = (b_1(\mathbf{s}_t, \mathbf{p}_t) \dots, b_N(\mathbf{s}_t, \mathbf{p}_t))$. Of course, one special case to consider is $\mathbf{b}_t = \mathbf{0}$.

Absence of arbitrage implies that there is a unique stochastic discount factor (SDF) M_t that prices all assets in the economy:

$$E_t[M_{t+1}R_{nt+1}] = 1 \quad \forall n \quad (2)$$

where E_t indicates the expectation conditioned on the full set of information at time t . I consider cases in which M_{t+1} is determined by the state variables in periods t and $t+1$ and, critically, by the return on wealth:

$$M_{t+1} = M(\mathbf{s}_t, \mathbf{s}_{t+1}, R_{t+1}^w) \quad (3)$$

where R_{t+1}^w is the return on wealth between t and $t+1$. Pricing kernels that have this property include, for example, those generated by the CAPM and by Epstein-Zin-Weil preferences.

Denote the instantaneous risk-free rate as r_t^f . The above assumptions imply that the following relationship applies to the vector of time- t asset prices:

$$\mathbf{p}_t = \exp[-r_t^f] E_t[\mathbf{q}_{t+1}] + \text{cov}_t[M_{t+1}, \mathbf{q}_{t+1}] + \mathbf{b}_t \quad (4)$$

where \mathbf{q}_{t+1} is the vector of asset payoffs and $\text{cov}_t[M_{t+1}, \mathbf{q}_{t+1}]$ is the vector of conditional covariances between $t+1$ payoffs and the SDF. If payoffs are known with certainty and $\mathbf{b}_t = \mathbf{0}$, then current prices are just future payoffs discounted at the risk-free rate. Outside of this case, equation (4) is not a closed-form solution for prices because, in general, M_{t+1} and \mathbf{b}_t can depend on \mathbf{p}_t .

Time- t wealth is given by $W_t = \mathbf{p}_t' \mathbf{X}_t$. For convenience, define the normalized par shares \mathbf{x}_t , where the n^{th} share $x_{nt} \equiv X_n(\mathbf{s}_t) / \mathbf{i}' \mathbf{X}_t$ and \mathbf{i} is a vector of 1's. The return on wealth can be written as

$$R_{t+1}^w = \frac{(\mathbf{q}_{t+1} + \mathbf{b}_t)' \mathbf{x}_t}{\mathbf{p}_t' \mathbf{x}_t} \quad (5)$$

It is through this equation that quantities can matter for asset prices. The equilibrium price vector must adjust such that the return on wealth and the return on any given asset satisfy both equation (2) and equation (5). Conditional on the payoff vector $\mathbf{q}_{t+1} + \mathbf{b}_t$, clearly different values of \mathbf{x}_t will be consistent with different equilibrium values of \mathbf{p}_t . However, one implication of equation (5) is that, so long as $b_{nt} = 0$, the quantity of asset n has no *direct* effect on the price of asset n ; only weighted sums of *all* of the x_{nt} 's appear in the pricing formula.⁶

Equations (2), (3), and (5) together represent the equilibrium conditions for no-arbitrage portfolio balance. These conditions are not by themselves sufficient to solve for the time- t prices \mathbf{p}_t because they have nothing to say about the distribution of asset payoffs, but in the case of assets that have a known terminal payoff at some date in the future prices can in principle be solved for recursively. In particular, if we impose the additional restrictions $q_{1t} = 1$ and $q_{n-t} = p_{n-t}$ for $n > 0$ we obtain the case of zero-coupon default-free bonds, where n now evidently indicates the bond's maturity at time t . (In this case, $p_{1t} = -\log r_t^f$ is the price of the one-period risk-free bond.) Perfect certainty (and $\mathbf{b}_t = \mathbf{0}$) corresponds to the "strong form" of the expectations hypothesis—long-term interest rates are the average of expected short-term interest rates (plus an adjustment reflecting Jensen's inequality). Otherwise, the covariance term in equation (4) adds a term premium to the yield curve. Even restricting attention to zero-coupon bonds, it is still not possible to solve analytically for \mathbf{p}_t in general, but the solution method described below works for most cases of interest.

⁶ Also note that only the relative quantities \mathbf{x}_t matter for prices, not the absolute amounts \mathbf{X}_t .

2.2 Solution method

Models satisfying the above conditions can be solved numerically for the time- t vector of asset prices \mathbf{p}_t using the following successive-approximation method. I first make explicit that prices and quantities depend on the state of the economy. Namely, we seek a vector-valued function $\mathbf{p}(\mathbf{s}_t)$ that describes how all asset prices depend on the state vector \mathbf{s}_t . It is assumed that the form of the pricing kernel in equation (3), the laws of motion for the states, and the dependence of quantities on the states are known—that is, we (and investors) have knowledge of the functions $\tau(\mathbf{s}_t | \mathbf{s}_{t-1})$, $M(\mathbf{s}_t, \mathbf{s}_{t+1}, R_t^w)$, $\mathbf{b}(\mathbf{s}_t)$, and $\mathbf{x}(\mathbf{s}_t)$.

Let $\mathbf{p}_d^i(\mathbf{s}_t)$ be a proposal for the pricing function on a discretization of the state space $d = (\mathbf{d}_1, \dots, \mathbf{d}_G) \in \mathbf{S}^G$, where G is the number of nodes and $i = 0, \dots, I$ indexes iterations. $\mathbf{p}_d^i(\mathbf{s}_t)$ and $\tau(\mathbf{s}_{t+1} | \mathbf{s}_t)$ are used to generate a proposal for the joint distribution of period $t+1$ states and prices. That distribution, in turn, generates an updated pricing function $\mathbf{p}_d^{i+1}(\mathbf{s}_t)$ through the solution to equation (2)—that is, by solving the system

$$\mathbf{p}_d^{i+1}(\mathbf{s}_t) = E_t \left[\left(\mathbf{q}_d^i(\mathbf{s}_{t+1}) + \mathbf{b}(\mathbf{s}_t) \right) M \left(\mathbf{s}_t, \mathbf{s}_{t+1}, \frac{\left(\mathbf{q}_d^i(\mathbf{s}_{t+1}) + \mathbf{b}(\mathbf{s}_t) \right)' \mathbf{x}(\mathbf{s}_t)}{\mathbf{p}_d^i(\mathbf{s}_t)' \mathbf{x}(\mathbf{s}_t)} \right) \right] \quad (6)$$

for the function $\mathbf{p}_d^{i+1}(\mathbf{s}_t)$ at each value $\mathbf{s}_t = \mathbf{d}_1, \dots, \mathbf{d}_G$ where

$$\mathbf{q}_d^i(\mathbf{s}_t) \equiv \left(1 \quad p_{d_1}^i(\mathbf{s}_t) \quad \dots \quad p_{d_{N-1}}^i(\mathbf{s}_t) \right) \quad (7)$$

Specifically, let the nodes be uniformly distributed over the state space, and approximate the conditional transition density from node j to node h as

$$\hat{\tau}(\mathbf{d}_h | \mathbf{d}_j) = \tau(\mathbf{d}_h | \mathbf{d}_j) \left[\sum_{g=1}^G \tau(\mathbf{d}_g | \mathbf{d}_j) \right]^{-1} \quad (8)$$

Then, at iteration $i+1$ of the algorithm, the price of an n -maturity bond when the state is at node j is calculated as

$$p_{dn}^{i+1}(\mathbf{d}_j) = \sum_{g=1}^G \hat{\tau}(\mathbf{d}_g | \mathbf{d}_j) \left(p_{dn-1}^i(\mathbf{d}_g) + b_n(\mathbf{d}_j) \right) M \left(\mathbf{d}_j, \mathbf{d}_g, \frac{\left(\mathbf{q}_d^i(\mathbf{d}_g) + \mathbf{b}(\mathbf{d}_j) \right)' \mathbf{x}(\mathbf{d}_j)}{\mathbf{p}_d^i(\mathbf{d}_j)' \mathbf{x}(\mathbf{d}_j)} \right) \quad (9)$$

This procedure converges in G and I , so long as the moments of the pricing kernel are well behaved. As long as $\mathbf{b} = 0$, for example, sufficient conditions for equation (9) to constitute a contraction mapping on p_{dn} are that (i) the short-term interest rate is always non-negative and (ii) $M(\cdot)$ is linear in R_t^w . (These two conditions will be satisfied by all of the models considered in this paper.) The Banach Theorem then

guarantees $\mathbf{p}_d^i(\mathbf{s}_t) \rightarrow \mathbf{p}_d(\mathbf{s}_t)$ for any given discretization d , where \mathbf{p}_d is the (unique) pricing function that obtains if $\hat{\tau}$ is the data-generating process. But continuity of τ ensures that, for any node j ,

$$\lim_{G \rightarrow \infty} P_{dn}(\mathbf{d}_j) = E_t \left[P_{n-1}(\mathbf{s}_{t+1}) M \left(\mathbf{d}_j, \mathbf{s}_{t+1}, \frac{\mathbf{q}(\mathbf{s}_{t+1})' \mathbf{x}(\mathbf{d}_j)}{\mathbf{p}(\mathbf{d}_j)' \mathbf{x}(\mathbf{d}_j)} \right) \right] \quad (10)$$

i.e., in the limit, the pricing function solves the no-arbitrage condition (2). In the more general case of non-zero \mathbf{b} , a similar argument applies, but convergence can also be achieved with negative rates (since there is nothing contradictory about negative rates when securities are paying unobserved benefits).

Finally, by construction, if the algorithm converges, any point of convergence is a rational-expectations equilibrium. This follows immediately, since convergence is *defined as* the fixed point at which the joint distribution of \mathbf{p} and M is consistent with the vector \mathbf{p} , for each point in the state space. Suggestive results on the speed of convergence are discussed in Section 3.3.

It is important to note that, although the algorithm only solves for the vector of prices at G points in the state space, once these solutions are in hand it is straightforward to calculate equilibrium prices at *any* point. In particular, take an arbitrary state value \mathbf{s}_t . For G large enough, we have

$$P_n(\mathbf{s}_t) \approx \left[\sum_{g=1}^G \tau(\mathbf{d}_g | \mathbf{s}_t) \left(P_{dn-1}(\mathbf{d}_g) + b_n(\mathbf{d}_g) \right) M \left(\mathbf{s}_t, \mathbf{d}_g, \frac{(\mathbf{q}_d(\mathbf{d}_g) + \mathbf{b}(\mathbf{s}_t))' \mathbf{x}(\mathbf{s}_t)}{\mathbf{p}_d(\mathbf{s}_t)' \mathbf{x}(\mathbf{s}_t)} \right) \right] \left[\sum_{g=1}^G \tau(\mathbf{d}_g | \mathbf{s}_t) \right]^{-1} \quad (11)$$

Once the algorithm has converged, the quantities on the right-hand side are all known. Thus, securities can be priced in at any point in \mathbf{S} .

3 Duration Effects in a Portfolio-Choice Model

To apply the above approach, I consider the case of Markowitz (1952) portfolio choice in which the spectrum Treasury bonds are the only assets available to investors. Of course, this is just one of many possible implementations of the framework laid out above. From an expositional standpoint, one virtue of this choice is its simplicity—mean-variance preferences over wealth are the easiest way to achieve a pricing kernel involving the market return as is required by equation (3). It thus makes it straightforward to analyze portfolio-balance effects. Partly for this reason, these assumptions have been used by most other papers in this literature, including Frankel (1985) and Vayanos and Vila (2009). As will be seen, the model presented below generally matches the data reasonably well despite its simplicity.

Another virtue of this framework, at least from the perspective of assessing policy, is that it effectively allows portfolio-balance effects to achieve their greatest possible strength. This is because it implicitly assumes that the Treasury market is completely isolated both from other asset markets and from

agents' consumption and savings decisions. This means, for example, that movements in Treasury duration outstanding affect the duration of investors' portfolios one-for-one (rather than being potentially damped by the presence of other assets in those portfolios). It also means that Wallace neutrality is assumed away. Thus, these assumptions effectively put an upper bound on how large portfolio-balance effects could be.

To focus on duration effects, I assume $\mathbf{b}_t = \mathbf{0}$ for the moment; this assumption will be relaxed in Section 5.

3.1 Equilibrium conditions

Suppose agents solve the portfolio-choice problem

$$\max_{\mathbf{w}_t} E_t[\mathbf{R}_{t+1} \mathbf{w}_t] - \frac{a_t}{2} \text{var}_t[\mathbf{R}_{t+1} \mathbf{w}_t] \quad (12)$$

where a_t is the coefficient of relative risk aversion at time t , and \mathbf{w}_t is the vector of dollar values allocated to each asset. The maximization is subject to $\mathbf{w}_t' \mathbf{i} \leq W_t$, where by definition $w_{nt} = p_{nt} X_{nt}$. It produces the standard first-order condition,

$$\frac{1}{W_t} \mathbf{w}_t = \frac{1}{a_t} \boldsymbol{\Sigma}_t^{-1} E_t[\mathbf{R}_{t+1}] \quad (13)$$

where $\boldsymbol{\Sigma}_t$ is the conditional covariance matrix of returns. Rearranging yields a system of quadratic equations in \mathbf{p}_t :

$$\mathbf{p}_t = \exp\left[-r_t^f \left(E_t[\mathbf{q}_{t+1}] - a_t \frac{\boldsymbol{\Omega}_t \mathbf{x}_t}{\mathbf{p}_t' \mathbf{x}_t} \right)\right] \quad (14)$$

where $\boldsymbol{\Omega}_t$ is the conditional covariance matrix of the asset payoffs \mathbf{q}_{t+1} . This is the equivalent of equation (4) for this model. Let $\omega_t^{m,n}$ be the m, n^{th} element of $\boldsymbol{\Omega}_t$. Since the first asset is the risk free bond, we have $\omega_t^{m,1} = \omega_t^{1,n} = 0$ for all m, n . Note that, in the general case, all other objects in equation (14) may be considered time-varying.

In the previous section it was pointed out that, if $\boldsymbol{\Omega}_t = \mathbf{0}$, then current prices are just future prices discounted at the rate r_t^f , regardless of the form of $M(\cdot)$. In the particular model considered here, equation (14) makes clear that risk neutrality, $a_t = 0$, produces the same result, regardless of the volatility of prices. Furthermore, note that this result, $p_{nt} = p_{1t} E_t[p_{nt+1}]$, also occurs under risk aversion in this model if the market portfolio consists only of the risk-free asset. This follows immediately from substituting $\mathbf{x}_t =$

(1 0 ... 0) into equation (14). On the other hand, as long as $a_t > 0$ and $x_1 < 1$, we have $p_{nt} < p_{1t}E_t[p_{nt+1}]$, for $n > 1$, so that all risky assets require a discount and all risky returns include a risk premium. Generally speaking, increasing the share of risky assets that investors have to hold lowers the price of those assets in period t .

Since (14) is the result of a Markowitz portfolio problem it has a conditional CAPM representation. Specifically, since $E_t[R_t^w] = E_t[\mathbf{q}_{t+1}]'\mathbf{x}_t / \mathbf{p}_t'\mathbf{x}_t$ is the expected return on wealth, we can write

$$E_t[\mathbf{R}_{t+1}] - R_{1t+1}\mathbf{i} = \boldsymbol{\beta}_t \left(E_t[R_t^w] - R_{1t+1} \right) \quad (15)$$

where $\boldsymbol{\beta}_t = \boldsymbol{\Omega}_t\mathbf{x}_t/\mathbf{x}_t'\boldsymbol{\Omega}_t\mathbf{x}_t$ and $R_{1t+1} = 1/p_{1t}$ is the return on a one-period, risk-free bond. The expected excess return on the market portfolio is $E_t[R_t^w] - R_{1t+1} = a_t \mathbf{x}_t'\boldsymbol{\Omega}_t\mathbf{x}_t$. From this, one can recover the SDF, which is conditionally linear in the market return:

$$M_{t+1} = \theta_{0t} + \theta_{1t}R_{t+1}^w \quad (16)$$

where $\theta_{0t} = a_t \exp[-r_t^f](1 + a_t^2 \mathbf{x}_t'\boldsymbol{\Omega}_t\mathbf{x}_t)$ and $\theta_{1t} = -a_t \exp[-r_t^f]$.⁷

In almost all versions of this model, the short rate r_t^f must be considered stochastic (otherwise, there is no risk and no term premium). If r_t^f is the *only* source of variation, there is a single “market price of risk” given by the Sharpe ratio

$$\lambda_t = \frac{E_t[R_{nt+1}] - R_{1t+1}}{\sigma_{nt}} \quad \forall n \quad (17)$$

where σ_{nt} is the standard deviation of the return on asset n . This, in turn, implies that all asset prices and returns are perfectly conditionally correlated, regardless of the value of \mathbf{x} or the dynamics of r_t^f . In particular, we can write $R_{nt} = A_{nt} + B_{nt} p_{1t+1}$ for any n , where $B_{nt} \equiv \sigma_{nt} / \sqrt{\omega_t^{2,2}}$ and $A_{nt} \equiv (1 - p_{2t} B_{nt})/p_{1t}$. (Recall that $\omega_t^{2,2}$ is the conditional variance of the one-period bond price, which is an exogenous quantity.) The coefficient B_{nt} represents the response of the return on a risky asset to the risk-free rate—the analogue of the “spot-rate duration” that appears in continuous-time term-structure models.⁸

⁷ In equation (16), the intercept also depends on quantities through the expected-return term $a_t\mathbf{x}_t'\boldsymbol{\Omega}_t\mathbf{x}_t$. However, this term has no asset-pricing consequences, since it cancels in the expectation of M_{t+1} and since it does not affect the covariance of M_{t+1} with other prices.

⁸ Prices in this model are perfectly correlated in levels. In contrast, in one-factor affine term structure models in continuous time (e.g., Vasicek, 1977), prices are perfectly correlated in logs. One can generate log correlation in the model here by (somewhat awkwardly) assuming investor preferences over instantaneous, rather than cumulative, returns in equation (12). For typical levels of yields, this choice makes little quantitative difference, and both models converge to the continuous-time case as the period length goes to zero.

The standard deviation of the return on the market portfolio can be expressed as $\sigma_t^w = \sqrt{\omega_t^{2,2} \mathbf{x}'\mathbf{B}_t}$, where \mathbf{B}_t is the vector of B_m coefficients. Combining this with equation (17), we have

$$\lambda_t = a\sqrt{\omega_t^{2,2} \mathbf{x}'\mathbf{B}_t} \quad (18)$$

analogous to a result produced by Vayanos and Vila (2009). Again, the individual asset share x_n does not matter for the individual asset price p_m . Only the weighted sum of asset shares $\mathbf{x}'\mathbf{B}_t$ is relevant, and it affects all prices in the same way. Consequently, this model will not produce local-supply effects from large quantity gluts or shortages in particular sectors of the market. Reducing the quantity $\mathbf{x}'\mathbf{B}_t$ is the model's version of "removing duration from the market." Doing so will lower both the total risk of the Treasury portfolio and the price of that risk. Note that the above exposition has made no assumptions about the process determining the short rate r_t^f other than that it is well-behaved enough that the first and second moments of p_{1t+1} exist.⁹

Comments on wealth and equilibrium

The pricing kernel that is solved for in the algorithm determines all asset prices in the economy, but it is only the quantities of financial assets that are in positive supply in the aggregate that affect the pricing kernel. Thus, one need not account for assets created by the private sector—the representative investor is short exactly as much of these securities as he is long. The extent to which government debt is "special," so that changes in its distribution matter for asset prices, depends on the degree to which it can be considered to be in positive net supply. This, in turn, depends on whether investors fully internalize the risks of future tax payments—that is, on whether Wallace (1981) neutrality holds. Violations of this condition can be motivated by market segmentation, investor inattention, or the government possessing a special revenue-generating technology. The model presented here is partial-equilibrium because (and only because) it does not specify how the government raises revenue to pay off its maturing debt, and it thus violates Wallace neutrality by assumption.

A related critique of portfolio-balance mechanisms, articulated by Eggertsson and Woodford (2003), does not presume frictionless markets as Wallace neutrality does but applies even to segmented markets so long as the pricing kernel depends only on consumption. Specifically, Eggertsson and Woodford describe the logic of their "neutrality proposition" as follows. In their model, "the marginal utility to the representative household of additional income in a given state of the world depends on the household's consumption in that state, not on the aggregate payoff of its asset portfolio in that state. And changes in

⁹ Also note that, technically, there is nothing in the development of this section that is specific to the Treasury market, because other than for the one-period bond the derivations have not made use of the restriction that $q_{n-1t} = p_{n-1t}$.

the composition of the securities in the hands of the public *don't* change the state-contingent consumption of the representative household..." [Emphasis in original.] This result follows from their assumption that preferences are time-separable and do not directly involve wealth. A consequence is that the pricing kernel (3) does not include the term R_t^w , which is why there are no portfolio-balance effects.

In contrast, in the more general case in which wealth is allowed to enter the pricing kernel, marginal utility does depend, at least in part, on the payoff of the asset portfolio, and changes in the composition of securities in the hands of the public *do* change this payoff. In these cases, portfolio-balance effects are not just a possibility—they are necessary for market clearing. (Of course, this is not to say that they are necessarily quantitatively important.) A key example is recursive preferences of the Epstein-Zin (1989) and Weil (1990) type, in which the pricing kernel can be written as a function of both consumption growth and the return on wealth. Recursive-preference models are beyond the scope of the present paper, but the results below ought to be suggestive of the kinds of results that could obtain in such models.¹⁰ Generally speaking, adding consumption to the model here would dampen—and possibly even reverse the sign of—duration effects by breaking the tight link between asset payoffs and marginal utility.

A final point to note is that, in the models considered here, I assume that default-free bonds are the only assets available to the representative investor. Strictly speaking, “wealth” should include nonfinancial assets, such as physical and human capital, that are in positive net supply, a critique that has long been leveled at empirical implementations of portfolio-based models (Roll, 1977). It would be interesting to incorporate such factors into the types of models presented below—and the solution algorithm described above would still apply to such models—but that is also beyond the scope of the present paper. Effectively, I consider only exchange-economy models here. Again, this will err on the side of attributing too much importance to duration effects that arise through government debt supply.

3.2 Calibration of the baseline one-factor model

I begin by considering the case in which the short-term interest rate is the only source of stochastic variation. Since the simple model in this section is only for exposition, I calibrate to loosely match the data. Parameters will be formally estimated in the multi-factor model developed in Section 6.

I take periods to be one year in length and consider $N = 30$ bonds. I assume that r_t^f follows an AR(1), calibrated to match the dynamics of the one-year rate in the U.S. data between December 1971 and December 2012. I also impose the zero lower bound. Specifically, let

$$r_t^f = -0.00041 + .964r_{t-1}^f + \varepsilon_t^r \quad (19)$$

subject to

¹⁰ See Piazzesi and Schneider (2007) for some related work along these lines.

$$\Pr[r_t^f < 0] = 0 \quad \forall t \quad (20)$$

where the standard deviation of ε_t^r is $\sigma^r = 0.017$. Equation (19) was estimated on the time series of one-year Treasury rates by maximum likelihood, using a truncated normal distribution for ε_t^r , where the truncation point varies in each period to enforce (20).¹¹ Figure 1 shows how the distributional forecast of r_{t+1}^f changes with r_t^f under this specification. The nonlinear process for the short rate captures the stylized facts of reduced volatility and increased skewness near the ZLB, while ensuring that investors always put zero probability on negative rates. It has the feature that, when the short rate is zero, its modal forecast for the following year is also zero. (To a degree, this captures the persistence of the ZLB outcome that has been observed since 2008.) Unlike some other specifications that restrict interest rates to be non-negative (e.g., Cox et al., 1985), it does not require the conditional variance of r_t^f to be zero at the ZLB.

To solve the model, one also needs to know the maturity distribution given by the relative quantities of each security, $\mathbf{x} = \{x_1, \dots, x_N\}$. In principle, one can specify this distribution in any way, so long as the shares are non-negative and sum to unity. For parsimony, I approximate the maturity structure of government debt with an exponential distribution:

$$x_n \propto \exp[-n/z] \quad (21)$$

The mean of the distribution z is calibrated to 2.7, which matches the average duration of U.S. government liabilities in the hands of the public since 1971, shown in the top panel of Figure 2.¹² The exponential approximation is likely to be a good one because the distribution of government liabilities typically consists overwhelmingly of shorter maturities. For example, as shown in the middle panel, liabilities with less than five years of duration (including currency and reserves) account, on average, for about 80 percent of total face value. The exponential functional form implies that, as in the data, the

¹¹ The Treasury yield data used in this paper are the Svensson (1994) zero-coupon yields discussed in Gurkaynak et al. (2007) and maintained on the Federal Reserve Board's website. Yields through 15-year maturities are available since 1971 (the date used for most of the exercises in this paper), and the 30-year yield becomes available in 1985. The parameters of (19) were estimated using nonoverlapping annual samples, with the log-likelihood given by

$$\log L = \sum_{t=2}^T \left(\frac{1}{\sigma^2} (r_t^f - \phi_0 - \phi_1 r_{t-1}^f)^2 - \log \int_0^{\infty} \exp \left[\frac{1}{\sigma^2} (r_t^f - \phi_0 - \phi_1 r_{t-1}^f)^2 \right] dr_t^f \right)$$

¹² All bonds in the model are zero-coupon, so duration and maturity are identical. When calibrating the model to the data, I focus on duration because this is the more appropriate measure of interest-rate risk that investors bear. Government liabilities in the hands of the public are computed as all Treasury securities outstanding, less the amount held by the Federal Reserve and other government entities, plus the monetary base. Data on Treasury securities outstanding are taken from CRSP[®] (Center for Research in Security Prices, Booth School of Business, The University of Chicago. Used with permission. All rights reserved. crsp.uchicago.edu.) CRSP reports the amount of each coupon security in public hands. I use Federal Reserve balance sheet data to calculate the corresponding quantity of bills, and I assume that the average duration of bills held by the Fed is equal to the average duration of all bills outstanding. The monetary base is assumed to have zero duration.

distribution is highly skewed toward securities with very low duration. Moreover, as suggested by equation (18) and illustrated further below, the exact shape of the distribution beyond its first moment is actually of little importance.

Finally, given the above specifications for r_t^f and \mathbf{x} , I calibrate the risk-aversion parameter a to match the average spread between 1- and 15-year yields in the data, which is 148 basis points. The 15-year yield is chosen as the reference point both because it is the midpoint of the 30-year yield curves that I will focus on and because it is the longest maturity for which yield data exist over the entirety of the sample. The calibration gives a value of $a = 12.8$.

3.3 Results of the baseline model

Convergence

Figure 3 displays some results on the convergence of the solution algorithm for this model. The top panel shows the computed 5-, 10-, 15-, and 30-year yields, shown for a short rate at its average value of 5.8%, across the first 50 iterations ($i = 1, \dots, 50$). The algorithm is initialized at a price vector $\mathbf{p}^0(\mathbf{s}_i) = (1, \dots, 1)$ for all values of \mathbf{s} , and uses $G = 65$ gridpoints across the state space. It is evident from this figure that, for each maturity n , the solution converges very quickly once $i > n$.

The middle panel shows convergence in the number of gridpoints by displaying the computed yield curve (after $I = 50$ iterations), again using $R_{1t} = 0.058$ for illustration. Yield curves are shown for $G = 5, 9, 17, 33,$ and 65 , in each case spaced equally across possible values of p_{1t} . The space is assumed bounded between 0 and 0.80, so this partitioning corresponds roughly to increments of between 38 basis points and 5 percentage points. While 5 nodes is clearly too few to achieve convergence, the solutions using 17 or more nodes are indistinguishable from each other.

For brevity, these results were shown for the average value of the short rate. Similar convergence results obtain for other points in the state space, although solutions will not be accurate near the bounds if the underlying state process itself is not actually bounded. For example, in the above case, we would not expect the procedure to generate correct solutions near $p_{1t} = 0.80$ (corresponding to a risk-free return of 25%). However, so long as the bound on the state space is imposed far enough away from the values of the states that are actually realized in practice, this limitation has a negligible effect on the results. The bottom panel of the figure illustrates this claim by comparing the yield curve computed above with the yield curve computed when the grid for p_{1t} is extended over the entire range $(0, 1)$, again at the average value of the short rate. (The latter computation used $G = 100$.) The two curves are virtually identical, differing by less than 1 basis point across maturities.

The above results suggest that convergence of the algorithm is fast in both the i and g dimensions and that results are reliable with a reasonable truncation of the theoretically admissible state space. With $I =$

50, and $G = 65$, the computation takes only seconds on a standard desktop computer. Although the curse of dimensionality kicks in when additional state variables are added, it remains practical to extend the approach to a few additional dimensions, as I will later.

Moments and Comparative Statics

The model developed above captures the duration channel. Indeed, the term premium in this model reflects *only* compensation for bearing duration risk. Figure 4 demonstrates how the duration of the Treasury portfolio affects the behavior of asset prices in the model by plotting the equilibrium values of the spot-rate duration B_{nt} and the price of risk λ_t across possible values of r_t^f , n , and z . (These are straightforward to compute using the formulas of the previous section, because they are functions only of the conditional moments of prices, which are generated by the solution algorithm.) Because r_t^f follows an AR(1) process away from the ZLB, the function B_{nt} has a similar form across n at each value of the short rate as in the Vasicek (1977) model. When r_t^f is near the ZLB, B_{nt} and λ_t are uniformly lower in magnitude because of the lower risk that prevails in that region of the state space. Decreasing the value of z used in the model also reduces risk exposures and risk prices across the term structure.

The blue lines in Figure 5 show the first two moments of the yield curve generated by the model, inputting actual monthly data on the one-year rate and assuming a constant duration of 2.7 years for the supply distribution. The sample moments of the data are shown in green for comparison. I perform the calculations over both the full 1971 – 2012 sample, and over the shorter 1985 – 2012 sample where the rate data are available for maturities beyond 15 years. The model does well at accounting for the average shape of the yield curve up through maturities of about 15 years (this is partly by construction), but it does not predict enough curvature at the long end of the curve. Furthermore, it underpredicts volatility, especially for longer-term yields. These results are not surprising since, apart from the ZLB, the model is essentially a Vasicek model, which is known to have trouble matching the data in these dimensions. The augmented versions of the model developed below will help to address these deficiencies, but this simple case is close enough to the data to begin to get a sense of the quantitative effects of changes in supply.

Specifically, consider an LSAP-type experiment that shifts the mean of the exponential Treasury supply distribution z from the level of 2.7 used above to a level of 2.0. The change of -0.7 years is approximately equivalent to the effects of the cumulative asset purchases by the Fed through 2012 (the first and second round of LSAPs and the Maturity Extension Program).¹³ The difference in the maturity distribution of securities is illustrated by the red bars in the top panel of Figure 6. The effect on the yield curve, evaluated at the sample average value of the short rate, is shown in the middle panel. Under the calibration used here, the ten-year yield is 56 basis points lower under the smaller-duration distribution.

¹³ The computation of the -0.7 value is explained in Section 7.

Again, this is due to reductions both in the price of risk (as was illustrated in the bottom panel of Figure 4) and in risk itself. The latter is shown in the bottom panel of Figure 6, which plots the standard deviation of one-year returns (the square-root of the diagonal of Σ_t , again conditional on $R_{1,t+1} = 0.058$) under the two supply distributions.

Now consider the case in which the short rate is at its zero lower bound. As noted, the principal consequence is a reduction in interest-rate volatility, particularly for short- to medium-term securities, as the lower tail of the conditional distribution of future interest rates is truncated. As a result, interest-rate risk is generally lower, and consequently the effects of changes in the distribution of Treasury supply are smaller. As shown by the solid lines in Figure 7, the same LSAP-type experiment that caused a 56 basis point reduction when the short rate was at 5.8% causes only a 29 basis point reduction when the short rate is at 0%.

The reduction in duration effects becomes even more pronounced if the short rate is anticipated to stay near zero with a high degree of certainty for multiple periods, because this results in a larger reduction in the variance of rates across the term structure. This type of thought experiment is relevant because the forward-guidance language that has been included in most FOMC statements since 2008 is widely viewed as having essentially committed the Fed to keeping rates near zero for a significant time into the future. To illustrate, the dashed lines in the figure show the model's outcomes when r_t^f is constrained to equal 0 for both periods t and $t+1$ (and investors know this with certainty at time t) and only then follows its usual AR(1) process. The forward guidance itself has a large effect on long-term rates; in this example, it reduces the 10-year yield by 52 basis points (the difference between the solid and dashed blue lines). However, with the forward guidance in place, the effect of the reduction in duration on the ten-year yield is only 23 basis points, rather than 29. The reason is that, as illustrated in the bottom panel, forward guidance in the one-factor model implies that the returns on all bonds in period $t+1$ are known with certainty. Consequently, the risk premium for the period over which the forward guidance is in place is zero, regardless of the amount of duration outstanding—i.e., $\omega_t^{2,2}$ is zero, and so, by equation (18), the price of risk λ_t is zero.¹⁴ Thus, although they both serve to reduce longer-term yields, forward guidance and asset purchases are somewhat offsetting, at least if duration effects are the only channel through which the supply distribution matters for prices.

¹⁴ This result does not imply that the *term premium* is zero in the presence of forward guidance, because investors still know that risk prices will be positive in future periods, and they factor that into the time- t prices of long-term bonds. Also note that the zero variance of bond returns follows from the fact that, here, the short rate is the only source of variation; thus when there is no uncertainty about r_t^f there is no uncertainty about anything. Bond-return variance could be positive under forward guidance in models in which returns depend on factors other than r_t^f .

4 Modifications to the Supply Distribution

This section takes a brief tour of possibilities for modifying the supply distribution \mathbf{x} . Namely, I consider (1) shapes of the distribution other than the exponential form used above, (2) endogeneity of the distribution in the Vayanos-Vila preferred-habitat model, (3) shocks to the distribution that are perfectly anticipated in advance, and (4) allowing the distribution to be stochastic over time. There are a variety of possibilities to consider with respect to each of these subjects, and for brevity I will confine the discussion to a few illustrative examples. A common theme will be that changing the properties of the supply distribution beyond its first moment does not have substantial effects on the average levels of yields. Although such adjustments can have effects on the second-order properties of the yield curve, at least in the examples here they do not enhance the importance of duration shocks.

4.1 Shape

As noted earlier, any individual value of x_n has little influence on the price of security n or on the price of any other security in the one-factor model—only weighted averages of the x 's are important. To illustrate this, consider a stark alternative to the exponential shape used for the maturity distribution above. In particular, suppose that the distribution was nearly degenerate around the same mean of $z = 2.7$, as shown in the top panel of Figure 8. This adjustment does slightly increase the duration risk of the portfolio, primarily because it completely eliminates the risk-free one-year bonds from investors' portfolios. However, after recalibrating a to again match the average slope of the empirical yield curve, there is virtually no difference between the term structures generated under the two distributions. This is illustrated in the bottom panel of the figure, where the solid blue line is the same yield curve shown in Figure 6, and the dashed red line is the (recalibrated) yield curve at the same value of the short rate under the degenerate supply distribution. Similar results obtain for other possible choices of the shape of \mathbf{x} . Since the shape of the supply distribution has little bearing on the shape of the yield curve, at least in the simple model considered here, we need not be overly concerned with the quality of the exponential approximation. Furthermore, modifications to the model that cause shifts in the shape of the distribution, such as the preferred-habitat extension considered next, are unlikely to have first-order effects, so long as they leave the mean unchanged.

4.2 Preferred habitat

In the model considered above, the supply of debt that must be held by investors is exogenous. More generally, this distribution itself could depend on interest rates, as in the model of Vayanos and Vila (2009). Apart from its affine structure, the essential feature of Vayanos-Vila that differs from the model considered above is that investors (“arbitraguers” in their terminology) face an elastic supply curve at

each maturity. This supply curve is assumed to arise from the presence of preferred habitat agents, each of whom deals in debt of only one maturity. Specifically, for each bond, preferred-habitat agents demand a (market-value) quantity that is a function h of maturity and price:

$$p_{nt}\xi_{nt} = h(n, p_{nt}) \quad (22)$$

where ξ_{nt} is the par value demanded. The par value left in the hands of the investors is simply $x_n - \xi_{nt}$, up to the constraints that no one can hold negative quantities.¹⁵ The model above effectively assumed $h(n, p_{nt}) = 0$ for all n . Equation (22) can be substituted directly into (5) to give the return on the investors' portfolio:

$$R_t^w = \frac{\sum_{n=1}^N \bar{x}_n(p_{nt}) p_{n-1t+1}}{\sum_{n=1}^N \bar{x}_n(p_{nt}) p_{nt}} \quad (23)$$

where

$$\bar{x}_n(p_{nt}) \equiv \min \left[x_n, \max \left[0, x_n - \frac{h(n, p_{nt})}{p_{nt}} \right] \right] \quad (24)$$

In Vayanos-Vila, the form of h is rather flexible, although they must restrict it somewhat to ensure an affine solution. The procedure here is not bound by this requirement. Nonetheless, I use a two-parameter specification similar to one that they adopt in an example:

$$h(n, p_{nt}) = \alpha \exp[\beta n] (p_{nt}^{-1/n} - 1) \quad (25)$$

with $\alpha > 0$. This form allows preferred-habitat demand for each bond to vary with its yield and for the slope of this demand curve to depend on the bond's maturity, while ensuring that the slope is always positive (i.e., for each maturity, a higher yield means greater preferred-habitat demand). The amount demanded when yields are zero is assumed to be zero, but this is without loss of generality since we have also complete freedom in choosing each x_n , and it is only \bar{x}_n that matters for prices.

Even with the three free parameters, α , β , and a , the range of shapes that the yield curve can take in the preferred-habitat model is severely restricted. For example, it is not generally possible to choose parameter values to match arbitrary values for both the average slope and average curvature. Therefore, for calibration purposes, I exploit the distinction between arbitrageurs and preferred-habitat agents in the

¹⁵ In Vayanos-Vila, assets are in zero net supply, and the amounts held by the arbitrageurs can be arbitrarily large or negative. Here I assume that only the government can create default-free debt, so that investors' share is limited to be between 0 and 1.

model and the available data on the relative quantities held by these types of agents in practice. Specifically, from the Flow of Funds data, I calculate the fraction of domestically, publicly held Treasuries that is owned by the broad classes of investors who are most likely to actively participate and attempt to exploit arbitrage opportunities. These categories are: households (because, by the construction of the Flow of Funds accounts, this designation includes hedge funds and private equity investors), broker-dealers, mutual funds, money-market funds, and depository institutions. The share of the market owned by these investors, shown in the bottom panel of Figure 2, is 55% on average since 1971 with an annual standard deviation of 12%. The α and β parameters are calibrated such that the model-produced equilibrium value of the ratio $\bar{x}_{nt} / \sum_n x_{nt}$ matches these moments. I continue to calibrate a to hit the average value of the 15-year slope. These criteria are achieved at $a = 36.8$, $\alpha = 0.53$, and $\beta = 0.19$.

Figure 9 compares the results of this model to the baseline case without preferred habitat. The solid blue line in the top panel is the yield curve under the baseline, evaluated at an average duration of 2.7 years and a short rate of 5.8% (the same as in Figure 6). The dotted blue line is the yield curve under the same conditions but including the preferred-habitat agents. Given the results of Section 4.1, it should not be surprising that the presence of these agents—which changes the shape of the distribution that arbitrage investors have to hold—has little effect on the shape of the yield curve, once the parameters are calibrated to achieve the same average level of long-term rates. However, as illustrated by the red lines, introducing preferred habitat does matter for the sensitivity of yields to duration. In particular, it dampens the yield response to a duration shock.

The bottom panel shows why this is so. In the model without preferred-habitat agents, represented by the solid bars, the downward shift in z causes investors to hold more short-term securities and fewer long-term securities, reducing their duration risk and resulting in lower term premia. (These bars represent the same change as the top panel of Figure 6.) But when preferred-habitat agents enter the model they end up holding a substantial fraction of long-term bonds in most states of the world—indeed, in the case considered here, preferred-habitat agents own the entirety of the market in maturities greater than four years. Consequently, arbitrageurs' holdings of those bonds do not change when the aggregate supply changes; they stay at zero. The preferred-habitat agents essentially absorb part of the change in duration risk.¹⁶ The change in the arbitrageurs' portfolio, illustrated by the dotted bars, is limited to fluctuations in securities with duration of four years and less, which has a relatively small impact on risk. Clearly, this result is potentially sensitive to the particular modeling choices made here, which are in some respects arbitrary. With some adjustment to functional forms—or even a somewhat different calibration of the

¹⁶ The preferred-habitat agents do demand less of all securities when yields fall, by assumption. But in the example here this effect is of relatively little importance because preferred-habitat demand for longer-term securities exceeds their supply for both values of z considered.

above model—it would be possible to generate situations in which preferred-habitat agents’ portfolios tilt toward more duration when yields fall, thus amplifying rather than offsetting the effects on arbitrageurs. Without more information on the behavior of these agents, however, it is difficult to come to firm conclusions in this regard.

4.3 Timing

The comparative-statics exercises of Section 3 can be viewed as one-time, unanticipated, permanent shifts in the supply distribution. What happens if these shifts are anticipated? For example, suppose that an upcoming shift is revealed on one date but not implemented until a later date. The model allows one to examine these “announcement effects” versus “implementation effects” by considering differences in timing between changes in the conditional moments ($E_t[\mathbf{q}_{t+1}]$ and $\mathbf{\Omega}_t$) and changes in the portfolio weights (\mathbf{x}). One reason this is relevant is that most central-bank asset purchase programs have been structured to conduct their operations over spans of several months. Furthermore, yields likely fluctuate in advance of these programs’ announcements (as, for that matter, they likely do in advance of Treasury auctions), as the expectations of changing supply are priced in. These types of dynamics and expectational effects will be particularly important when we consider models with stochastic supply in the following section.

Figure 10 shows the effect of the same experiment as in Figure 6, with the adjustment that, although investors know with certainty that the new distribution of Treasury maturities will be in place next year, they still face the old distribution this year. Thus, the experiment illustrates how risk premiums evolve leading up to perfectly anticipated duration shocks. The violet line shows the change in the yield curve in the first period (i.e., after the announcement but before the actual purchases). The red line shows the cumulative change after purchases have occurred, the same total change that was illustrated in Figure 6. Most of the effects of the program occur on the announcement, as expectations of the supply reduction are priced in. However, about 11 of the total 56 basis point reduction in the 10-year yield occurs with the purchases themselves. The reason that the perfectly anticipated supply withdrawals are not fully priced in on the announcement is that investors must continue to bear the greater interest-rate risk of long-term securities during the course of the year between the announcement and the purchases. Note, in particular, that the two-period yield does not react to the announcement at all. This is because a decrease in duration one period in the future does nothing to benefit the investors that currently hold that bond—next period, the current two-period bond will be a one-period bond, which bears no interest-rate risk.

4.4 Stochastic supply

As the previous section suggested, prices will typically change in advance of supply fluctuations if those fluctuations are predictable. It therefore makes sense to consider models in which this is the case, that is,

in which the distribution \mathbf{x} follows a known stochastic process. This requires adding a second factor to the model.

Specifically, I calibrate the dynamics of the maturity distribution by estimating an AR(1) model on the yearly time-series of the average duration of U.S. government obligations in the hands of the public (see Figure 2). This gives the following result:

$$z_t = 0.11 + .972z_{t-1} + \varepsilon_t^z \quad (26)$$

where the standard deviation of ε_t^z is 0.16. I assume that ε_t^z is uncorrelated with ε_t^r . Note that, although z_t has the interpretation of the average duration of debt outstanding, the entire distribution \mathbf{x}_t is still the input in the model, and the shape of this distribution at each point in time depends on z_t through equation (21). Furthermore, it is important to recognize that investors price in the stochastic variation in z_t in such a way that no-arbitrage continues to hold. The equilibrium condition (14) is unchanged, except that now \mathbf{x} has a time subscript. One consequence of the additional risk factor is that the level of risk aversion needed to match the average 15-year slope of the yield curve is a bit lower than in the one-factor model. Specifically, I now set $a = 12.5$.

Figure 11 shows the dynamic response of the yield curve to one-standard-deviation positive shocks to r_t^f and z_t in this model. Because of nonlinearities, the responses depend on initial conditions. I illustrate the typical effects by evaluating both at the sample means of the two state variables and at the nonstochastic steady states of the two state variables. For the short rate the steady state is zero, so the latter case is also illustrative of how responses at the ZLB differ from their average.

In response to a positive shock to the short rate, all yields increase, and the size of the increase is smaller for higher maturities. Over time, the responses decline monotonically to zero, as one would expect given the assumed AR-type process for r_t^f . (These results were true in the one-factor model as well.) A shock to average duration has almost the opposite effect across maturities: short rates do not move at all, and 30-year rates move by 15 to 20 basis points. The initial response is qualitatively similar to the comparative-statics exercises depicted in Figures 6 and 7. In the dynamic model here, these responses decay over time, as the shock to z_t dies out. The responses to duration shocks are about 25% smaller at the ZLB than at the sample average, again reflecting the lower volatility of interest rates in that state of the world. Having stochastic supply in the model adds a second source of variation to longer-term interest rates and thus, all else equal, increases their volatility. Nonetheless, as the figure suggests, the effect of the typical duration shock on, say, the 10-year yield is only about one-tenth of the effect of the typical short-rate shock. Thus, at least under this calibration, duration shocks account for only a modest fraction of the overall variation in longer-term yields. (This is why the calibrated value of a only has to be moved down by a small amount relative to the one-factor model.)

5 Imperfect Substitutability

I now consider ways of introducing imperfect substitutability among assets by allowing non-zero values of the convenience yield b . In addition to imparting greater realism (in the sense of capturing anecdotal and other evidence on investor behavior) these extensions will be valuable in allowing the models to fit certain features of the yield-curve data. For illustration here, I continue with the one-factor version of the model, in which the supply distribution is static. The following section incorporates these effects into a model that also includes stochastic supply.

It may be that investors receive a benefit from holding securities of particular maturities. Such a formulation could be motivated, for example, by a planning horizon longer than one period, by a desire to match long-duration liabilities, or by supposing that short-term assets provide liquidity services. (Cox et al., 1981, Krishnamurthy and Vissing-Jorgenson, 2012, and Greenwood et al., 2013, provide models in which this can occur in different forms.) In particular, suppose that the return on an individual security is given not just by its price appreciation but also by an unobserved benefit that depends on its maturity. For parsimony, I assume that this dependence is linear in the end-of-period maturity, $n-1$:

$$b^{maturity}(n) = b(n-1) \quad (27)$$

for some parameter b . From a modeling perspective, adding $b^{maturity}$ to the model will generally help to better match the curvature of the empirical yield curve.

In a recent paper, Hanson and Stein (2012) propose that preferences over yields, rather than returns, could explain some of the apparently large movements of longer-term yields over the business cycle. In particular, they sketch a two-period model with arbitrageurs and “yield-oriented” investors, where the latter agents have preferences over yields. With heterogeneous agents, this model is similar to the preferred-habitat case considered in Section 4.2. An alternative way of incorporating the same idea, while maintaining the representative-investor framework, is to allow investors to receive benefits on bonds based on their current yields, in excess of the short-term yield:

$$b^{yield}(n) = c \left(p_m^{-1/n} - \exp[-r_t^f] \right) \quad (28)$$

where $c > 0$ is a parameter reflecting the additional benefit generated by the current yield. As argued by Hanson and Stein, this benefit could reflect gains due to tax, regulatory, or accounting distortions. As in their model, its practical effect is to increase the volatility of longer-term yields. In particular, it

magnifies the response of long-term yields to short-rate shocks, a phenomenon reminiscent of “reach for yield”-type behavior.¹⁷

I combine the yield and maturity preference specifications into the single convenience-yield function $b(n) = p_{nt}(b^{maturity}(n) + b^{yield}(n))$, where the multiplication by the price of bond n imposes that the benefit is proportional to the market value, rather than the face value, of the security held in the portfolio at time t . This gives the following specification for returns:

$$R_{nt} = \frac{P_{n-t+1}}{P_{nt}} + b(n-1) + c(p_{nt}^{-1/n} - \exp[-r_t^f]) \quad (29)$$

For illustration in the one-factor mean-variance model, I calibrate the three parameters a , b , and c to match the average 15-year slope of the yield curve, the average curvature of the yield curve (as measured by the difference between the 1-15 and 15-30-year slopes over the more limited sample for which the 30-year yield is available), and the sample variance of the 15-year yield. This gives $a = 29.1$, $b = 0.0056$, and $c = 0.97$. Taken literally, these parameter values are quite large. For example, the calibrated value for b implies that, holding risk constant, investors would be willing to forego a 16% return on a short-term bond in order to get the unobserved maturity benefits of holding a 30-year bond. The reason that such large values are required is that, as noted earlier, the one-factor mean-variance model has difficulty explaining the curvature and volatility of the long end of the yield curve with duration effects alone. This will be less true in the multi-factor model below, where inflation and supply dynamics naturally contribute to the curvature and volatility of longer-term rates, and consequently the imperfect substitutability parameters in that model are estimated to take significantly smaller values.

Figure 12 reconsiders the LSAP-type comparative statics using the model with the imperfect-substitutability terms as calibrated above. I examine both the case in which the short rate is at its sample-average value (top) and in which it is at the ZLB with one year of forward guidance (bottom). The solid lines show the effect on the yield curve of reducing z by 0.7 years in the baseline model without imperfect substitutability (the same as in Figures 6 and 7). The dashed lines show the corresponding yield curves using the imperfect-substitutability model. The inclusion of these terms increases the curvature of the yield curve in both panels. It also increases the sensitivity of long-term yields to the level of the short rate, as can be seen by comparing the blue lines in the bottom panel. The sensitivity of yields to the supply distribution is also higher under this calibration. The removal of 0.7 years of duration reduces the 10-year yield by 103 basis points when the short rate is at its average value and by 54 basis points at the ZLB—effects that are two to three times larger than in the baseline model.

¹⁷ Importantly, b^{yield} depends on the *relative* levels of short- and long-term yields. If it is instead modeled as depending on the absolute level of yields, the effect is to *dampen* short-rate shocks.

This result must be interpreted with caution, however. Almost the entirety of the difference between the effect of the LSAP in this model and the baseline is due to the much larger value of risk aversion used here, not to the effects of the convenience-yield terms themselves. If risk aversion is set to its baseline value of $a = 12.8$, maintaining the above values of b and c , the LSAP effect is 57 basis points at the sample average and 18 basis points at the ZLB, almost identical to the effect in the model in which b and c were zero. Thus, the primary effect of the convenience yield is to change the average shape of the yield curve and its response to short-rate shocks, without having any major consequences for its sensitivity to the supply distribution.

6 A Three-Factor Model with Inflation Risk

The models above illustrate several aspects of the portfolio-balance approach that might be relevant, and one can certainly imagine a variety of other possible extensions. One final case that is particularly important to consider is the addition of inflation, since inflation—or, more precisely, time-varying inflation *risk*—has been shown to be an important determinant of nominal term premiums. (See the literature surveyed by Gurkaynak and Wright, 2012.) This section adds inflation to the model and brings it to the data. To study the importance of supply fluctuations properly, I also include the stochastic-supply factor discussed in Section 4.4 and the convenience-yield terms discussed in Section 5. Rather than struggling to calibrate the several free parameters of this model, I estimate them using the data on long-term Treasury yields.

5.1 Model Setup

Investors now solve the problem

$$\max_{w_t} E_t [\mathbf{R}'_{t+1} \mathbf{w}_t - \pi_{t+1}] - \frac{a}{2} \text{var}_t [\mathbf{R}'_{t+1} \mathbf{w}_t - \pi_{t+1}] \quad (30)$$

where π_{t+1} represents the gross rate of inflation between periods t and $t+1$. I allow for the convenience yields discussed in the previous section, so each nominal asset return R_{nt+1} is given by equation (29). It is straightforward to show that prices solve

$$p_{nt} = \exp[-r_t^f] \left(E_t [p_{n-1t+1}] - a \left(\frac{\boldsymbol{\Omega}_t \mathbf{x}_t}{\mathbf{p}'_t \mathbf{x}_t} - \text{cov}_t [p_{n-1t+1}, \pi_{t+1}] \right) + b(n-1) + c \left(p_{nt}^{-1/n} - \exp[r_t^f] \right) \right) \quad (31)$$

for $n = 2, \dots, N$, and that the nominal pricing kernel takes the form

$$M_{t+1} = \theta_{0t} + \theta_{1t} (R_{t+1}^w - \pi_{t+1}) \quad (32)$$

It is clear from these equations that changes in \mathbf{x}_t have no *direct* effect on inflation-risk premia through portfolio rebalancing. They may change the covariance of prices with inflation through their effect on the overall volatility of asset prices, but these effects will generally be second-order.

5.2 Estimation

I assume that inflation and the short rate are jointly determined by the process

$$r_t^f = \phi_0 + \phi_1 r_{t-1}^f + \phi_2 \pi_t + \varepsilon_t^r \quad (33)$$

$$\pi_t = \gamma_0 + \gamma_1 \pi_{t-1} + \varepsilon_t^\pi \quad (34)$$

subject to

$$\Pr[r_t^f < 0] = 0 \quad \forall t$$

where ε_t^r is a truncated-normal disturbance with shape parameter σ_r , and ε_t^π is distributed normally with time-varying variance:

$$\sigma_\pi^2 = \rho_0 + \rho_1 |\pi_{t-1}| \quad (35)$$

where ρ_0 and ρ_1 are both non-negative. This system produces behavior of the short rate similar to that depicted in Figure 1, except that now the configuration of the curves in that figure depends on the level of inflation. The dependence of the short rate on contemporaneous inflation, together with the heteroscedasticity of inflation shocks, provides a simple way of capturing time-varying covariance between r_t^f and π_t .¹⁸ The process for z_t is assumed to be the same as in Section 3.3.

Estimation proceeds by fitting the measurement equations

$$y_t^n = \hat{y}_t^n + e_t^n \quad (36)$$

where y_t^n is the n -period yield in the data, \hat{y}_t^n is the corresponding yield generated by the model, and e_t^n is an iid normally distributed measurement error, for $n = \{5, 10, 15\}$. Estimation of this model without further restrictions indicated that it was not well identified from the data on longer-term Treasury yields alone. This is perhaps unsurprising, given that yields are highly serially and cross-sectionally correlated and that several of the model parameters (e.g., risk and risk aversion) are likely to have very similar effects. I therefore take two measures to aid identification. First, as is common in the term-structure literature, I restrict the covariance matrix of the measurement errors to be diagonal. (This is equivalent to

¹⁸ The assumption that inflation does not depend on the interest rate in equation (34) is motivated by the empirical observation that lags of r_t^f have essentially no marginal predictive power for π_t at an annual frequency, once the lag of π_t is controlled for.

minimizing the weighted mean-squared errors over the three long-term yields.) Second, I bring in information from the observable dynamics of the short rate and inflation themselves, rather than just the information contained in the Treasury bond yields. Because the data are observed monthly while the periods in the model are taken to be annual, joint estimation is not straightforward, but the mixed frequencies are easily dealt with using Bayesian methods. Specifically, I first estimate the system (33) – (35) by maximum likelihood on non-overlapping annual samples (similarly to the way the calibration was set for the smaller model in Section 3), using only the one-year Treasury rate and core PCE inflation. The results are reported in the top line of Table 1. I then use these estimates to form a prior distribution which I combine with the likelihood generated by the measurement equations (36). I put weak priors on the parameters a , b , and c , which are not identified from the first-stage estimation. All of the parameters were then estimated using 20,000 Metropolis-Hastings draws from the posterior distribution, with a burn-in sample of 2,000 discarded. Because the forward guidance issued by the Federal Reserve after the short rate hit the ZLB affected expectations of the path of rates in an uncertain way, I exclude this period from the sample and use only the data from 1971 – 2008.

5.3 Results

Model fit

The parameter estimates (means and standard deviations of the posterior distributions) over the 1971 – 2008 period are shown in the second row of Table 1. Since $\phi_2 > 0$, inflation is correlated with the short rate. In addition, the variance of inflation varies significantly with its level ($\rho_1 > 0$). This implies a time-varying inflation risk for all bonds. However, since inflation itself is not highly persistent ($\gamma_1 = 0.72$), this risk is greatest for short- and medium-term securities, adding curvature to the yield curve. The maturity preference parameter b is 0.0009, implying a mild preference for holding long-term over short-term securities. The log of the yield-preference parameter c is estimated to be -3.1, which is inconsequential.

Table 2 reports various statistics summarizing where the model succeeds and fails in fitting the data, when evaluated at the posterior mode. I assess the fit over spot yields (which were used in the estimation), forward rates, and excess returns. The one-year forward rate ending n years ahead is calculated from bond prices, in both the model and the data, as

$$f_{nt} = (p_{n-t}/p_{nt}) - 1 \quad (37)$$

Excess returns are computed annually as

$$exret_{nt} = (p_{n-t}/p_{nt-1}) - R_{1t} \quad (38)$$

(To avoid overlapping data, these returns are only calculated using the December data in each year.) I report statistics for 5-, 10-, and 15-year bonds over the entire sample and for 30-year bonds over the post-1985 sample when these data are available. Since the 30-year bonds were not used in the estimation, they provide a sort of out-of-sample check on the results.

Given the tight structural restrictions of the model and the relatively parsimonious parameterization, we would not expect it to fit the data as well as, say, a reduced-form multifactor affine model. Nonetheless, it generally does well, coming close to matching most of the first and second moments of yields within-sample. It still predicts a bit too little volatility (of course, it is typically true that fitted values are less volatile than the data they are fitting) and a bit too much correlation of longer-term yields with the short rate. However, the model explains about 90% of the variation in yields of up to 15 years, and it even explains 73% of the variation in the out-of-sample 30-year yield. Perhaps most remarkably, the model generally matches the features of excess bond returns through the 15-year maturity, with R^2 's of 0.33 to 0.63. Although it misses substantially on the overall features of the 30-year bond returns, it still manages to explain 14% of their variation as well.

Impulse-response functions

Figure 13 examines the effects of shocks in the model. Again, I consider one-standard-deviation shocks (in a direction that leads to positive yield responses). I consider initial conditions both at the sample mean (left column) and at a set of values that reflect conditions during the LSAP period (right column)—specifically, $r_0^f = 0.003$, $\pi_0 = 0.014$, and $z_0 = 2.8$.

Roughly speaking, short-rate shocks and inflation shocks have similar effects on long-term yields. Short-rate shocks have larger initial impact, but they die out rather quickly. The inflation shocks have more-persistent effects on short rates and also increase inflation uncertainty, both of which serve to amplify their impact on the long end of the curve. The dynamics of the system impart a hump shape to the short-rate response to inflation shocks over time, and this is reflected in the shape of the yield curve reaction to such shocks across maturities. There is not much qualitative difference between the responses at the steady state and during the LSAP period. The responses to duration shocks have a similar shape as in the model of Section 4.4. However, they are about one-third smaller near the ZLB than at the sample mean, again because of the lower interest-rate volatility in that environment.

Term premium decomposition

The top panel of Figure 14 shows how the model decomposes the actual time series of ten-year yields into an expectations component, a term premium, and a model-error term. (Since only the period 1971 – 2008 was used in the estimation, the last four years in the Figure are out-of-sample estimates.) The

expectations component is calculated as the value of the ten-year yield that would be implied by the model at each point in time if a , b , and c were equal to zero. The term premium is then just the difference between the model-implied yield on each date and the expectations component. As noted above, the model generally tracks the behavior of longer-term interest rates well over this time. The most notable exceptions are the early 1980s, during the Volcker disinflation when the market-perceived inflation risk may have been higher than the level captured by the model, and the period from about 2005 – 2007, corresponding to the so-called “conundrum” in long-term interest rates (see Gurkaynak and Wright, 2012). The latter is a well known anomaly in the yield curve that is not likely to be reproduced by the factors considered explicitly here. However, one possible explanation is that a large quantity of foreign investment increased the demand for longer-term Treasuries—the “global savings glut” introduced by Chairman Bernanke.¹⁹ Conceivably, this type of mechanism could be incorporated into the portfolio-balance model, perhaps through the preferred-habitat version developed earlier.

The second panel shows the term premium. In broad strokes, it follows a trajectory that is similar to the term premium produced by reduced-form models, such as Kim and Wright (2005), rising through the 1970s and early 1980s, and falling since the late 1980s. (Rudebusch et al., 2007, estimate the Kim-Wright term premium for a sample that covers this entire period.) Because of nonlinearities, the decomposition of the term premium into its structural contributing factors is not uniquely defined. However, to get a sense of the relative contributions, I calculate counterfactual yields holding z_t and π_t equal to their sample means in all periods and under the assumption that investors *know* that they will always be equal to these values. The latter assumption serves to eliminate the contribution of the uncertainty about each variable to the term premium. A non-trivial fraction of the movement in the term premium since the late 1990s is attributed to the reduced interest rate risk associated with the short rate approaching the ZLB. This can be seen by the dashed line in the middle panel, which isolates the contribution of the short rate to the term premium by setting both duration and inflation to their sample averages. It is also clear from this figure that these two factors have jointly contributed to significant variation in the term premium relative to what would have been implied by variation in the short rate alone.

The bottom panel shows the contributions of inflation and Treasury supply to the term premium, measured by subtracting the counterfactual yield in which both series are held constant from the counterfactual yield in which only one or the other is held constant. Overall, fluctuations in inflation are estimated to have moved term premiums within a range of about 100 basis points over the course of the sample. Consistent with previous research (e.g., Campbell et al., 2009), these fluctuations have mostly

¹⁹ See “The Global Savings Glut and the U.S. Current Account Deficit.” Speech delivered at the Sandridge Lecture, Virginia Association of Economists, Richmond VA, March 10, 2005.

been in the form of a downward trajectory since the early 1980s. On the other hand, the Treasury supply distribution shifted significantly to longer maturities during the period from about 1975 to 1985, and the model suggests that this resulted in an upward movement in term premiums of about 150 basis points. Since the 1980s, supply-related movements have been relatively small because shifts in Treasury supply itself have been small. Broadly speaking, the results are consistent with the finding of Li and Wei (2012) that measures of Treasury supply do affect term premiums over time. The net result of all of these effects, returning to the middle panel, has been an increase in the overall term premium from just over 100 basis points at the beginning of the sample to a peak of almost 300 basis points around 1990, followed by a decline to near its original levels by the end of the sample. The early run-up was due primarily to an increase in duration risk, while the subsequent decline was due to a decrease in inflation risk, together with the approach of the ZLB.²⁰

Table 3 summarizes these findings by comparing the monthly change in the model-implied ten-year yield and term premium to the changes that would have occurred in the counterfactual scenarios. By this measure, inflation accounts for 39% of the overall variance in term premiums and fluctuations in Treasury supply account for 70%. The two factors are correlated in the sample, so excluding them both reduces the volatility of the term premium by 94%. Although these fractions are large, it is important to bear in mind that most interest-rate fluctuations in the model are not due to term premiums but rather to expected short rates. Indeed, as shown in the top row of the Table, holding both inflation and average duration at their sample averages reduces the variance of 10-year yield changes by just 6%.

7 Assessment of the LSAP programs

In this section, I consider the Federal Reserve's LSAP programs using the three-factor model presented above. The relevant aspects of the programs are summarized in Table 5. A large component of the Federal Reserve's purchases consisted of agency mortgage-backed securities. Throughout the paper I have ignored agency MBS when calculating the duration of government liabilities, but this exclusion is likely inappropriate because these securities have typically been perceived to carry an implicit government guarantee. Two technical measurement problems arise when trying to extend the analysis to agency MBS. The first is that comprehensive data on the maturity structure of outstanding MBS are not available. The second is that, unlike the case of zero-coupon Treasuries, duration for MBS depends on economic conditions—in particular, it depends on the level of interest rates through the negative convexity induced by the prepayment option. However, according to the Barclays MBS index, the

²⁰ One might wonder to what degree these results are due to the inclusion of the imperfect-substitutability terms in the model. The answer is very little. The parameter values associated with these terms are estimated to be small, and they affect the term premium estimates by less than 10 basis points in most periods.

average duration of MBS since 1989 is about three years, only slightly lower than that of Treasury debt. (See Hanson, 2012.) In performing the calculations below, I continue to use the approximation of an exponential distribution for the duration distribution. Moreover, I assume that asset purchases did not materially change the average duration of MBS outstanding, but I do incorporate the *amount* of MBS outstanding, so that Fed purchases swap their duration for zero-duration reserves.²¹

The net effect of the LSAP programs, as of December 2012, was to reduce the outstanding supply of Treasury and MBS securities by \$2 trillion and increase reserves by a similar amount. Moreover, the Treasury securities removed from the market were, after the duration twist induced by the Maturity Extension Program (MEP), almost all of maturity greater than five years. Consequently, the average duration of Treasury securities in the hands of the public after all was said and done was about 0.5 years lower than it otherwise would have been. Taking all of these facts together, we can estimate how investors' duration changed as a result of the programs. Specifically, as of Q4 2008, Treasury debt held by the U.S. public was \$5.3 trillion, agency MBS was \$5.0 trillion, and the monetary base was \$1.4 trillion. The average duration of Treasuries (excluding the monetary base) in the hands of the public was 3.1 years as of the beginning of the program. Therefore, assuming that MBS have an average duration of three years (and again counting the base as zero duration), the weighted average of Treasuries, MBS, and the monetary base in the hands of the public was 2.7 years when the first LSAP was announced. Holding all else constant, this value would have fallen to 2.0 years as a result of the programs.²²

The final column of the table presents the range of values for the empirical effects of the LSAPs on the ten-year term premium, culled from the literature that has examined this question. In particular, this range draws on estimates from Gagnon et al. (2010), Krishnamurthy and Vissing-Jorgensen (2011), Ihrig et al (2012), D'Amico et al. (2012), and Rosa (2013). (In some cases, additional minor calculations were required to make the results comparable.) Of course, all of these estimates are subject to a high degree of uncertainty surrounding both parameter values and specification. Furthermore, since most of these papers employ an event-study methodology using program-announcement dates, the results of Section 4.3 suggest that the estimates may be understated. Nonetheless, they provide a rough guide to what we should expect for the combined effects of the programs—likely on the order of 100 to 250 basis points altogether. To be clear, this is (according to the authors of the studies) only the effect on the term premium, controlling for changes in expectations of the short rate, and it only reflects the initial impact of the programs, ignoring any subsequent dynamic effects.

²¹ Data on MBS quantities come from SIFMA (www.sifma.org/research/statistics.aspx).

²² I ignore the effect of the Fed's \$165 billion of agency debt purchases on the outstanding duration of government securities, although the resulting increase in reserves is included.

How do these estimates compare to the results of the portfolio-balance model? The black line in Figure 15 shows the effect of removing 0.7 years of duration from the hands of investors, starting from a level of duration that is equal to what was observed prior to the LSAP's introduction (2.7 years) and a configuration of short rates and inflation that approximates the situation on average since that time. Specifically, inflation is set at a level of 1.4 percent (its average over the period 2009 – 2012), and the short-term interest rate is taken to be near the zero lower bound with two years of forward guidance in place. Although the market's perception of the length of time that the ZLB would bind surely fluctuated over this period, the assumption of two years is likely close to the average that prevailed, perhaps erring on the conservative side (see Femia et al., 2013). As a reality check, the first row of Table 5 shows that the slope of the yield curve produced by the model under these assumptions over the LSAP period was reasonably close to the average slope of the yield curve in the data during this time. (Recall that this part of the sample was not used in the estimation, and that the model was estimated under the assumption of no forward guidance.) The effect of the LSAPs on the ten-year yield of the LSAP shock in this environment is a mere 18 basis points. This is the initial impact—given the estimated dynamics of z_t , the effect would decay to zero over time, similarly to the response shown in Figure 13.²³

The other lines in the Figure check the robustness of this result to the most likely sources of measurement error. First, although the two-year period assumed for the forward guidance may be a reasonable approximation in terms of the time interval involved, the assumption that the Fed committed to the level of the short rate with absolute certainty for this entire time could be too strong, given that the forward guidance issued by the FOMC has always left open the possibility that rates could rise sooner than expected if economic conditions warranted. Indeed, empirical measures of short-term uncertainty about the short rate have consistently remained significantly above zero while forward guidance has been in place. The green line examines the sensitivity in this dimension by instead imposing that investors believe that the ZLB will bind with certainty for only one year following the LSAP shock. This has the effect of raising the standard deviation of the one-year-ahead short rate from zero to 1.2%, and it brings the 10-year slope into even closer alignment with the data. However, it increases the effect of the LSAPs only to 24 basis points. Second, given that the calculation of the size of the 0.7-year duration shock was somewhat back-of-the-envelope, one might worry that it is less than the true amount of duration removed by the LSAPs, perhaps understating the duration outstanding at the beginning of the programs and overstating the amount left at the end. The orange line shows the effect of a shock that is twice as large as

²³ Since the persistence of z_t is high, there is not much difference in the impact if one models the LSAP shocks as permanent, rather than transitory, withdrawals of duration. Of course, it may well be that investors perceived the persistence of the LSAPs as being lower than that of the typical Treasury-supply shock, in which case the impact of the LSAPs on yields would be smaller than the value reported here. The decay in the effect over time is qualitatively consistent with the evidence presented in Ihrig et al. (2012) and Wright (2012).

the baseline, going from an average duration of 3.05 years to 1.65 years; the effect on the ten-year yield is about double the previous estimate, but it is still only about a third of the low end of the empirical range of LSAP effects.

One might also object that, during much of the LSAP period, the average level of risk aversion was higher than usual. The second row of Table 5 shows the effect of doubling the estimated risk-aversion coefficient to a value of 25.4. The impact of the LSAP shock in this case is still just 29 basis points under the assumption of two-year forward guidance. This value rises to 60 basis points—perhaps half of the empirical value—if one assumes shorter forward guidance in addition to higher risk aversion, but that scenario does not seem likely to be apposite, as it implies a counterfactually steep slope for the yield curve (420 bp, rather than 268 bp, on average). Finally, one might conjecture that the convenience-yield parameters b and c also rose to higher-than-normal levels during the crisis. As suggested in Section 5, changes in these parameters are generally of little importance for the magnitude of duration effects, holding risk aversion constant. The last row of the table corroborates that result in the current model—doubling the values of both parameters (maintaining $a = 25.4$) leaves the impact of LSAP shocks about unchanged, though it does help somewhat with the steepness of the slope. While it is possible to find alternative parameter values or other modifications that could achieve an LSAP effect in the model on the order of what has been estimated in the literature, I conclude that this would require a departure from the data in other respects that would be well beyond the bounds of statistical uncertainty.

7 Conclusion

This paper has presented a new method for studying the term structure of interest rates and, in particular, the ways in which the term structure is affected by supply fluctuations. The type of model considered is a rational-expectations version of portfolio-balance models that have been in use, with varying degrees of formality, for decades. It may be viewed as a generalization of the preferred-habitat model of Vayanos and Vila (2009), with investors that potentially have a broader class of objective functions and in which nonlinearities may be important. It is hoped that this approach may be of use in studying other asset-pricing phenomena as well.

One feature of the particular set of models to which I have applied the approach is that, for various reasons discussed earlier, they likely overstate the effect that supply fluctuations could have through a duration channel. The reasonably strong fit of the model suggests that the exaggeration may not be large, although it does call for some caution when thinking about results like the term-premium decomposition. However, the upper-bound nature of the models can be viewed as an advantage when trying to gauge the likely effects of asset-purchase programs, since even under their potentially strong assumptions the models suggest that those effects are fairly small. I interpret these results, not as evidence that the LSAPs

were ineffective, but as evidence that they probably had their effects primarily through mechanisms other than the removal of duration risk. It seems possible that phenomena not easily captured by a no-arbitrage model, such as market dislocations, liquidity shocks, and capital constraints, were important during and after the financial crisis and that LSAPs had additional effects through those channels. This reading is broadly consistent with the empirical conclusions of Krishnamurthy and Vissing-Jorgensen (2011), Cahill et al. (2013), and D'Amico and King (2013). More structural modeling of the behavior driving such scarcity effects is needed to determine whether they would significantly alter the results presented above.

Table 1. Parameter estimates for the three-factor model

	ϕ_0	ϕ_1	ϕ_2	$\log \sigma^r$	γ_0	γ_1	$\log \rho_0$	ρ_1	a	b	$\log c$
<i>Prior</i>	-0.005 (0.009)	0.8 (0.1)	0.4 (0.2)	-4.2 (0.4)	0.003 (0.005)	0.9 (0.2)	-20 (21)	0.003 (0.002)	12 (5)	0 (0.001)	-2 (2)
<i>Posterior</i>	0.0005 (0.0010)	0.89 (0.01)	0.21 (0.03)	-4.0 (0.05)	0.0047 (0.0015)	0.72 (0.05)	-12.4 (2.1)	0.0015 (0.0007)	12.6 (1.0)	0.0009 (0.0002)	-3.1 (0.3)

Notes: The table reports parameter estimate—means with standard deviations in parentheses—for the model in equations (33) through (35). The prior distribution is jointly normal and is based on maximum-likelihood estimation of the system using annual data on the one-year Treasury rate and core PCE inflation. The reported posterior values are based on 20,000 Metropolis draws, using monthly data on 5-, 10-, and 15-year yields from 1971 – 2008.

Table 2. Comparison of the three-factor model to the data.

A. Yields (%)

		1yr	5yr	10yr	15yr	30yr (beg. 12/85)
Data	Mean	6.4	7.1	7.5	7.7	6.6
	Std. Dev.	3.0	2.6	2.4	2.3	1.5
	1-year autocor.	0.81	0.87	0.88	0.88	0.91
	Corr w/short rate	1	0.96	0.92	0.89	0.71
Model	Mean	6.4	7.0	7.4	7.6	7.0
	Std. Dev.	3.0	2.7	2.3	2.1	1.3
	1-year autocor.	0.81	0.82	0.83	0.84	0.87
	Corr w/short rate	1	1.00	0.99	0.97	0.88
	R^2	--	0.92	0.87	0.84	0.73

B. Forward rates (%)

		1yr	5yr	10yr	15yr	30yr (beg. 12/85)
Data	Mean	6.4	7.5	8.0	8.1	6.0
	Std. Dev.	3.0	2.4	2.1	2.2	1.9
	1-year autocor.	0.81	0.88	0.87	0.88	0.83
	Corr w/short rate	1	0.88	0.82	0.81	0.61
Model	Mean	6.4	7.5	8.0	8.0	6.9
	Std. Dev.	3.0	2.4	1.8	1.4	1.1
	1-year autocor.	0.81	0.83	0.85	0.86	0.88
	Corr w/short rate	1	0.98	0.94	0.87	0.69
	R^2	--	0.80	0.77	0.68	0.54

See next page for notes.

Table 2 continued.

C. Annual excess returns (%)

		1yr	5yr	10yr	15yr	30yr (beg. 12/85)
Data	Mean	0	2.1	3.4	4.4	15.3
	Std. Dev.	0	6.5	12.4	17.9	29.8
	1-year autocor.	--	0.04	-0.09	-0.12	-0.58
	Corr w/short rate (end of period)	--	-0.48	-0.45	-0.40	-0.20
Model	Mean	0	1.7	2.9	3.6	2.7
	Std. Dev.	0	6.3	12.2	16.4	5.6
	1-year autocor.	--	0.20	0.19	0.18	0.01
	Corr w/short rate (end of period)	--	-0.44	-0.43	-0.43	-0.52
	R^2	--	0.63	0.40	0.33	0.14

Notes: Data are monthly averages of daily values, Dec. 1971 – Dec. 2008, except for 30-year yield. Model results are generated using the mean values of the parameters reported in Table 1, which were estimated over the 5-, 10-, and 15-year spot yields. Reported forward rates are one-year rates ending 1, 5, 10, 15, and 30 years ahead. Excess returns are computed on non-overlapping samples, December to December of each year, and are calculated relative to the initial one-year yield. Excess returns in the data are calculated from the Gurkaynak et al. (2007) zero-coupon yields. Shading indicates model output that exactly matches the data by construction.

Table 3. Contributions to the ten-year yield in the three-factor model

	Inflation factor	Supply factor	Both factors
Yield	0.07	0.06	0.06
Term premium	0.39	0.70	0.94

Notes: Table reports the percentage by which the total variance of the ten-year yield or term premium falls when the indicated factors are held at their sample-mean values, and investors price these levels in with certainty. Variances are computed over monthly first differences, 1971 – 2012.

Table 4. Summary of Federal Reserve asset-purchase programs

	Dates	Net quantity of Treasuries purchased (\$bil)	Net quantity of MBS purchased (\$bil)	Net quantity of reserves created (\$bil)	Assumed change in average Treasury duration outstanding	Empirical effect on ten-year term premium
LSAP I	Dec. 2008 – May 2010	\$300	\$941	\$1,318	0	40-100 bp
MBS reinvestment	Aug. 2010 – Oct. 2011	\$285	\$0	\$285	0	10-25 bp
LSAP II	Nov. 2010 – July 2011	\$600	\$0	\$600	0	15-55 bp
MEP	Oct. 2011 – Dec. 2012	\$0	\$0	\$0	-0.5	30-65 bp
Total		\$1,185	\$941	\$2,203	-0.5	95-245 bp

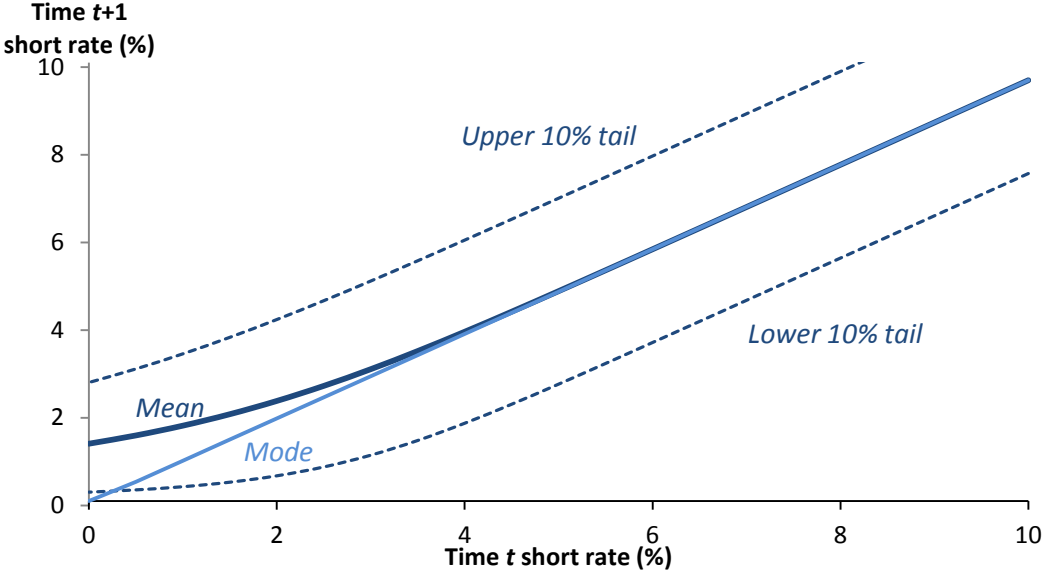
Notes: All quantities are net of redemptions and principal payments through December 2012. The assumed duration change is only that in Treasury securities (i.e., excluding agency debt, MBS, and the monetary base). The empirical effect of each program on the ten-year nominal Treasury yield is taken from the literature discussed in Section 7 of the text.

Table 5. Effects of alternative parameter assumptions on the LSAP impact

		Effect of LSAPs on 10y term premium (bp)	Ave. 10y slope, Dec. 2008 – Dec. 2012 (bp)
Data		95 - 245	268
<i>Parameters</i>	<i>Forward guidance</i>		
Estimated Model	2 years	18	214
($a = 12.7$; $b = 0.0009$; $c = 0.05$)	1 year	24	273
$a = 25.4$	2 years	46	328
	1 year	60	420
$a = 25.4$; $b = 0.0018$; $c = 0.10$	2 years	45	280
	1 year	59	371

Notes: The estimated model is the three-factor model described in Section 6 of the text, using data from 1971 – 2008, with parameters evaluated at the posterior mean. The LSAP shock is assumed to be a decrease in average duration outstanding from 2.7 to 2.0 years. “Forward guidance” means that investors know that the short-term interest rate will be equal to its current value for the indicated time after the shock.


Figure 1. Estimated short-rate processes




Notes: The figure shows the features of the conditional distribution of r_{t+1}^f , as a function of r_t^f near the ZLB, given by the truncated autoregressive process estimated in equations (19) and (20). Data are at an annual frequency.

Figure 2. Characteristics of U.S. government liabilities in public hands


A. Average duration



B. Percentage with duration < 5 years




C. Share held by potential "arbitrage" investors



Notes: Includes coupon notes, bonds, and bills issued by the U.S. Treasury, less the amount held in the Federal Reserve's SOMA portfolio, plus reserves and currency in circulation, which are assumed to have a duration of zero. Panel C is the percentage of domestically, publicly held Treasury and bank reserves debt held by households, depository institutions, broker-dealers, and mutual funds. Sources: CRSP, Federal Reserve, Flow of Funds.

Figure 3. Solution convergence in the one-factor model




Notes: The graphs show how the solution algorithm converges in the number of iterations (top), number of nodes in the grid (middle), and truncation point of the state space (bottom). All calculations are illustrated at the average value of the short rate.


Figure 4. Risk and risk prices in the one-factor model

A. Short-rate sensitivity (B_{nt})

$z = 2.0$

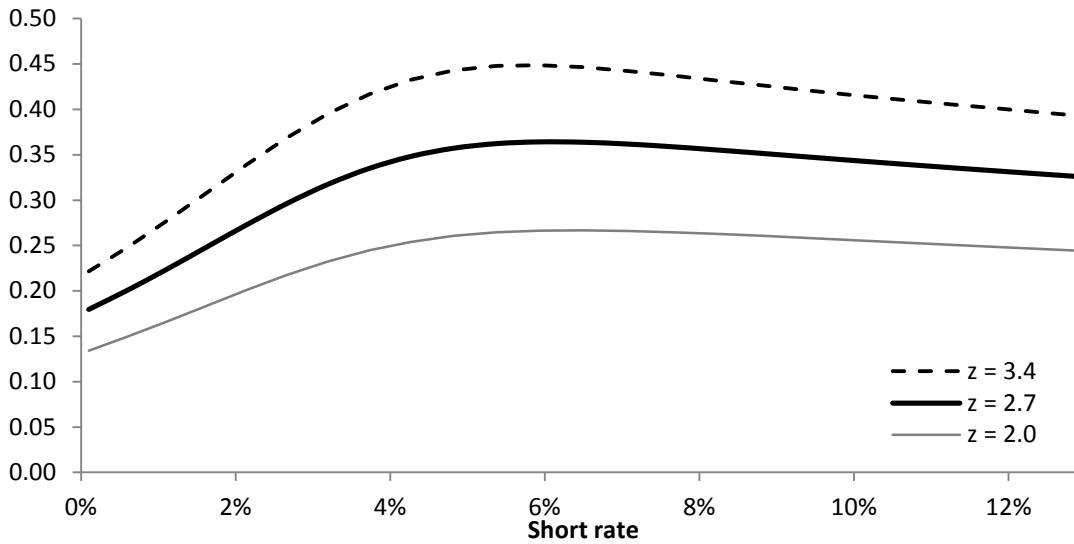


$z = 3.4$



Note: The figures plot the coefficient of the returns on bonds of each maturity (R_n) with respect to the price of the risk-free bond (p_1) for various values of the short-term interest rate (R_1) and aggregate duration (z) in the one-factor model.

B. Market price of risk (λ_t)




Note: The figure shows the price of short-rate risk (expected excess returns divided by their standard deviation) in the one-factor model across various levels of the short-term interest rate (R_1) and for various values of aggregate duration (z).


Figure 5. Moments of the yield curve in the one-factor model

Average yield curve

1971 – 2012




1985 - 2012




Standard deviation

1971 – 2012




1985 - 2012




Notes: Green lines show sample moments of monthly, zero-coupon nominal yields based on the Gurkaynak et al. (2007) data; blue lines show the corresponding moments generated by the one-factor portfolio-balance model, inputting the observed one-year rate.

Figure 6. An LSAP-type shift in the one-factor model


Duration distribution of outstanding debt



Yield curve



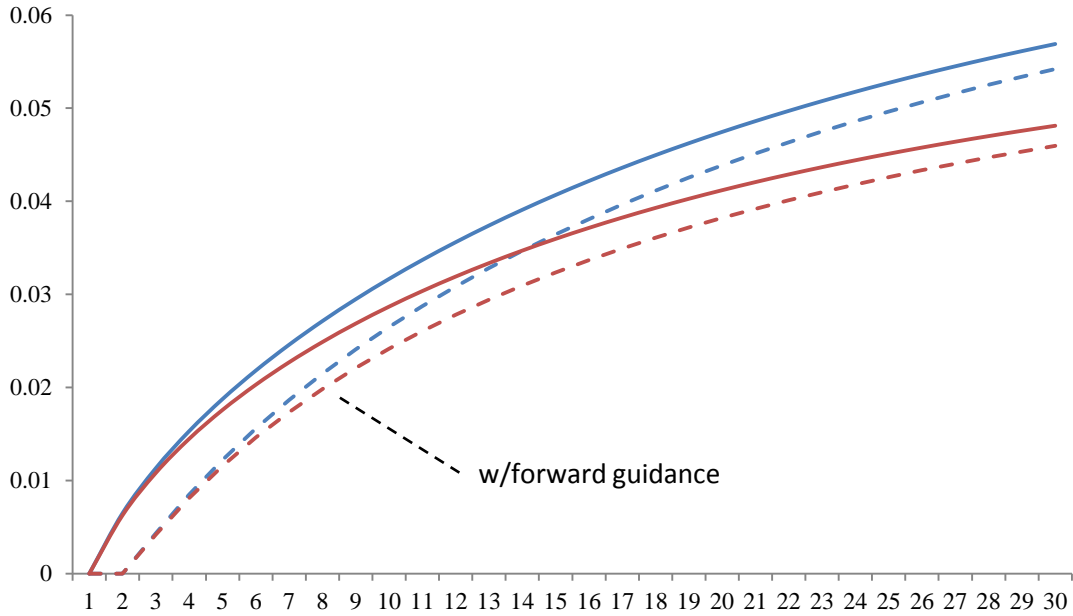
Return Volatility (One-Year)




Notes: The figure shows the effect of moving from an average duration of 2.7 to 2.0 years in the one-factor model, evaluated at a short rate of 5.8% (the sample average).

Figure 7. An LSAP-type shift in the one-factor model at the zero lower bound

Yield curve




Return Volatility (One-Year)




Notes: The figure shows the effect of moving from an average duration of 2.7 years (blue) to 2.0 years (red) in the one-factor model, evaluated at a short rate of 0%. The solid lines represent the case in which the short rate is expected to evolve according to its usual truncated-AR(1) process. The dashed lines (“forward guidance”) represent the case in which the short rate is anticipated to remain at 0% with certainty for one year and only then to follow its usual truncated-AR(1) process.

Figure 8. Effect of the maturity-distribution shape in the one-factor model

Duration distribution of outstanding debt



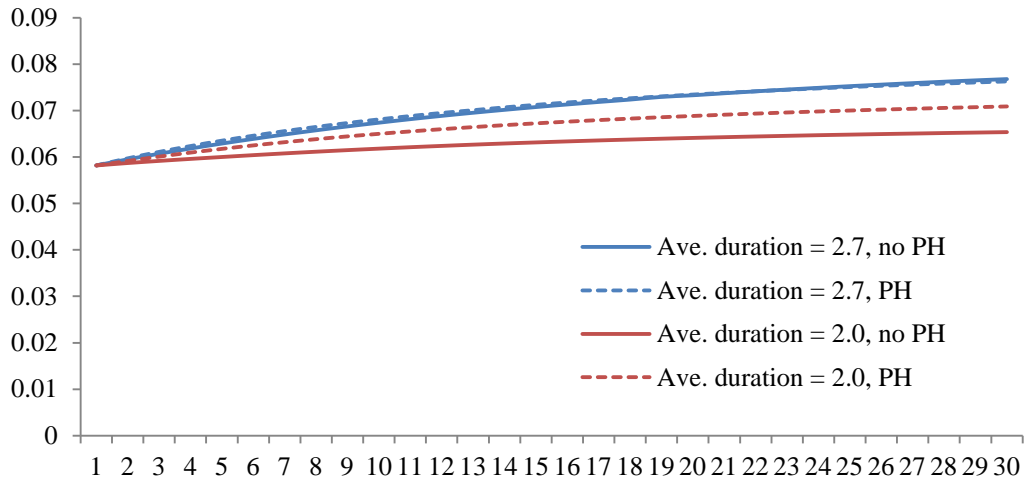
Yield curve




Notes: The top panel shows two possible distributions for the maturity of outstanding Treasury debt, one exponential and one nearly degenerate, both with means of 2.7 years. The bottom panel shows the corresponding yield curves generated by the one-factor model when the short rate is at its average value of 5.8%. In both cases, the risk-aversion coefficient is calibrated to match the average value of the 15-year yield.

Figure 9. Duration effects in the preferred-habitat model

Yield curve




Change in arbitrageur holdings when duration changes from 2.7 to 2.0



Notes: The top panel shows the yield curve produced by the one-factor model, without (solid lines) and with (dotted lines) the presence of preferred-habitat agents, evaluated at a short rate of 5.8%. The blue lines represent yields when the average duration of maturities outstanding is 2.7 years; the red lines represent yields when average duration is 2.0 years. The distribution of outstanding securities is assumed to be exponential in both cases. The bottom panel shows the difference in the arbitrage investor's holdings of each bond, as a fraction of the total bonds outstanding, again between the cases when the average duration is 2.7 and at 2.0. The solid bars show this difference under the baseline model without preferred-habitat agents; the dotted bars show this difference when the preferred-habitat agents are present.

Figure 10. Announcement and purchase effects of an LSAP-type shift in the one-factor model




Notes: The red line shows how the yield curve would change under the baseline one-factor model if aggregate duration moved unexpectedly and permanently from 2.7 and 2.0, evaluated at the average value of the short rate in both cases. The violet line shows the initial price reaction if this change is anticipated with certainty one-year before it occurs.


Figure 11. Impulse-response functions for the stochastic-supply model

At the sample mean

Response to short rate shock




Response to supply shock




At the steady state / ZLB

Response to short rate shock




Response to supply shock



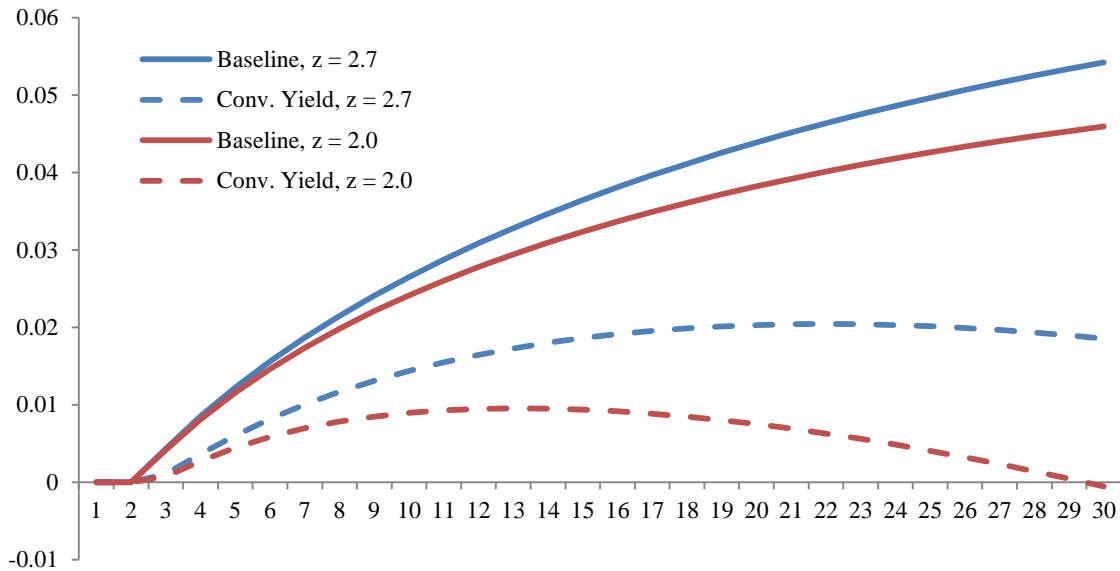
Notes: The figures show responses of the yield curve to one standard deviation positive shocks to the short-term interest rate and average Treasury duration in the two-factor model, evaluated at the sample average values (top panels) and at the steady-state values (lower panels) of the state variables. Maturities are plotted along the upper-left axis, and the lower-left axis represents calendar time (in years) after the shock.

Figure 12. Supply effects in the model with convenience yield

At the sample mean



At the steady state / ZLB





Notes: The figures shows the effect of moving from an average duration of 2.7 years (blue) to 2.0 years (red) in both the baseline one-factor model (solid lines) and the model including the convenience-yield term (dashed lines). The top panel shows the case in which the short rate is at its sample mean of 5.8%. Bottom panels shows the case in which the short rate is at 0% and is anticipated to remain there with certainty for one year and then to follow its usual truncated-AR(1) process. Parameters are calibrated to the sample moments of yields, as described in the text.

Figure 13. Impulse-response functions for the estimated three-factor model



At sample means

At LSAP-period values



Response to short-rate shock



Response to inflation shock




Response to supply shock





Notes: The figures show responses of the yield curve to one standard deviation positive shocks to the short-term interest rate, core PCE inflation, and average Treasury duration in the three-factor model, initialized at the 1971 - 2008 average values of the state variables (left panels) and at a set of state values representative of the 2008-2012 period (right panels). Maturities are plotted along the upper-left axis, and the lower-left axis represents calendar time (in years) after the shock.

Figure 14. Decomposition of the 10-year yield in the three-factor model

Yield components




Term premium components



Notes: The model is estimated over the period Dec. 1971 – Dec. 2008. The term premium is calculated as the difference between the model-implied rate and the expectations component. Term premium shares are calculated as the difference between the term premium implied in the full model and that implied in counterfactual models in which inflation or Treasury supply are always equal to their sample-average values with certainty. The “short-rate contribution” to the term premium is calculated holding both Treasury supply and inflation equal to their average.

Figure 15. The effect of the LSAPs in the three-factor model



Notes: The black line shows the initial impact of a duration shock on the scale of the LSAP programs in the estimated three-factor model, evaluated at initial values approximating conditions experienced during the LSAP period, assuming that forward guidance leads market participants to expect short rates to remain at zero for two years. The green line shows the effect of the same shock when the forward guidance is for one year, rather than two. The orange line shows the effect of a duration shock of double the size.

References

- Bauer, M. D. and Rudebusch, G. D., forthcoming. “The Signaling Channel for Federal Reserve Bond Purchases.” *International Journal of Central Banking*.
- Bernanke, B. S.; Reinhart, V., R.; and Sack, B. P., 2004. “Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment.” *Brookings Papers on Economic Activity* (2): 1 – 78.
- Bernanke, B., 2010. The Economic Outlook and Monetary Policy. Remarks delivered at Jackson Hole, Wyoming, August.
- Cahill, M.; D’Amico, S.; Li, C; and Sears, J., 2013. “Duration Risk versus Local Supply Channel in Treasury Yields: Evidence from the Federal Reserve’s Asset Purchase Announcements.” Working paper (April).
- Campbell, J. Y.; Sunderam, A.; and Viceira, L. M., 2009. “Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds.” Working paper, Harvard University, January.
- Chabot, B. and Herman, G., forthcoming. “A History of Large-Scale Asset Purchases before the Fed.” FRB Chicago *Economic Perspectives*.
- Cox, J. C.; Ingersoll, J. E.; and Ross, S. A., 1985. “A Theory of the Term Structure of Interest Rates.” *Econometrica* 53(2): 385-408.
- Cox, J. C.; Ingersoll, J. E.; and Ross, S. A., 1981. “A Re-Examination of Traditional Hypotheses about the Term Structure of Interest Rates.” *Journal of Finance* 36(4): 769-99.
- Culbertson, J. M., 1957. “The Term Structure of Interest Rates.” *Quarterly Journal of Economics* 71(4): 485-517.
- D’Amico, S.; English, W.; Lopez-Salido, D.; and Nelson, E., 2012. “The Federal Reserve’s Large-Scale Asset Purchase Programmes: Rationale and Effects.” *Economic Journal* 122 (564): F415-46.
- D’Amico, S. and King, T. B., 2013. “Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply.” *Journal of Financial Economics* 108(2): 425-48.
- Doh, T., 2010. “The Efficacy of Large-Scale Asset Purchases at the Zero Lower Bound.” FRB Kansas City *Economic Review* (Q2): 5-34.
- Eggertsson, G. and Woodford, M., 2003. “The Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity* (1): 139-211.
- Epstein, L. G. and Zin, S. E., 1989. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework.” *Econometrica* 57(4): 937-69.
- Femia, K.; Friedman, S.; and Sack, B., 2013. “The Effects of Policy Guidance on Perceptions of the Fed’s Reaction Function.” FRB New York Staff Report 652 (November).
- Frankel, J. A., 1985. “Portfolio Crowding-out, Empirically Estimated.” *Quarterly Journal of Economics* 100: 1041-65.
- Friedman, B. M., 1980. “The Effect of Shifting Wealth Ownership on the Term Structure of Interest Rates: The Case of Pensions.” *Quarterly Journal of Economics* 94(3): 567-90.

- Gagnon, J.; Raskin, M.; Remache, J.; and Sack, B., 2010. "Large-Scale Asset Purchases by the Federal Reserve: Did They Work?" *International Journal of Central Banking*.
- Greenwood, R.; Hanson, S. G.; and Stein, J. C., 2013. "A Comparative Advantage Approach to Government Debt Maturity." Working paper, September.
- Greenwood, R. and Vayanos, D., 2010. "Price Pressure in the Government Bond Market." *American Economic Review* 90(2): 585-90.
- Greenwood, R. and Vayanos, D., forthcoming. "Bond Supply and Excess Bond Returns." *Review of Financial Studies*.
- Gurkaynak, R. S.; Sack, B.; and Wright, J. H., 2006. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics* 54(8): 2291-394.
- Gurkaynak, R. S. and Wright, J. H., 2012. "Macroeconomics and the Term Structure." *Journal of Economic Literature* 50(2): 331-67.
- Hamilton, J. D. and Wu, J. C., 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit and Banking* 44(s1): 3 – 46.
- Hanson, S. G., 2012. "Mortgage Refinancing, the Quantity of Risk, and Bond Risk Premia." Harvard Business School, Working paper (September).
- Hanson, S. G. and Stein, J. C., 2012. "Monetary Policy and Long-Term Real Rates." FEDS Working Paper 2012-46.
- Ihrig, J.; Klee, E.; Li, C.; Schulte, B.; and Wei, M., 2012. "Expectations about the Federal Reserve's Balance Sheet and the Term Structure of Interest Rates." FEDS working paper 2012-57.
- Kim, D. H. and Wright, J. H., 2005. "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates." FEDS paper 2005-33.
- Krishnamurthy, A. and Vissing-Jorgensen, A., 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity* (2): 215-65.
- Krishnamurthy, A. and Vissing-Jorgensen, A., 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120(2): 233-67.
- Kuttner, K., 2006. "Can Central Banks Target Bond Prices?" NBER Working Paper 12454.
- Li, C. and Wei, M., 2012. "Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs." *International Journal of Central Banking*.
- Markowitz, H. M., 1952. "Portfolio Selection." *Journal of Finance* 7(1): 77-91.
- Meaning, J. and Zhu, F., 2011. "The Impact of Recent Central Bank Asset Purchase Programmes." *BIS Quarterly Review*, Dec.
- Piazzesi, M. and Schneider, M., 2007. "Asset Prices and Asset Quantities." *Journal of the European Economic Association* 5(2-3): 380-89.
- Piazzesi, M. and Schneider, M., 2008. "Bond Positions, Expectations, and the Yield Curve." FRB Atlanta working paper 2008-2.

- Roley, V. V., 1982. "The Effect of Federal Debt Management Policy on Corporate Bond and Equity Yields." *Quarterly Journal of Economics*
- Roll, R., 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(2): 129-76.
- Rosa, C., 2012. "How 'Unconventional' Are Large-Scale Asset Purchases? The Impact of Monetary Policy on Asset Prices." FRB NY Staff Report 560, May.
- Rudebusch, G. D.; Sack, B. P.; and Swanson, E. T., 2007. "Macroeconomic Implications of Changes in the Term Premium." FRB St. Louis *Review* 89(4): 241-69.
- Sack, B., 2009. Remarks at the Money Marketmakers of New York University, New York City, December 2: <http://www.newyorkfed.org/newsevents/speeches/2009/sac091202.html>.
- Svensson, L. E. O., 1994. "Estimating and Interpreting Forward Rates: Sweden 1992-4." NBER Working Paper 4871.
- Vasicek, O., 1977. "An Equilibrium Characterisation of the Term Structure." *Journal of Financial Economics* 5(2): 177-88.
- Vayanos, D. and Vila, J.-L.. 2009. "A Preferred-Habitat Model of the Term Structure of Interest Rates." Working paper.
- Wallace, N., 1981. "A Modigliani-Miller Theorem for Open-Market Operations." *American Economic Review* 71(3): 267-74.
- Weil, P., 1990. "Non-Expected Utility in Macroeconomics." *Quarterly Journal of Economics* 105(1): 29-42.
- Wright, J., 2012. "What Does Monetary Policy Do to Long-Term Interest Rates at the Zero Lower Bound?" Working paper (May).

Working Paper Series

A series of research studies on regional economic issues relating to the Seventh Federal Reserve District, and on financial and economic topics.

Comment on “Letting Different Views about Business Cycles Compete” <i>Jonas D.M. Fisher</i>	WP-10-01
Macroeconomic Implications of Agglomeration <i>Morris A. Davis, Jonas D.M. Fisher and Toni M. Whited</i>	WP-10-02
Accounting for non-annuitization <i>Svetlana Pashchenko</i>	WP-10-03
Robustness and Macroeconomic Policy <i>Gadi Barlevy</i>	WP-10-04
Benefits of Relationship Banking: Evidence from Consumer Credit Markets <i>Sumit Agarwal, Souphala Chomsisengphet, Chunlin Liu, and Nicholas S. Souleles</i>	WP-10-05
The Effect of Sales Tax Holidays on Household Consumption Patterns <i>Nathan Marwell and Leslie McGranahan</i>	WP-10-06
Gathering Insights on the Forest from the Trees: A New Metric for Financial Conditions <i>Scott Brave and R. Andrew Butters</i>	WP-10-07
Identification of Models of the Labor Market <i>Eric French and Christopher Taber</i>	WP-10-08
Public Pensions and Labor Supply Over the Life Cycle <i>Eric French and John Jones</i>	WP-10-09
Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash <i>Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein</i>	WP-10-10
Prenatal Sex Selection and Girls’ Well-Being: Evidence from India <i>LuoJia Hu and Analía Schlosser</i>	WP-10-11
Mortgage Choices and Housing Speculation <i>Gadi Barlevy and Jonas D.M. Fisher</i>	WP-10-12
Did Adhering to the Gold Standard Reduce the Cost of Capital? <i>Ron Alquist and Benjamin Chabot</i>	WP-10-13
Introduction to the <i>Macroeconomic Dynamics</i> : Special issues on money, credit, and liquidity <i>Ed Nosal, Christopher Waller, and Randall Wright</i>	WP-10-14
Summer Workshop on Money, Banking, Payments and Finance: An Overview <i>Ed Nosal and Randall Wright</i>	WP-10-15
Cognitive Abilities and Household Financial Decision Making <i>Sumit Agarwal and Bhashkar Mazumder</i>	WP-10-16

Working Paper Series *(continued)*

Complex Mortgages <i>Gene Amromin, Jennifer Huang, Clemens Sialm, and Edward Zhong</i>	WP-10-17
The Role of Housing in Labor Reallocation <i>Morris Davis, Jonas Fisher, and Marcelo Veracierto</i>	WP-10-18
Why Do Banks Reward their Customers to Use their Credit Cards? <i>Sumit Agarwal, Sujit Chakravorti, and Anna Lunn</i>	WP-10-19
The impact of the originate-to-distribute model on banks before and during the financial crisis <i>Richard J. Rosen</i>	WP-10-20
Simple Markov-Perfect Industry Dynamics <i>Jaap H. Abbring, Jeffrey R. Campbell, and Nan Yang</i>	WP-10-21
Commodity Money with Frequent Search <i>Ezra Oberfield and Nicholas Trachter</i>	WP-10-22
Corporate Average Fuel Economy Standards and the Market for New Vehicles <i>Thomas Klier and Joshua Linn</i>	WP-11-01
The Role of Securitization in Mortgage Renegotiation <i>Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff</i>	WP-11-02
Market-Based Loss Mitigation Practices for Troubled Mortgages Following the Financial Crisis <i>Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff</i>	WP-11-03
Federal Reserve Policies and Financial Market Conditions During the Crisis <i>Scott A. Brave and Hesna Genay</i>	WP-11-04
The Financial Labor Supply Accelerator <i>Jeffrey R. Campbell and Zvi Hercowitz</i>	WP-11-05
Survival and long-run dynamics with heterogeneous beliefs under recursive preferences <i>Jaroslav Borovička</i>	WP-11-06
A Leverage-based Model of Speculative Bubbles (Revised) <i>Gadi Barlevy</i>	WP-11-07
Estimation of Panel Data Regression Models with Two-Sided Censoring or Truncation <i>Sule Alan, Bo E. Honoré, LuoJia Hu, and Søren Leth-Petersen</i>	WP-11-08
Fertility Transitions Along the Extensive and Intensive Margins <i>Daniel Aaronson, Fabian Lange, and Bhashkar Mazumder</i>	WP-11-09
Black-White Differences in Intergenerational Economic Mobility in the US <i>Bhashkar Mazumder</i>	WP-11-10

Working Paper Series *(continued)*

Can Standard Preferences Explain the Prices of Out-of-the-Money S&P 500 Put Options? <i>Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein</i>	WP-11-11
Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture <i>Ezra Oberfield</i>	WP-11-12
Equilibrium Bank Runs Revisited <i>Ed Nosal</i>	WP-11-13
Are Covered Bonds a Substitute for Mortgage-Backed Securities? <i>Santiago Carbó-Valverde, Richard J. Rosen, and Francisco Rodríguez-Fernández</i>	WP-11-14
The Cost of Banking Panics in an Age before “Too Big to Fail” <i>Benjamin Chabot</i>	WP-11-15
Import Protection, Business Cycles, and Exchange Rates: Evidence from the Great Recession <i>Chad P. Bown and Meredith A. Crowley</i>	WP-11-16
Examining Macroeconomic Models through the Lens of Asset Pricing <i>Jaroslav Borovička and Lars Peter Hansen</i>	WP-12-01
The Chicago Fed DSGE Model <i>Scott A. Brave, Jeffrey R. Campbell, Jonas D.M. Fisher, and Alejandro Justiniano</i>	WP-12-02
Macroeconomic Effects of Federal Reserve Forward Guidance <i>Jeffrey R. Campbell, Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano</i>	WP-12-03
Modeling Credit Contagion via the Updating of Fragile Beliefs <i>Luca Benzoni, Pierre Collin-Dufresne, Robert S. Goldstein, and Jean Helwege</i>	WP-12-04
Signaling Effects of Monetary Policy <i>Leonardo Melosi</i>	WP-12-05
Empirical Research on Sovereign Debt and Default <i>Michael Tomz and Mark L. J. Wright</i>	WP-12-06
Credit Risk and Disaster Risk <i>François Gourio</i>	WP-12-07
From the Horse’s Mouth: How do Investor Expectations of Risk and Return Vary with Economic Conditions? <i>Gene Amromin and Steven A. Sharpe</i>	WP-12-08
Using Vehicle Taxes To Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden <i>Thomas Klier and Joshua Linn</i>	WP-12-09
Spending Responses to State Sales Tax Holidays <i>Sumit Agarwal and Leslie McGranahan</i>	WP-12-10

Working Paper Series *(continued)*

Micro Data and Macro Technology <i>Ezra Oberfield and Devesh Raval</i>	WP-12-11
The Effect of Disability Insurance Receipt on Labor Supply: A Dynamic Analysis <i>Eric French and Jae Song</i>	WP-12-12
Medicaid Insurance in Old Age <i>Mariacristina De Nardi, Eric French, and John Bailey Jones</i>	WP-12-13
Fetal Origins and Parental Responses <i>Douglas Almond and Bhashkar Mazumder</i>	WP-12-14
Repos, Fire Sales, and Bankruptcy Policy <i>Gaetano Antinolfi, Francesca Carapella, Charles Kahn, Antoine Martin, David Mills, and Ed Nosal</i>	WP-12-15
Speculative Runs on Interest Rate Pegs The Frictionless Case <i>Marco Bassetto and Christopher Phelan</i>	WP-12-16
Institutions, the Cost of Capital, and Long-Run Economic Growth: Evidence from the 19th Century Capital Market <i>Ron Alquist and Ben Chabot</i>	WP-12-17
Emerging Economies, Trade Policy, and Macroeconomic Shocks <i>Chad P. Bown and Meredith A. Crowley</i>	WP-12-18
The Urban Density Premium across Establishments <i>R. Jason Faberman and Matthew Freedman</i>	WP-13-01
Why Do Borrowers Make Mortgage Refinancing Mistakes? <i>Sumit Agarwal, Richard J. Rosen, and Vincent Yao</i>	WP-13-02
Bank Panics, Government Guarantees, and the Long-Run Size of the Financial Sector: Evidence from Free-Banking America <i>Benjamin Chabot and Charles C. Moul</i>	WP-13-03
Fiscal Consequences of Paying Interest on Reserves <i>Marco Bassetto and Todd Messer</i>	WP-13-04
Properties of the Vacancy Statistic in the Discrete Circle Covering Problem <i>Gadi Barlevy and H. N. Nagaraja</i>	WP-13-05
Credit Crunches and Credit Allocation in a Model of Entrepreneurship <i>Marco Bassetto, Marco Cagetti, and Mariacristina De Nardi</i>	WP-13-06

Working Paper Series *(continued)*

Financial Incentives and Educational Investment: The Impact of Performance-Based Scholarships on Student Time Use <i>Lisa Barrow and Cecilia Elena Rouse</i>	WP-13-07
The Global Welfare Impact of China: Trade Integration and Technological Change <i>Julian di Giovanni, Andrei A. Levchenko, and Jing Zhang</i>	WP-13-08
Structural Change in an Open Economy <i>Timothy Uy, Kei-Mu Yi, and Jing Zhang</i>	WP-13-09
The Global Labor Market Impact of Emerging Giants: a Quantitative Assessment <i>Andrei A. Levchenko and Jing Zhang</i>	WP-13-10
Size-Dependent Regulations, Firm Size Distribution, and Reallocation <i>François Gourio and Nicolas Roys</i>	WP-13-11
Modeling the Evolution of Expectations and Uncertainty in General Equilibrium <i>Francesco Bianchi and Leonardo Melosi</i>	WP-13-12
Rushing into American Dream? House Prices, Timing of Homeownership, and Adjustment of Consumer Credit <i>Sumit Agarwal, Luojia Hu, and Xing Huang</i>	WP-13-13
The Earned Income Tax Credit and Food Consumption Patterns <i>Leslie McGranahan and Diane W. Schanzenbach</i>	WP-13-14
Agglomeration in the European automobile supplier industry <i>Thomas Klier and Dan McMillen</i>	WP-13-15
Human Capital and Long-Run Labor Income Risk <i>Luca Benzoni and Olena Chyruk</i>	WP-13-16
The Effects of the Saving and Banking Glut on the U.S. Economy <i>Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti</i>	WP-13-17
A Portfolio-Balance Approach to the Nominal Term Structure <i>Thomas B. King</i>	WP-13-18