

The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment*

James D. Hamilton[†]

Department of Economics

University of California, San Diego

Jing (Cynthia) Wu[‡]

Department of Economics

University of California, San Diego

August 25, 2010

Revised: February 12, 2011

Abstract

This paper reviews alternative options for monetary policy when the short-term interest rate is at the zero lower bound and develops new empirical estimates of the effects of the maturity structure of publicly held debt on the term structure of interest rates. We use a model of risk-averse arbitrageurs to develop measures of how the maturity structure of debt held by the public might affect the pricing of level, slope and curvature term-structure risk. We find these Treasury factors historically were quite helpful for predicting both yields and excess returns over 1990-2007. The historical correlations are consistent with the claim that if in December of 2006, the Fed were to have sold off all its Treasury holdings of less than one-year maturity (about \$400 billion) and use the proceeds to retire Treasury debt from the long end, this might have resulted in a 14-basis-point drop in the 10-year rate and an 11-basis-point increase in the 6-month rate. We also develop a description of how the dynamic behavior of the term structure of interest rates changed after hitting the zero lower bound in 2009. Our estimates imply that at the zero lower bound, such a maturity swap would have the same effects as buying \$400 billion in long-term maturities outright with newly created reserves, and could reduce the 10-year rate by 13 basis points without raising short-term yields.

*We thank Christiane Baumeister for assistance with obtaining some of the data for this project, and thank Michael Bauer, Gauti Eggertsson, Jeff Hallman, Eric Swanson, Dimitri Vayanos, Kenneth West, Michael Woodford, and seminar and conference participants at the Federal Reserve Board, Federal Reserve Bank of Boston, and Federal Reserve Bank of San Francisco for helpful comments on earlier versions of this paper.

[†]jhamilton@ucsd.edu

[‡]jingwu@ucsd.edu

1 Introduction.

The key instrument of monetary policy is the interest rate on overnight loans between banks, which in normal times is quite sensitive to the quantity of excess reserves. However, since December 2008, the Fed's target for the fed funds rate has been essentially zero. The level of reserves, which had typically been around \$10 billion prior to the financial crisis, has been maintained in the neighborhood of a trillion dollars. Trying to lower the short-term interest rate or increase the volume of reserves any further offers little promise of boosting aggregate demand. With the Fed's traditional tools incapable of providing further stimulus to the economy, it is of considerable interest to ask what other options might be available to the central bank.

Our study begins by briefly reviewing some of the available options and the Fed's experience with using them. That analysis leads us to focus on one strategy in particular, which is to try to influence the term structure of interest rates through the maturity structure of securities acquired by open-market purchases.

A number of previous studies have reported evidence that the relative supplies of Treasury securities of different maturities are correlated with yield spreads; see for example Roley (1982), Bernanke et al. (2004), Kuttner (2006), Gagnon et al. (2010), Doh (2010), Greenwood and Vayanos (2010), D'Amico and King (2010), and Swanson (2011).¹ But using those correlations to infer potential effects of nonstandard open-market operations raises questions from the perspective of both economic theory, in terms of the proposed mechanism whereby the effects

¹Other closely related research includes Krishnamurthy and Vissing-Jorgensen (2010), Baumeister and Benati (2010), Kitchen and Chinn (2010), and Hancock and Passmore (2011).

could possibly be generated, as well as from the perspective of econometric methodology, in terms of whether it is reasonable to place a causal interpretation on the correlations. Our paper makes contributions in both areas.

Our theoretical motivation follows Vayanos and Vila (2009), who developed a promising framework for understanding how the supplies of assets of different maturities might influence their respective yields. Vayanos and Vila postulate the existence of two groups of investors. The willingness of preferred-habitat investors to buy securities of maturity n is presumed to be an increasing function of the yield on that asset. A second group, known as arbitrageurs, is willing to hold any assets based on a simple tradeoff between expected return and risk. The behavior of the second group generates no-arbitrage conditions relating the yields on different securities.

Our empirical analysis follows Doh (2010) and Greenwood and Vayanos (2010) in using the Vayanos and Vila (2009) framework to try to quantify the ability of nonstandard open-market operations to change the yields on assets of different maturities. We differ from these earlier researchers in making more use of the details of the framework to inform the empirical estimates, developing a discrete-time version of the model and relating it directly to maximum-likelihood estimates of the dynamic behavior of the term structure of interest rates. We develop specific historical measures of how the maturity structure of debt issued to the public might be expected to affect the pricing of level, slope, and curvature risk according to this framework, and show that our inferred Treasury risk factors were historically quite helpful in predicting yields and excess returns. For example, we find that over 1990-2007, the excess one-year return from holding 2-year Treasuries over 1-year Treasuries can be predicted

with an R^2 of 71% on the basis of traditional term-structure factors along with our proposed Treasury risk factors.

One of the challenges for estimating potential policy effects on the basis of historical correlations is the problem of endogeneity, in that the correlation between bond supplies and interest rates may reflect the response of the Treasury or the Fed to interest rates. We try to minimize this endogeneity bias by looking at forecasting rather than contemporaneous regressions and including the current level, slope, and curvature as additional explanatory variables in the regression. Our impact estimates are based on the incremental contribution of the Treasury maturity structure to a one-month-ahead forecast of interest rates beyond the information already contained in the current term structure, so that insofar as the maturities of debt issued by the Treasury or purchased by the Fed are responding to current interest rates, that response could not account for our estimated effects. Our dynamic formulation also avoids the potential spurious regression problem that could arise in simple contemporaneous regressions that make no allowance for near-unit-root dynamics.

We use our estimated forecasting relations to analyze the outcome of the following policy change. Suppose the Federal Reserve were to sell off all of its holdings of Treasury securities of less than one-year duration, and use the proceeds to buy up all the outstanding Treasury debt it could at the long end of the yield curve. For example, in 2006 this would have involved a \$400 billion asset swap that would have retired all Treasury debt of more than 10-years duration. Our estimates imply that, in an environment not affected by the zero lower bound, this would have decreased the 10-year yield by 14 basis points and increased the 6-month yield by 11 basis points.

We next develop a framework for analyzing the behavior of interest rates when the short-term interest rate hits the zero lower bound. Our basic approach is to postulate that movements in longer-term yields in such a setting are explained by arbitrageurs' assumption that the economy will eventually break out of the zero lower bound, and that, once it does, short-term interest rates would again fluctuate in response to the same kind of forces as they did historically. We propose a very parsimonious description in which arbitrageurs assume that, apart from a possible downward shift in the average level, the post-ZLB dynamics will be the same as those observed in the pre-ZLB experience. Given an exogenous probability of exiting the ZLB in any given period, we then develop a no-arbitrage theory of how the term structure evolves dynamically when at the ZLB. We find this model provides a reasonable empirical description of the behavior of the term structure during 2009 and 2010.

We then use this model to revisit the question analyzed for the pre-2007 data. We find that, at the ZLB, an asset swap could continue to depress long-term yields by the same amount that it would in normal times, without producing any rise in short-term yields. Thus, whereas swapping short-term for long-term assets has no consequences for the overall level of interest rates in normal times, it is an available tool for lowering the overall level at the ZLB. Moreover, since at the ZLB newly created reserves are essentially equivalent to short-term T-bills, direct large-scale asset purchases are a feasible tool that the Fed could use to lower long-term interest rates when at the ZLB.

The plan of the paper is as follows. Section 2 reviews alternative mechanisms whereby monetary policy might still be able to influence interest rates for an economy at the ZLB, and explains our reason for focusing in particular on the possible effects arising through changes

in the maturity composition of outstanding debt. Section 3 develops a discrete-time version of the Vayanos and Vila (2009) framework for analyzing the nature of preferred-habitat asset markets and the pricing of term-structure risk. Section 4 provides details of our method for obtaining maximum-likelihood estimates of parameters, while Section 5 reviews the data set assembled for this study. In Section 6 we analyze the effects of nonstandard open-market operations in an environment of fluctuating short-term interest rates, while Section 7 extends the analysis to an economy in which the short-term rate is temporarily stuck at some lower bound. We briefly compare our results with other recent estimates in Section 8. Section 9 concludes.

2 Options for monetary stimulus at the zero lower bound.

When the short-term interest rate gets all the way to zero, an open-market purchase of a short-term Treasury security with newly created base money represents an exchange of essentially equivalent assets. Such an exchange is obviously incapable of lowering the short-term rate any further, and it's not clear how the exchange could affect any economic magnitude of interest. Eggertsson and Woodford (2003) describe this as a situation in which the demand for money is completely satiated. With over a trillion dollars in excess reserves, the United States presently appears to be well past the satiation point for Federal Reserve deposits.

Even if the demand for reserve balances is presently satiated, as long as the situation is not permanent, at some future date the Fed will regain its ability to influence overnight rates. Thus even at the zero lower bound, Krugman (1998) and Eggertsson and Woodford

(2003) propose that the central bank could mitigate the current problems by successfully communicating its commitment to reverse any decreases in the price level, embracing the higher future inflation rates necessary to achieve that. Although such a strategy holds appeal in theory, in practice it appears to be quite hard to achieve. For example, the top panel of Figure 1 plots the 5-year expected inflation rate implied by the difference between nominal and inflation-indexed U.S. Treasuries. This plunged in the fall of 2008, and has yet to recover to its pre-crisis levels. Five-year expected inflation has also declined according to the average response to the Survey of Professional Forecasters (bottom panel). The failure of the Fed to follow the theoretical policy prescription of trying to increase inflationary expectations in response to the crisis is not so much an indictment of the Fed as it is a clear demonstration that these expectations are far more difficult to control in practice than simple theoretical treatments might sometimes suppose.

If buying T-bills with newly created reserves has no effect, the Fed could buy some other assets which clearly are not perfect substitutes for cash. One obvious class of assets to consider purchasing would be those denominated in foreign currencies. If the Fed announced a commitment to buy such assets without limit until the dollar depreciated, it is hard to imagine real-world market forces that could prevent the goal from being achieved. In terms of theoretical models, the ability of the Fed to make good on such a commitment could arise from a portfolio balance effect (McCallum (2000)), or the announcement could serve as an expectations coordinating mechanism (Svensson (2001)). In either case, it certainly seems one practical tool for preventing deflation even if no others are available.

If the private sector were indeed indifferent between holding freely created reserves and

long-term Treasury debt, one wonders why the Federal Reserve wouldn't want to buy up the entire stock of outstanding public debt, thereby eliminating the need for future taxes to service that debt. A related question is why the government would choose to use taxes rather than money creation as the means to pay for any of its current or projected future expenditures. Auerbach and Obstfeld (2005) explore the possible expansionary effects resulting from reducing the distortionary effect of taxes.

In the actual U.S. experience over 2008-2010, the Federal Reserve doubled the size of its balance sheet, buying two broad classes of assets (see Figure 2). In the first year of the crisis, the Fed was aggressively extending loans through a variety of new facilities such as the Term Auction Facility (essentially a term discount window open to all depository institutions on an auction basis), foreign currency swaps (used to assist foreign central banks in lending dollars), and the Commercial Paper Funding Facility (which helped provide loans for issuers of commercial paper). These measures could matter both in terms of making these markets more liquid (in the sense of reducing bid-ask spreads) as well as potentially absorbing some default risk onto the Fed's balance sheet. Christensen et al. (2009), McAndrews et al. (2008), Taylor and Williams (2009), Adrian et al. (2010) and Duygan-Bump et al. (2010) provide empirical assessments of the effectiveness of such measures.

Beginning in March 2009, these lending facilities began to be unwound and replaced by the gradual purchase of up to \$1.1 trillion in mortgage-backed securities, along with \$160B in agency debt and \$300B in new holdings of Treasury bonds with greater than one year maturity. Although rates on MBS and agency debt might be argued to include a default premium, with the de facto nationalization of Fannie and Freddie, it seems most natural to regard the effect

of these purchases as coming from a change in the relative supply of longer-term assets.² As this has become the most important tool going forward, our analysis in this paper focuses on the potential of such operations to alter the term structure of interest rates.

The mechanism by which such asset purchases might have an effect is very different from that characterizing traditional open-market operations. The Federal Reserve is the monopoly supplier of reserves held by depository institutions and currency held by the public, and the supply it creates of these assets unquestionably has consequences under normal economic conditions. However, when the demand for such assets is satiated, it is not clear that anything the Fed does could affect the pricing kernel determining other yields. For example, Eggertsson and Woodford (2003) elaborate conditions under which an open-market purchase of any asset whatever would have zero consequences for variables such as real output and the price level provided that it has no implications for the future conduct of monetary or fiscal policy. Woodford (2010) notes that if the operations have no affect on the asset's state-contingent income stream or on the state-contingent aggregate supply of goods available for consumption, they should have no effect on the price of the asset.

Certainly from the perspective of an individual investor, a 10-year Treasury bond has different risk characteristics from a 6-month T-bill, and these differences get priced by the market. If an individual investor changes her relative holdings of these assets, she perceives herself to have a different risk exposure, and perceives the U.S. Treasury to be the counterparty. By focusing on this aspect of bond pricing, as our paper does, our answer to the Eggertsson-Woodford critique is that, if the government changes the maturity structure of its outstanding

²Hancock and Passmore (2011) nevertheless found evidence that the MBS purchases did lower the premium on MBS relative to Treasury securities.

debt, it is in fact committing to a different state-contingent path for spending, taxes, or inflation in order to maintain intertemporal budget balance under the altered debt structure.

Changing the risk exposure of the holders of government debt appears to be the key mechanism whereby changes in the maturity structure of government debt would be able to influence the term structure of interest rates in a class of formal descriptions of the portfolio-balance effect. We illustrate this point by developing in the following section a discrete-time version of the framework recently proposed by Vayanos and Vila (2009). This exercise both clarifies the mechanism whereby nonstandard open market operations could affect the term structure, and also suggests particular empirical measures that prove to be helpful for quantifying plausible sizes for these effects.

3 Preferred-habitat investing and market arbitrage.

Vayanos and Vila (2009) propose that the investors we will refer to as “arbitrageurs” care only about the mean and variance of $r_{t,t+1}$, the rate of return between t and $t + 1$ on their total portfolio³:

$$E_t(r_{t,t+1}) - (\gamma/2)\text{Var}_t(r_{t,t+1}). \tag{1}$$

If y_{1t} denotes the return on a risk-free asset, arbitrageurs will choose portfolio weights such

³Vayanos and Vila (2009) assume that arbitrageurs maximize an objective function that is quadratic in the change in wealth rather than in the rate of return as here. Although their specification may have more theoretical appeal, their parameterization would be more difficult to bring to the data in the manner we propose here for an economy in which there is a trend in the level of wealth.

that for any asset with a risky yield $r_{i,t,t+1}$,

$$y_{1t} = E_t(r_{i,t,t+1}) - \gamma \vartheta_{it} \quad (2)$$

where ϑ_{it} is $(1/2)$ the derivative of total portfolio variance with respect to holdings of asset i .

Consider a pure-discount n -period bond that is free of default risk, the log of whose price at date t (denoted p_{nt}) is conjectured to be an affine function of a vector of J different macroeconomic factors (denoted f_t),

$$p_{nt} = \bar{a}_n + \bar{b}'_n f_t. \quad (3)$$

The risk-free one-period rate is a function of the same factors,

$$y_{1t} = a_1 + b'_1 f_t,$$

where $y_{1t} = -p_{1t}$, $a_1 = -\bar{a}_1$, and $b_1 = -\bar{b}_1$. Although these bonds have no default risk, the future pricing factors f_{t+s} are not known with certainty at date t , and so there is an uncertain one-period holding yield associated with buying the n -period bond at date t and selling the resulting $(n - 1)$ -period bond at date $t + 1$ given by

$$r_{n,t,t+1} = \exp\left(\bar{a}_{n-1} + \bar{b}'_{n-1} f_{t+1} - \bar{a}_n - \bar{b}'_n f_t\right) - 1. \quad (4)$$

Suppose that the pricing factors follow a VAR(1) process,

$$f_{t+1} = c + \rho f_t + \Sigma u_{t+1} \quad (5)$$

with $u_t \sim \text{i.i.d. } N(0, I_J)$, and that the arbitrageurs hold a fraction z_{nt} of their portfolio in the bond of maturity n , so that the return on their portfolio is given by

$$r_{t,t+1} = \sum_{n=1}^N z_{nt} r_{n,t,t+1}.$$

Then, as we detail in Appendix A, an approximation to the portfolio optimization problem results in the following implication of (2) for each maturity n :

$$-\bar{a}_1 - \bar{b}'_1 f_t = \bar{a}_{n-1} + \bar{b}'_{n-1} (c + \rho f_t) + (1/2) \bar{b}'_{n-1} \Sigma \Sigma' \bar{b}_{n-1} - \bar{a}_n - \bar{b}'_n f_t - \bar{b}'_{n-1} \Sigma \lambda_t \quad (6)$$

$$\lambda_t = \gamma \Sigma' d_t \quad (7)$$

$$d_t = \sum_{n=2}^N z_{nt} \bar{b}_{n-1}. \quad (8)$$

If the number of maturities N is greater than the number of factors J , equation (6) implies a set of restrictions that bond prices must satisfy as a result of the actions of arbitrageurs, who will price factor j risk the same way no matter which bonds it may be reflected in.

Vayanos and Vila close the model by postulating that other credit market participants may have a particular preference for bonds of a given maturity. They present examples in which the borrowing demand from these participants for bonds of maturity n , denoted ξ_{nt} , is

a decreasing affine function of the yield y_{nt} . In our application, we will express these demands relative to W_t , the net wealth of the arbitrageurs:

$$\xi_{nt}/W_t = \zeta_{nt} - \alpha_n y_{nt}.$$

Thus ζ_{nt} reflects the overall level of preferred-habitat borrowing of bonds of maturity n and α_n the sensitivity of this demand to the interest rate. Equilibrium then requires that the net borrowing by the preferred-habitat sector equals the net lending from the arbitrage sector:

$$z_{nt} = \zeta_{nt} - \alpha_n y_{nt}. \quad (9)$$

Suppose that ζ_{nt} is also an affine function of f_t . We show in Appendix B that in equilibrium,

$$\lambda_t = \lambda + \Lambda f_t. \quad (10)$$

Substituting (10) into (6), we see that

$$\bar{b}'_n = \bar{b}'_{n-1} \rho^Q - b'_1 \quad (11)$$

$$\rho^Q = \rho - \Sigma \Lambda \quad (12)$$

$$\bar{a}_n = \bar{a}_{n-1} + \bar{b}'_{n-1} c^Q + (1/2) \bar{b}'_{n-1} \Sigma \Sigma' \bar{b}_{n-1} - a_1 \quad (13)$$

$$c^Q = c - \Sigma \lambda. \quad (14)$$

4 Estimation of Affine-Term-Structure Models.

Equations (11) through (14) will be recognized as the no-arbitrage conditions for a standard affine-term-structure model (e.g., equations (17) in Ang and Piazzesi, 2003). Thus the Vayanos-Vila formulation can be viewed as one explanation for the origins of affine prices of risk. In this section we describe how we estimated parameters for this class of models; for further details see Appendix C.

Let y_{nt} denote the yield and p_{nt} the log price on an n -period pure discount bond, which are related by $y_{nt} = -n^{-1}p_{nt}$. From (3),

$$y_{nt} = a_n + b_n' f_t \quad (15)$$

with $a_n = -\bar{a}_n/n$ and $b_n = -\bar{b}_n/n$. In the models we estimate, the factors f_t are represented by a $(J \times 1)$ vector of observed variables, whose dynamic parameters c and ρ can be obtained from OLS estimation of (5). We suppose that we have available a set of M different observed yields $Y_{2t} = (y_{n_1,t}, y_{n_2,t}, \dots, y_{n_M,t})'$ whose values differ from the theoretical prediction (15) by measurement error

$$Y_{2t} = A + Bf_t + \Sigma_e u_t^e \quad (16)$$

with $u_t^e \sim N(0, I_M)$. We assume that the measurement error u_t^e is independent of the factor innovation u_t in (5), but otherwise the structure of Σ_e does not affect the estimation procedure—full-information maximum-likelihood estimates of all parameters other than Σ_e will be numerically identical regardless of whether the matrix Σ_e is assumed to be diagonal.

Our estimates come from the minimum-chi-square estimation algorithm proposed by Hamil-

ton and Wu (2010) which allows OLS to do the work of maximizing the joint likelihood function and uses the theoretical model to translate those OLS estimates back into the asset-pricing parameters of interest. Note that the structure of (5) and (16) implies that OLS equation by equation is the most efficient procedure for estimation of these reduced-form parameters. In the special case of a just-identified model, in which the number of observed yields M is one more than the number of factors J , there is an exact solution for the parameters of interest in terms of these OLS coefficients, and the resulting estimates are numerically identical to those that would be obtained by maximization of the joint likelihood function $f(Y_{2T}, f_T, Y_{2,T-1}, f_{T-1}, \dots, Y_{21}, f_1 | Y_{20}, f_0)$ with respect to the parameters of the affine-term-structure model, namely, $c, \rho, \Sigma, c^Q, \rho^Q, b_1, a_1$ and Σ^e .

Among other advantages, this approach allows us to recognize instantly whether estimates represent a local rather than a global maximum to the likelihood function, and makes it feasible to calculate small-sample confidence intervals for any function of the parameters of interest, by simulating a thousand different samples for $\{f_t, Y_{2t}\}_{t=1}^T$ from a postulated structure and calculating the estimates that result from the proposed procedure on each separate artificial sample.

5 Data.

Our baseline estimates use weekly observations for y_{nt} , based on constant-maturity Treasury yields as of Friday or the last business day of the week as reported in the FRED database of

the Federal Reserve Bank of St. Louis.⁴ We supplement this with monthly analysis of holding yields on securities of nonstandard maturities, for which we construct constant-maturity yields from the daily term-structure parameterization of Gürkaynak et al. (2007) as of the last day of the month.⁵

We also constructed estimates of the face value of outstanding U.S. Treasury debt at each weekly maturity as of the end of each month between January 1990 and December 2009 as detailed in Appendix E. For purposes of the pure theory sketched above, we would want to interpret each semiannual coupon on a given bond as its own separate zero-coupon security (paying $\$C$ at some time $t + s$) and construct the market value of the bond as the sum of the market value of its individual components, each coupon viewed as a separate pure-discount bond. However, converting the face value into a market value by this device would be quite unsatisfactory for our larger purpose of identifying exogenous sources of variation in the supply of outstanding securities at different maturities. The true market value of a given security would be highly endogenous with respect to changes in interest rates, whereas the face value, by construction, is not.⁶ Note moreover that, when issued, the face value of the original coupon bond should be close to the market value of the sum of its individual stripped components. For these reasons, we regard the face value as reported by the Treasury and

⁴The 30-year yields are unavailable for 2002/2/19 to 2006/2/8. Over this interval we used instead the 20-year rate minus 0.21, which is the amount by which the 20-year rate exceeded the 30-year rate both immediately before and after the gap.

⁵Specifically, we calculated y_{nt} from their equations (6) and (9) as

$$y_{nt} = \beta_{0t} + n^{-1}\beta_{1t}\tau_{1t}[1 - \exp(-n/\tau_{1t})] + \beta_{2t}\tau_{1t}\{1 - [1 + (n/\tau_{1t})]\exp(-n/\tau_{1t})\} \\ + \beta_{3t}\tau_{2t}\{1 - [1 + (n/\tau_{2t})]\exp(-n/\tau_{2t})\}$$

using daily values for the parameters $\{\beta_{0t}, \beta_{1t}, \beta_{2t}, \tau_{1t}, \tau_{2t}\}$ downloaded from <http://www.federalreserve.gov/econresdata/researchdata.htm>.

⁶Greenwood and Vayanos (2010) deal with this issue by stripping coupons off and converting from face value to present value using the historical average short rate.

the Fed to be the better measures to use for our purposes, and simply use the number of remaining weeks to maturity on any given series as the value for n .

We separately constructed rough estimates of how much of the security of each maturity was held by the Federal Reserve, as detailed in Appendix E. The resulting data structures for outstanding Treasury debt and Fed holdings take the form of (240×1577) matrices, with rows corresponding to months (ranging from January 31, 1990 to December 31, 2009) and columns corresponding to maturity in weeks up to 30 years. Figure 3 displays the information from the December 31, 2006 rows of these two matrices. Figure 4 provides a sense of some of the time-series variation, plotting the average maturity of debt held by the public for each month.⁷ Average maturity dropped temporarily in the mid-1990s and began a more significant and sustained decrease after 2001. Average maturity dropped sharply between September 2007 and October 2008, but has since reverted back to September 2007 levels.

6 The term structure of interest rates prior to the financial crisis.

In our baseline specification, we took the $J = 3$ observed factors to be the deviations from the sample mean of the level, slope, and curvature of the term structure implied by the 6-month, 2-year, and 10-year Treasuries⁸, sampled weekly from January 1990 through the end of July, 2007. These yields and the 3 implied factors are plotted in Figure 5. The level factor trended

⁷The graph plots $\sum_{n=1}^N nz_{nt}$ for each t .

⁸That is, if maturities were measured in weeks, prior to demeaning we would have $f_{1t} = (1/3)(y_{26,t} + y_{104,t} + y_{520,t})$, $f_{2t} = y_{520,t} - y_{26,t}$, and $f_{3t} = y_{520,t} - 2y_{104,t} + y_{26,t}$.

down over this period, with pronounced dips after the recessions of 1990-91 and 2001. During these episodes, the term structure also sloped up more than usual and the curvature increased as the 2-year yield fell away from the 10-year. The parameters c , ρ and Σ reported in Table 1 were estimated by OLS regressions of each factor on a constant and lagged values of the other three factors. We chose $M = 4$ other yields⁹ (the 3-month, 1-year, 5-year, and 30-year) in the vector Y_{2t} in order to estimate the parameters c^Q, ρ^Q, a_1, b_1 and Σ_e from equation (16). We measured f_t in annual percentage points to keep reporting units natural and measured y_{nt} in weekly discount units so that the asset-pricing recursions all hold as written; for example, a 5.2% continuously compounded annual rate would correspond to $f_{1t} = 5.2$ and $y_{nt} = 0.001$.

The model described in Section 3 implies that an objective forecast (sometimes referred to as the P -measure expectation) of the 3 factors is given by

$$E_t^P(f_{t+1}) = c + \rho f_t.$$

⁹Note that this approach does not make full use of all the available information, in that we do not impose any connection between the model-implied value for

$$y_{520,t} - y_{26,t} = a_{520} - a_{26} + b'_{520}f_t - b'_{26}f_t$$

and the observed value of f_{2t} itself. However, the smooth structure of the ATSM causes these restrictions to be approximately satisfied even without imposing them, that is, the estimates reported below are characterized by

$$\begin{bmatrix} \hat{b}_{26} \\ \hat{b}_{104} \\ \hat{b}_{520} \end{bmatrix} \approx \begin{bmatrix} (1/3) & (1/3) & (1/3) \\ -1 & 0 & 1 \\ 1 & -2 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & -(1/2) & (1/6) \\ 1 & 0 & -(1/3) \\ 1 & (1/2) & (1/6) \end{bmatrix}.$$

It is possible instead to apply the minimum-chi-square algorithm to a system imposing restrictions such as the above equation directly. The effect of adding this restriction (along with the analogous expressions for level and curvature) is to fix the values of ρ^Q and b_1 up to the eigenvalues of ρ^Q , which eigenvalues are then estimated from (16). We applied this approach to several of the systems examined below and obtained almost identical results to those from the simpler approach that ignores these restrictions. To minimize the computational and expositional burden, we only report here the estimates from the unrestricted version of the model.

However, as a result of risk aversion, arbitrageurs value assets the way a risk-neutral investor would if that investor believed that the forecast was instead characterized by the Q -measure expectation

$$E_t^Q(f_{t+1}) = c^Q + \rho^Q f_t.$$

The risk premium is the difference between these two forecasts,

$$\begin{aligned} E_t^P(f_{t+1}) - E_t^Q(f_{t+1}) &= \Sigma\lambda + \Sigma\Lambda f_t \\ &= \Sigma\lambda_t \end{aligned}$$

which is plotted in Figure 6. At the beginning of the sample, investors behaved as if they expected next week's level of interest rates to be about 3 basis points higher (at an annual rate) than an objective forecast would imply, though this risk premium had mostly vanished by the end of the sample. Throughout the sample, arbitrageurs acted as if they expected the slope to be flatter than it usually turned out to be, and often expected the 10-year-2-year spread to move closer to the 2-year-6-month spread.

We next consider how the term-structure risk factors would be priced according to the Vayanos-Vila framework under the following special case. Suppose that (1) the preferred-habitat sector consisted solely of the U.S. Treasury and Federal Reserve, (2) the arbitrageurs comprise the entire private sector, and (3) U.S. Treasury debt is the sole asset held by arbitrageurs. These are obviously extreme assumptions, but they have the benefit of implying a clear answer to how changes in the maturity structure of outstanding Treasury debt would influence the price of risk in one highly stylized case. Under these conditions, the arbitrageurs'

portfolio weights z_{nt} could be measured directly from the ratio of debt held by the public of maturity n to the total outstanding publicly held debt at that date. From equations (7) and (8), we would then predict that $\Sigma\lambda_t = \gamma\Sigma\Sigma' \sum_{n=2}^N z_{nt}\bar{b}_{n-1}$. Figure 7 plots the three elements of

$$q_t = 100\Sigma\Sigma' \sum_{n=2}^N z_{nt}\bar{b}_{n-1} \quad (17)$$

where a value of $\gamma = 100$ was assumed in order to bring the series roughly on the same scale as $\Sigma\lambda_t$. This series for q_t was calculated with the values \bar{b}_n calculated from equation (11) for ρ^Q , and b_1 reported in Table 1. The values for the 3 elements of q_t are highly correlated, though as we shall see shortly, there is statistically useful information in the difference between them.

If the strong assumptions detailed above were literally true, then the vector q_t would be proportional to the corresponding series plotted in Figure 5, and indeed the level, slope, and curvature of the term structure could be described solely in terms of changes in the maturity composition of the public debt as summarized by these three factors. Obviously the assumptions do not hold, and the maturity composition of outstanding Treasury debt is just one of many factors potentially contributing to interest rate moves. However, it is interesting to look at what connections there may be in the data between q_t and pricing of interest-rate risk. Before doing so, we emphasize that although the above theory suggests that q_t might be related to the behavior of interest rates, in terms of how the series is constructed mechanically from the data, the time-series variation in q_t is driven solely by changes in the composition of Treasury debt z_{nt} and not at all by changes in interest rates. We accordingly propose the vector q_t as a possible 3-dimensional summary statistic of how the maturity composition of Treasury debt changes over time, where the simple theory sketched above suggests that this

might be a summary statistic of interest for purposes of analyzing changes over time in the term structure of interest rates.

We begin by examining the ability to predict excess holding yields for bonds of different maturities. Let p_{mt} denote the log price of a pure-discount m -month bond purchased on the last day of month t .¹⁰ The k -month holding yield for the bond (quoted at an annual rate) is $(12/k)(p_{m-k,t+k} - p_{mt})$. This compares with the holding yield for a k -month bond of $(12/k)(p_{0,t+k} - p_{kt}) = (12/k)(-p_{kt})$. Let h_{mkt} denote the excess holding yield for an m -month relative to a k -month bond:

$$h_{mkt} = (12/k)(p_{m-k,t+k} - p_{mt} + p_{kt}).$$

We explored regressions to predict these holding yields on the basis of information available at date t :

$$h_{mkt} = c_{mk} + \beta'_{mk}f_t + \gamma'_{mk}x_t + u_{mkt}. \quad (18)$$

If investors were risk-neutral, all the coefficients in (18) would be zero. Our finding of nonzero elements for λ and Λ in Table 1 (and a huge literature before us) suggests nonzero values for c_{mk} and β_{mk} , though if the market pricing of risk were fully captured by the 3-factor affine-term-structure model, no other variables x_t should enter statistically significantly.¹¹

Table 2 reports the results from OLS estimation of (18), giving the R^2 of the regression and Newey-West tests of the hypothesis that γ_{mk} or subsets of γ_{mk} are zero for various specifications

¹⁰We inferred these prices from the daily term-structure summaries of Gürkaynak et al. (2007).

¹¹Although u_{mkt} is uncorrelated with the regressors in (18), it is not independent of the regressors, and thus OLS is subject to the small-sample problems highlighted by Stambaugh (1999). Moreover, given that risk-neutrality does not hold, both the left-hand and right-hand variables in (18) are highly serially correlated, raising potential spurious regression concerns if these are near-unit-root processes.

of x_t .¹² The first row reproduces the well-known result that the traditional level, slope, and curvature factors f_t can predict a significant amount of the excess holding yield on assets of assorted maturities, with for example an R^2 of 0.33 in the case of predicting the excess returns from holding a 2-year bond for one year. The second row adds the average maturity of outstanding debt,

$$z_t^A = \sum_{n=1}^N n z_{nt},$$

which was one of the summary statistics examined by Greenwood and Vayanos (2010),¹³ but which we find in our sample usually does not have statistically significant additional predictive power beyond that contained in f_t . On the other hand, the other measure they propose, the fraction of outstanding debt of more than 10-year maturity,

$$z_t^L = \sum_{n=521}^N z_{nt},$$

does statistically significantly predict excess returns.

One could consider various other linear combinations of $\{z_{nt}\}_{n=1}^N$ as possible predictors, such as the first three principal components. We find in the fourth row of Table 2 that these are helpful for forecasting the holding returns on short-maturity assets, but are generally inferior to z_t^A or z_t^L .

¹²Note that even though the excess holding yield would follow an $MA(k-1)$ process under the null hypothesis of risk neutrality, one would still need to let the Newey-West lag parameter go to infinity as the sample size grows in order to get a consistent estimate. The Newey-West approach is helpful under the alternative hypothesis of a possibly more complex serial correlation, and generates a positive-definite variance-covariance matrix by construction. We also performed these calculations using Hansen-Hodrick (1980) standard errors based on $k-1$ lags. These produced the same results except for one case in which the Hansen-Hodrick standard error was negative.

¹³Greenwood and Vayanos (2010) use duration rather than maturity.

The theory sketched above suggests three particular linear combinations of $\{z_{nt}\}_{n=1}^N$ that should matter for term premia, namely the three elements of the vector q_t in (17). The sixth row of Table 2 shows that these turn out to be incredibly useful for predicting holding returns, with an R^2 as high as 0.71 in the case of predicting the 2-over-1 excess return. The contribution of q_t is statistically significant for every maturity, even if the regression already includes both f_t and the first three principal components of $\{z_{nt}\}_{n=1}^N$.

Cochrane and Piazzesi (2005) propose a particular yield pricing factor that they have found very helpful for forecasting excess holding returns. In our application, we confirm that this factor¹⁴ provides a statistically significant improvement over using just f_t alone (row 5 of Table 2). Nevertheless, our Treasury factors q_t still provide a very dramatic improvement in forecasting ability beyond that contained in f_t and the Cochrane-Piazzesi factor v_t (row 8).

We next examine the ability of the Treasury factors q_t to help predict the yields themselves, examining OLS regressions of the form

$$f_{t+1} = c + \rho f_t + \phi q_t + \varepsilon_{t+1} \tag{19}$$

for ϕ a (3×3) matrix. The first column of Table 3 reports that the vector q_t makes a useful contribution to predicting each of the term-structure factors, with the hypothesis that the i th row of ϕ is zero being rejected for each i .

It is then tempting to use (19) to draw tentative conclusions about what the effects on yields of different maturities might be of a change in the composition of publicly held debt. Such

¹⁴In our application, we constructed v_t from the fitted value of a regression of $(1/4)(h_{24,12,t} + h_{36,12,t} + h_{48,12,t} + h_{60,12,t})$ on a constant and the 1- through 5-year forward rates at date t .

calculations are subject to a well-understood endogeneity problem: historical variations in z_{nt} may have represented a response by the Treasury or the Fed to overall economic conditions or to term-structure developments in particular. Although this is also a potential concern for (19), our formulation has three advantages over traditional regressions which simply examine the contemporaneous correlations. First, any contemporaneous response of q_t to f_t could not account for a nonzero value of ϕ in (19). We are explicitly asking about the ability of q_t to forecast future f_{t+1} over and above any information contained in f_t itself.¹⁵ Second, because the statistics we report represent the answer to well-posed forecasting questions, the results have independent interest as objective summaries of those forecasting relations, regardless of what the underlying dynamic structural relations may be. Third, because we include lags of the dependent variable in the regression, we avoid the potential spurious regression problem that could plague other popular approaches such as trying to use OLS to estimate a relation of the form $f_t = \alpha + \beta z_t^A$.

For purposes of focusing on a particular forecasting question that might be of interest to policy makers, we consider the following exercise. Suppose that at the end of month t , the Federal Reserve were to sell all its Treasury securities with maturity less than 1 year, and use the proceeds to buy up all of the outstanding nominal Treasury debt of maturity greater than n_{1t} , where n_{1t} would be determined by the size of the Fed's short-term holdings and outstanding long-term Treasury debt at time t . For example, if implemented in December of 2006, this would result in the Fed selling about \$400 B in short-term securities and buying about \$400 B in long-term securities, effectively retiring all the federal debt of ten-year and

¹⁵On the other hand, if q_t only matters for f_{t+1} through its effect on f_t , we might understate the contribution of q_t using our approach.

longer maturity. We then calculated what q_t^A would be under this counterfactual scenario, and calculated the average historical value of $q_t^A - q_t$, which turns out to be

$$\Delta = \begin{bmatrix} 0.026101 \\ 0.022712 \\ -0.00780 \end{bmatrix}. \quad (20)$$

We then asked, by how much would one expect f_{t+1} to change according to (19) if q_t were to change by Δ ? As should be clear from the description of the exercise, we are talking about a quite dramatically counterfactual event. If one considers the analogous forecasting equations of the form $q_{t+1} = c_q + \rho_q f_t + \phi_q q_t + \varepsilon_{q,t+1}$, a change of q_t of the size of Δ would represent a 36σ event, obviously something so far removed from anything that was attempted during the historical sample as to raise doubts about interpreting the parameter estimates as telling policy makers what would happen if they literally implemented a change of this size.

The second column of Table 3 reports how a forecast of the traditional term-structure factors would be affected by this change. We find that changing q_t by this amount could flatten the slope of the yield curve by 25 basis points, with no effect on the level of interest rates themselves. If it reduces the slope but has no effect on the level, that means it would reduce long-term yields and raise short-term yields. Indeed, our 3-factor ATSM has a prediction¹⁶ as to how much any given interest rate would change if the factors were to change by the amount specified in Table 3, which predicted responses we plot in Figure 8. Yields on maturities longer than 2-1/2 years would fall, with those at the long end decreasing by up to 17 basis

¹⁶The predicted change in y_{nt} is given by $b'_n \hat{\phi} \Delta$ for $b_n = -\bar{b}_n/n$, \bar{b}_n calculated from equation (11) using the values of ρ^Q and b_1 reported in Table 1, $\hat{\phi}$ the OLS estimates from equation (19), and Δ given by (20).

points. Yields on the shortest maturities would increase by almost as much.¹⁷

A separate question from the feasibility for the Federal Reserve to achieve such effects is the desirability of its attempting to do so. Although we have described this as a Fed operation, it is probably more natural to think of it as a Treasury operation, implemented by the Treasury doing more of its borrowing at the shorter end of the yield curve. According to the simple framework that motivated our definition of q_t , the average slope of the yield curve arises from the preference of the U.S. Treasury for doing much of its borrowing with longer-term debt. For reasons presumably having to do with management of fiscal risks, the Treasury is willing to pay a premium to arbitrageurs for the ability to lock in a long-term borrowing cost. If the Treasury has good reasons to avoid this kind of interest-rate risk, it is not clear why the Federal Reserve should want to absorb it.

Our conclusion is that, although it appears to be possible for the Fed to influence the slope of the yield curve in normal times through the maturity of the System Open Market Account holdings, very large operations are necessary to have an appreciable immediate impact. If there is no concern about a zero-lower-bound constraint, this potential tool should clearly be secondary to the traditional focus of open-market operations on the short end of the yield curve.

¹⁷Our estimates would also allow us in principle to answer dynamic questions, though we are much less comfortable with using the framework for this purpose. One problem is that the standard errors for dynamic responses turn out to be quite large. Another challenge is trying to infer the permanent consequences of changes whose time-series variation has been transitory.

7 The term structure of interest rates at the zero lower bound.

The above analysis ended prior to the first stages of the financial crisis in August 2007. As discussed in Section 2, we divide subsequent developments into two phases. The first phase was characterized by high default premiums, failures of some leading financial institutions, and serious disruption of traditional lending patterns. Gürkaynak and Wright (2010) documented that under the financial strains, significant arbitrage opportunities between yields on different Treasury securities often persisted between October 2008 and February 2009. We will not attempt to address the many important issues having to do with monetary policy under those circumstances, but instead begin our analysis here with the second phase which began in March of 2009, and during which policy makers have confronted the longer-term issue of how to provide stimulus to aggregate demand when the short-term interest rate had essentially reached zero.

Figure 9 plots assorted yields over this period. The 3-month yield has remained stuck near zero over this period, and the 1-year, although higher, has also displayed little variability. Nonetheless, there has continued to be considerable fluctuation in longer-term yields. What is the nature of the developments driving long-term yields in this environment?

The natural answer is that investors do not believe the U.S. will remain at the zero lower bound forever. When the U.S. escapes from the ZLB, interest rates at all maturities will again respond as they always have to changes in economic fundamentals. Any news today that leads to revisions in the expectations of those future fundamentals shows up as changes

in those longer-term yields.

We propose that one way to interpret current long-term yields is to postulate the existence of latent factors, denoted f_t , which would determine what interest rates would currently be doing if the ZLB were not binding, along with probabilities that arbitrageurs assign to escaping from the ZLB at various future dates. For the first task, what should we assume about the dynamic behavior of these latent factors? The most parsimonious hypothesis would obviously be that, when the economy escapes from the ZLB, the factor dynamics would revert to their historic behavior as represented by equations like (5) or (19). The difference is that, when we originally introduced these equations, we were treating the factors f_t as directly observed from the level, slope, and curvature of the term structure, whereas we are proposing now to interpret them as latent factors characterizing what the level, slope, and curvature would be if we were not stuck at the ZLB. For the second task, we again adopt the simplest possible hypothesis, which is that arbitrageurs assign a constant Q -measure probability π^Q that the economy will remain at the ZLB next week.

To develop this idea in more detail, we postulate that, once the economy escapes from the ZLB, the short rate will return to being determined by the factors according to the structure

$$\tilde{y}_{1t} = a_1 + b'_1 f_t$$

$$\tilde{p}_{nt} = \bar{a}_n + \bar{b}'_n f_t$$

where the sequences $\{\bar{a}_n, \bar{b}_n\}_{n=1}^N$ can be calculated as before using the recursions (11) and (13).

However, as long as the economy remains at the ZLB, we instead have

$$y_{1t}^* = a_1^*$$

$$p_{nt}^* = \bar{a}_n^* + \bar{b}_n^{*'} f_t.$$

If the zero lower bound were interpreted literally, then a_1^* would be zero. We represent it instead with some number slightly above zero to match the U.S. experience in which an interest rate paid on reserves has prevented the rate from falling all the way to zero.

Let $q_{n,t+1}$ denote the holding return on an n -period bond purchased at t and sold at $t + 1$. Note that if t is characterized by the ZLB, the Q -measure expectation of this return is given by

$$\begin{aligned} E_t^Q(q_{n,t+1}) &= E_t^Q \left[\frac{(P_{n-1,t+1} - P_{nt})}{P_{nt}} \right] \\ &= \pi^Q E_t^Q \left[\frac{(P_{n-1,t+1}^* - P_{nt}^*)}{P_{nt}^*} \right] + (1 - \pi^Q) E_t^Q \left[\frac{(\tilde{P}_{n-1,t+1} - P_{nt}^*)}{P_{nt}^*} \right] \\ &\approx \pi^Q \left[\bar{a}_{n-1}^* + \bar{b}_{n-1}^{*'} (c^Q + \rho^Q f_t) \right] + (1 - \pi^Q) \left[\bar{a}_{n-1} + \bar{b}_{n-1}' (c^Q + \rho^Q f_t) \right] \\ &\quad + (1/2) \pi^Q \bar{b}_{n-1}^{*'} \Sigma \Sigma' \bar{b}_{n-1}^* + (1/2) (1 - \pi^Q) \bar{b}_{n-1}' \Sigma \Sigma' \bar{b}_{n-1} - \bar{a}_n^* - \bar{b}_n^{*'} f_t. \end{aligned}$$

No-arbitrage requires the Q -measure expected one-period holding yield for an n -period bond to equal y_{1t} ,

$$a_1^* = E_t^Q(q_{n,t+1}).$$

This requires

$$\bar{b}_n^{*'} = \pi^Q \bar{b}_{n-1}^{*'} \rho^Q + (1 - \pi^Q) \bar{b}_{n-1}^{*'} \rho^Q \quad (21)$$

$$\begin{aligned} \bar{a}_n^* &= \pi^Q \bar{a}_{n-1}^* + (1 - \pi^Q) \bar{a}_{n-1} + \pi^Q \bar{b}_{n-1}^{*'} c^Q + (1 - \pi^Q) \bar{b}_{n-1}^{*'} c^Q \\ &+ (1/2) \pi^Q \bar{b}_{n-1}^{*'} \Sigma \Sigma' \bar{b}_{n-1}^* + (1/2) (1 - \pi^Q) \bar{b}_{n-1}^{*'} \Sigma \Sigma' \bar{b}_{n-1} - a_1^*. \end{aligned} \quad (22)$$

Given $c^Q, \rho^Q, a_1, b_1, \Sigma$ we can calculate $\{\bar{a}_n, \bar{b}_n\}_{n=1}^N$ from (11) and (13). Given these and $\bar{b}_1^* = 0$, we can calculate $\{\bar{a}_n^*, \bar{b}_n^*\}_{n=1}^N$ as functions of π^Q and a_1^* . Predicted bond yields under the ZLB are then given by

$$y_{nt}^* = a_n^* + b_n^{*'} f_t \quad (23)$$

where $a_n^* = -\bar{a}_n^*/n$ and $b_n^* = -\bar{b}_n^*/n$.

As a first pass, we propose to use the same values for $c^Q, \rho^Q, a_1, b_1, \Sigma$ as estimated from the earlier historical sample. Note that even though these parameters are the same as before, the implied mapping from factors f_t into observed yields has changed. Let $Y_{1t} = (y_{26,t}, y_{104,t}, y_{520,t})'$ denote the 6-month, 2-year, and 10-year yields observed at time t . In our historical sample, these were related to the factors f_t according to

$$Y_{1t} = A_1 + B_1 f_t \quad (24)$$

$$A_1 = \begin{bmatrix} a_{26} \\ a_{104} \\ a_{520} \end{bmatrix} \quad B_1 = \begin{bmatrix} b'_{26} \\ b'_{104} \\ b'_{520} \end{bmatrix}.$$

Because we treated the factors in normal times as directly observed from the 6-month, 2-year, and 10-year level, slope, and curvature, and because of the smoothness of the ATSM term structure, our estimates were characterized by

$$B_1 \approx \begin{bmatrix} (1/3) & (1/3) & (1/3) \\ -1 & 0 & 1 \\ 1 & -2 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & -(1/2) & (1/6) \\ 1 & 0 & -(1/3) \\ 1 & (1/2) & (1/6) \end{bmatrix}$$

where the approximation would have been exact if we had imposed the restriction that Y_{1t} is observed without error.

By contrast, under the ZLB, the relation is

$$Y_{1t} = A_1^* + B_1^* f_t \tag{25}$$

$$A_1^* = \begin{bmatrix} a_{26}^* \\ a_{104}^* \\ a_{520}^* \end{bmatrix} \quad B_1^* = \begin{bmatrix} b_{26}^{*'} \\ b_{104}^{*'} \\ b_{520}^{*'} \end{bmatrix} .$$

Let Y_{2t} denote the four other yields used in the estimation, namely the 3-month, 1-year, 5-year, and 30-year yields. The model implies that

$$Y_{2t} = A_2^* + B_2^* f_t + \varepsilon_t^e \tag{26}$$

$$A_2^* = \begin{bmatrix} a_{13}^* \\ a_{52}^* \\ a_{260}^* \\ a_{1560}^* \end{bmatrix} \quad B_2^* = \begin{bmatrix} b_{13}^{*'} \\ b_{52}^{*'} \\ b_{260}^{*'} \\ b_{1560}^{*'} \end{bmatrix}$$

where $\varepsilon_t^e \sim N(0, \Omega_e)$ denotes measurement error. Substituting (25) into (26),

$$Y_{2t} = A_2^\dagger + B_2^\dagger Y_{1t} + \varepsilon_t^e \quad (27)$$

$$A_2^\dagger = A_2^* - B_2^*(B_1^*)^{-1}A_1^* \quad B_2^\dagger = B_2^*(B_1^*)^{-1}. \quad (28)$$

We applied the minimum-chi-square estimation approach developed by Hamilton and Wu (2010) to weekly interest rate data from March 6, 2009 to August 4, 2010 to infer the values of π^Q and a_1^* from the OLS estimates of \hat{A}_2^\dagger and \hat{B}_2^\dagger , taking $c^Q, \rho^Q, a_1, b_1, \Sigma$ as given by the pre-2007 parameter estimates, as detailed in Appendix D.

This procedure resulted in estimates $5200\hat{a}_1^* = 0.037$ and $\hat{\pi}^Q = 0.9834$, implying that the ZLB is characterized by a one-week interest rate of 4 basis points (at an annual rate) and that arbitrageurs expect the ZLB to persist for $1/(1 - \pi^Q) = 60$ weeks. We used these two parameters along with the pre-crisis values for $c^Q, \rho^Q, a_1, b_1, \Sigma$ reported in Table 1 to calculate b_n^* and a_n^* from (21) and (22), and used these to infer a value for f_t on the basis of the observed 6-month, 2-year, and 10-year yield using (25). With this f_t we then have from (26) predicted values for each week's 3-month, 1-year, 5-year, and 30-year yields, which predictions are plotted as dashed lines of Figure 9. The R^2 for each relation is reported in the first column of Table 4. We might compare these with the best possible fit as represented by an

unrestricted OLS regression of each yield on a constant and the 6-month, 2-year, and 10-year yields, whose R^2 is reported in the second column of Table 4. Particularly for the longer-term yields, the predictions from our simple restricted parameterization are not far from what is actually observed during the ZLB period.

A tougher test of the framework is whether it can successfully predict yields in advance. Here we used the f_t constructed as above, formed the one-week-ahead forecast $E_t^P(f_{t+1}) = c + \rho f_t$ again on the basis of the pre-crisis parameters reported in Table 1, and calculated the implied yields $y_{n,t+1}$ using (23). Again, particularly for the longer maturities, these forecasts are reasonably close to the best possible in-sample fit as represented by an unrestricted OLS regression of $y_{n,t+1}$ on a constant and Y_{1t} (see columns 3 and 4 of Table 4).

Although the post-sample fit is good, the model could nevertheless still be improved. Hamilton and Wu (2010) propose a test of the overidentifying restrictions, which is basically a test of the statistical significance of the difference in R^2 between the first and second columns of Table 4. This leads to quite strong rejection, with a $\chi^2(14)$ test statistic of 344.5.

We made one further simple adjustment to improve the fit further. We postulated that when the economy escapes from the ZLB, arbitrageurs anticipate a different average level of interest rates (as governed by the parameter a_1) compared to that observed in the pre-crisis episode. The estimated value of $5200a_1$ is 2.19, meaning arbitrageurs expect the post-ZLB average short rate to be below the 4.12 level observed over 1990-2007. The new estimate of $5200a_1^*$ is 0.068 and of π^Q is 0.9907, implying an expected ZLB duration of 108 weeks. These changes improve the fit relative to that of the model summarized in Figure 9 and Table 4, though the specification would still be rejected ($\chi^2(13) = 176.0$).

Although one could relax other restrictions of the model until a perfect fit is achieved, we regard this as an attractive parsimonious framework that successfully captures the broad features of how interest rates have been observed to behave under the ZLB regime to date. Another benefit is that this framework gives us an immediate basis for drawing conclusions about how the effects of monetary policy differ under the ZLB from normal times.

Figure 10 plots the factor loadings, which summarize how the yield of any maturity n is predicted to respond to changes in any of the three factors. The main difference is that, under the ZLB, short-term yields are essentially unresponsive to any macroeconomic developments, with all three elements of b_n^* near zero for small n . This is because arbitrageurs see very little probability of escaping from the ZLB over most of the term of the security. As n increases, the response of the yield to macroeconomic factors becomes larger and approaches the response observed in normal conditions, because there is an increasing probability that the economy will be away from the ZLB for most of the security's duration.

This framework allows us to revisit the consequences of a shift in the maturity of the Fed's Treasury holdings. Given our assumption that the latent factors f_t are responding in the same way as they would when away from the ZLB, we can still use the prediction that a change in the maturity composition of publicly held debt that changes the Treasury risk factor vector by Δ would change f_{t+1} by $\phi\Delta$. But whereas in normal times we premultiplied this vector by b'_n to see what the change Δ implied for a yield of maturity n , at the ZLB we would instead premultiply $\phi\Delta$ by b_n^{*} . These predicted impacts are compared in Figure 11. The policy continues to depress long-term yields by the same amount as in normal times, but, because of the ZLB, it has very little effect on short-term yields. Cumulative effects on short-term

yields are also negligible, while the ability to bring long yields down is the same as without the ZLB, as seen in Figure 11.

We have analyzed here the effects of a swap by the Federal Reserve of short-term assets for longer-term assets. An alternative strategy, which might be characterized as quantitative easing, is for the Fed to buy longer-term assets outright with newly created reserves. At the ZLB, interest-bearing reserves are essentially indistinguishable from zero-risk 1-week bonds. The effect of quantitative easing is to reduce the available supply of longer-term securities without changing the private-sector's exposure to the risk associated with holding short-term securities. But at the ZLB, changes in the supply of short-term securities have essentially no effects. Thus, the economic consequences of quantitative easing would be identical to those of the maturity swap just described if the economy were at the ZLB.

8 Discussion.

8.1 Comparison with other estimates.

Here we compare our estimates with those obtained by other researchers. For this purpose, we standardize on the basis of the two scenarios analyzed above. The first scenario is a simultaneous sale by the Fed of \$400 B in securities at the short end and purchase of \$400 B in securities at the long end, implemented in December of 2006. The second scenario is an outright purchase of \$400 B in long-term securities, implemented at the zero lower bound.

Gagnon et al. (2010) used as an explanatory variable the face value of privately-held debt of more than one-year maturity as a percent of GDP, and as dependent variable the 10-year

yield or 10-year term premium. They estimated the effect of debt supply on yields using regressions estimated 1986:M12 to 2008:M6 that included several other explanatory variables, and obtained a coefficient relating the 10-year yield to bond supply of 0.069. Since \$400 B would represent about 2.9% of U.S. GDP in 2006:Q4, their estimates imply a predicted decline in the 10-year yield under scenario 1 of $(2.9)(0.069) = 20$ basis points. This is close to our estimate of a decline of 14 basis points, as reported in the first row of Table 5.¹⁸

In the analysis of Greenwood and Vayanos (2010), the right-hand variable was the fraction of privately-held debt with duration greater than 10 years, and the left-hand variable was assorted yield spreads. They found that a one-percentage-point increase in the share resulted in a 4-basis-point increase in the 5-year-1-year spread over the period 1952-2006. In the sample we studied (1990-2007), a maturity swap of the size contemplated in scenario 1 would have lowered the share of debt with maturity greater than 10 years by 9.8 percentage points. This gives an effect implied by the Greenwood-Vayanos estimates of $(9.8)(4) = 39$ basis points. For comparison, our estimate of the size of the effect is 17 basis points for scenario 1, but only 9 basis points for scenario 2. The reason for the difference between the two scenarios is that, in our framework, part of the drop in the spread if the policy had been implemented over the period studied by Greenwood and Vayanos (2010) would have come from an increase in short-term yields, something that would not happen if the same purchase were implemented at the zero lower bound.

¹⁸Gagnon et al. (2010)'s regressions in which the term premium rather than the yield is the left-hand variable would imply estimates as low as 12 basis points. However, these are harder to compare directly with those for our scenario. In our conception of the question being asked, we assume that the supply of securities with maturity less than one year increases by \$400 B, driving up the yield on those securities and making the decrease in the term premium larger than the decrease in the yield. This effect is not captured by the Gagnon et al. (2010) regressions.

Another recent analysis comes from D’Amico and King (2010), who look at the change in yields of different maturities during the Fed’s purchase of \$300 billion in long-term securities between March and October of 2009. They conclude that these purchases lowered the yield on 10-year Treasuries by about 50 basis points, which would translate into an effect of $(4/3)(50) = 67$ basis points for the \$400 B purchase analyzed in Table 5, a somewhat larger effect than implied by our estimates. However, the 10-year yield was where these purchases were concentrated and where D’Amico and King found the biggest effects, and large standard errors are associated with any of these estimates.

Deutsche Bank (2010) attempted to synthesize the estimates of Gagnon et al. (2010), Macroeconomic Advisers, and their own research staff, and estimated that \$1 trillion in long-term purchases in the current setting might produce a 50-basis-point decline in long-term yields, which we’ve translated as a 20-basis-point decline for the \$400 billion purchase reported in Table 5.

In 1961, the U.S. attempted to use Treasury and Fed operations to lower the fraction of publicly-held long-term debt in what was referred to as “Operation Twist.” Swanson (2011) used a daily event study of announcements pertaining to the Operation Twist and found effects on bond yields that, when scaled by the change in size of outstanding Treasury debt, are broadly consistent with those summarized in our Table 5.

Although our estimates of the effects are the smallest in this group, they are generally in the same ballpark, which is somewhat surprising given the very different ways in which these estimates are derived. There is overall agreement that sufficiently large asset purchases could achieve a modest reduction in long-term yields. There is nevertheless considerable uncertainty,

both in terms of the econometric standard errors and possible specification errors, in any of the estimates reported.

8.2 Effects on non-Treasury securities.

Here we sketch a generalization of the theoretical framework in Section 3 to allow arbitrageurs also to hold other securities with a nonzero probability of default.

Let P_{1t}^\ddagger denote the price paid at t for a one-period bond whose value next period will be

$$P_{0,t+1}^\ddagger = \begin{cases} 1 & \text{with probability } \exp(-\psi_t) \\ 0 & \text{with probability } 1 - \exp(-\psi_t) \end{cases}.$$

If the arbitrageurs hold a fraction z_{1t}^\ddagger in the risky asset and if the probability of default ψ_t is independent of risk factors f_t , then using a similar approach to that in Appendix A, the contribution of the risky asset to the variance can be approximated¹⁹ by $z_{1t}^{\ddagger 2}\psi_t$ and the no-arbitrage condition (2) becomes

$$y_{1t}^\ddagger = y_{1t} + \psi_t(1 + \gamma z_{1t}^\ddagger). \quad (29)$$

¹⁹If we conjecture that $p_{1t}^\ddagger = h(\bar{a}_1^\ddagger + \bar{b}_1^{\ddagger'} f_t + \bar{c}_\psi^\ddagger \psi_t + \bar{c}_\zeta^\ddagger \psi_t \zeta_{1t}^\ddagger)$ for ζ_{1t}^\ddagger independent factors affecting the supply of risky assets,

$$\begin{aligned} E_t \left[z_{1t}^\ddagger \left(\frac{P_{0,t+1}^\ddagger}{P_{0t}^\ddagger} - 1 \right) \right]^2 &= z_{1t}^{\ddagger 2} \{ \exp(-\psi_t h) \exp[-2h(\bar{a}_1^\ddagger + \bar{b}_1^{\ddagger'} f_t + \bar{c}_\psi^\ddagger \psi_t + \bar{c}_\zeta^\ddagger \psi_t \zeta_{1t}^\ddagger)] \\ &\quad - 2 \exp(-\psi_t h) \exp[-h(\bar{a}_1^\ddagger + \bar{b}_1^{\ddagger'} f_t + \bar{c}_\psi^\ddagger \psi_t + \bar{c}_\zeta^\ddagger \psi_t \zeta_{1t}^\ddagger)] + 1 \} \\ &= z_{1t}^{\ddagger 2} \psi_t h + o(h). \end{aligned}$$

In the absence of risk aversion ($\gamma = 0$), in equilibrium the risky security will offer the same expected return as the risk-free security, which requires a premium of ψ_t to compensate for the probability of default. With risk aversion ($\gamma > 0$) and a positive exposure of arbitrageurs to this risk ($z_{1t}^\ddagger > 0$), the risky asset will offer a higher expected return to compensate for the risk.

If the factors that govern ψ_t and determine equilibrium z_{1t}^\ddagger are independent of the factors f_t that determine the risk-free yield, the one-period risky rate would have identical loadings as y_{1t} on fluctuations in the level, slope, and curvature factors, as well as additional loadings on separate default-risk factors. A parallel result can be derived for risky assets of longer maturity, with p_{nt}^\ddagger loading on f_t with the same coefficients \bar{b}_n as for risk-free bonds, along with separate loadings on the default-risk factors.

Although the independence of Treasury and default risk factors is a highly stylized assumption, there is no question that risky yields of different maturities respond in a similar way to the factors driving Treasury yields. Figure 12 displays the comovement between the 10-year Treasury rate and that on 30-year mortgages and Aaa-rated and Baa-rated corporate debt²⁰.

Rather than impose a particular loading of non-Treasury yields on the level, slope, and curvature factors, we can estimate the empirical loading directly by OLS estimation of

$$y_{jt}^\ddagger = a_j^\ddagger + b_j^{\ddagger'} f_t + u_t$$

²⁰Aaa and Baa yield represent values as of the last day of the month, while the 30-year mortgage rate is for the last week of the month, from the FRED database of the Federal Reserve Bank of St. Louis.

over $t = 1990:M1$ to $2007:M7$ for assorted securities j . Note that if there is a correlation between the default risk factors and f_t , this will be incorporated in the estimated values of b_j^\ddagger . Table 6 reports the empirical factor loadings for these three risky yields, which, not surprisingly given Figure 12, turn out to be similar to those for 10-year Treasury bonds.

In the next-to-last column we use these estimated values of b_j^\ddagger to calculate the predicted effect in normal times of a shift in the maturity composition of Fed holdings.²¹ Based on the historical correlations between bond yields, in the pre-crisis period, if the Fed were to sell \$400 billion of short-term Treasuries and buy \$400 billion in long-term Treasuries, the 10-year T-bond and the Aaa and Baa corporate yields would each be expected to decline by 14 basis points, and the 30-year fixed mortgage rate by 11 basis points.

We can also get a quick impression of what might be expected at the zero lower bound as follows. The predicted change in the 6-month, 2-year, and 10-year yields of this \$400 billion maturity swap when at the ZLB are given by the corresponding elements of the vector $B_1^* \phi \Delta$. If y_{jt}^\ddagger tracks these as estimated historically (namely, by $b_j^{\ddagger'} B_1^{-1}$), then we get a predicted effect on y_{jt}^\ddagger at the ZLB of $b_j^{\ddagger'} B_1^{-1} B_1^* \phi \Delta$. These estimates are reported in the last column of Table 6. Interestingly, buying long-term Treasuries might if anything have an even bigger effect on risky yields when at the ZLB than it does in normal circumstances. Again, at the ZLB, in our framework the effects are the same whether the Fed finances the purchases with sales of short-term T-bills or with newly created reserves.

If the Fed were instead to purchase risky securities directly, the resulting reduction in arbitrageurs' holdings of these securities z_{nt}^\ddagger would both reduce the default risk premium

²¹These were calculated as $b_j^{\ddagger'} \phi \Delta$ for ϕ the matrix of OLS coefficients in (19) and Δ given by (20).

(through equation (29)) as well as affect the pricing of Treasury level, slope or curvature risk (because by holding these risky securities an investor is also exposed to the conventional term structure factors). For example, the Fed's MBS purchases could both flatten the slope of the Treasury yield curve and narrow the spread between MBS and Treasury yields.

We should also comment on how arbitrageurs' holding of risky securities would influence our empirical estimates of the matrix ϕ itself. If Treasuries represent only a subset of arbitrageurs' holdings, then Treasury holdings as a fraction of their total wealth z_{nt} would be a smaller number than we have assumed. If, for example, each z_{nt} were divided by 2, our vector q_t and therefore the magnitude Δ would be divided by two, while the OLS estimates $\hat{\phi}$ would be multiplied by two. Notice that a change in scale of this type would leave the estimated product $\phi\Delta$ unchanged and have no effect on any of the estimates reported. This invariance results from the fact that ultimately our estimates are simply an empirical summary of the historical relations between observed yields and maturity shares z_{nt} defined as a percentage of total publicly held federal debt, and it is the historical covariation of yields with outstanding Treasury debt that determined these estimates. If we were trying to make an inference about structural coefficients such as the risk aversion parameter γ , getting the scale right would be important. But for the purposes for which the estimates are used here, the scale of z_{nt} does not matter for any of the reported results.

8.3 Application: Evaluation of QE2.

On November 3, 2010, the Federal Reserve announced its intention to implement additional measures to stimulate the economy, which was described in the financial press as a second

round of quantitative easing (QE2). The plan was to purchase an additional \$600 billion of longer-term Treasury securities by the end of the second quarter of 2011, a pace of about \$75 billion per month. This differed in several details from the scenarios analyzed above.

The first difference is that, as implemented, the purchases were concentrated not on the longest-maturity securities, but instead focused primarily on securities between 2-1/2 and 10 years. Over the period 1990-2006, if the Fed had sold all its holdings of less than 1 year and used the proceeds to purchase outstanding Treasury debt evenly over the 2-1/2 to 10 year range, the resulting average change in q_t would be not the value reported in expression (20) but instead

$$\Delta_2 = (0.006898, 0.004479, -0.004406)'. \quad (30)$$

Whereas the estimated effects of Δ on the term structure are statistically significantly distinguishable from zero in our framework, those resulting from Δ_2 are not. Figure 13 compares the estimated effects on yields of Δ and Δ_2 if implemented at the zero lower bound. The dashed curve summarizes the predicted effects if the Fed were to sell all its holdings of less than 1-year maturity, and use the proceeds to retire debt of the longest outstanding maturities. Note this is identical to the dashed curve in Figure 11. The solid curve summarizes the predicted effects if the Fed were to sell all its holdings of less than 1-year maturity, and spread the proceeds evenly to purchase outstanding Treasury debt in the 2-1/2 to 10 year range. The latter has a significantly smaller effect on long-term rates. Again we interpret an outright purchase of a comparable quantity of securities as having similar effects to a debt swap when the economy is at the zero lower bound.

A second important difference between QE2 as it's been implemented by the Fed and the

scenarios analyzed here is in the timing, with the purchases associated with QE2 spread over a period of 8 months. Between November 2009 and November 2010, non-TIPS Treasury debt increased by \$127.3 billion per month, of which \$71.4 billion was in the 2-1/2 to 10 year maturity range. Hence, the proposed QE2 would barely absorb the newly issued medium-term debt, and debt of greater than 10 years would continue to increase rather than decline. The top panel of Figure 14 shows that the average maturity of publicly-held Treasury debt has been higher in each of the first three months of QE2 than it had been in any month over the preceding 2 years. The bottom panel shows that the fraction of publicly-held debt of more than 10 years maturity continued to increase even as the Fed was implementing its QE2 bond purchases. Our conclusion is that QE2 as implemented had little potential to lower long-term interest rates via the mechanism explored in this paper.

9 Conclusion.

We have found statistically significant forecasting relations over 1990-2007 between the maturity structure of Treasury debt held by the public and the behavior of U.S. interest rates. These relations suggest that in normal times, the Federal Reserve has some potential to flatten the yield curve, though not to reduce the overall level of interest rates, by selling short-term securities and buying long-term securities. Our estimates of the effect on impact suggest that quite massive operations would be necessary to have a measurable effect on interest rates.

We proposed that altering the maturity structure of publicly held Treasury debt would be equally effective at lowering long-term yields when the economy is at the zero lower bound.

But because there are negligible consequences for short-term yields in such a setting, the policy of reducing public holdings of long-term bonds has the potential to bring the overall level of interest rates down for an economy at the ZLB, whereas it could not do so in a normal environment. Quantitative easing, defined as buying the long-term bonds with newly created reserves, has the identical potential in this model.

One might suppose that the potential small magnitude of the effect is not a concern as far as the latter policy is concerned— if hundreds of billions are not enough to make much difference, then perhaps purchases in the trillions, such as the Fed has embarked upon with its holdings of mortgage-backed securities, might do the trick. However, we would emphasize that, in the model of the ZLB proposed here, the entire ability to influence long-term yields comes from investors' perceptions of what fundamentals are going to be *after* normal conditions have returned. A policy that only kept the supplies off the market during the ZLB episode itself would have much more limited potential. In this sense, this particular form of nonstandard monetary policy could end up having limited effectiveness for the same reasons as policies that hope to influence the public's expectation of what the target will be for short-term interest rates once the economy escapes from the ZLB.

Our estimated effects are linear— twice as big a purchase is predicted to have twice as big an effect on yields. But this is simply an assumption of our empirical estimation strategy and not a proposition we have tested directly in the data. Particularly since the magnitudes under discussion are so different from the observed historical variations from which our estimates were inferred, extrapolation of these effects to larger and larger policy measures is of necessity an uncertain exercise.

We also noted that, although we have framed the discussion here in terms of options available to the Federal Reserve, this policy tool could in many ways more naturally be implemented by the Treasury itself altering the term structure of debt that it issues. If the Treasury has sound reasons not to do so, it is unclear why the Federal Reserve should try to undo the Treasury's attempted hedging of the unified government's balance sheet with respect to interest rate risk. Conversely, if the Fed has good reasons to try to flatten the slope of the yield curve, it is unclear why the Treasury should resist being the agent to implement the plan.

References

- Adrian, Tobias, Karin Kimbrough, and Dina Marchioni**, “The Federal Reserve’s commercial paper funding facility,” *FRBNY Economic Policy Review*, 2010.
- Auerbach, Alan J. and Maurice Obstfeld**, “The case for open-market purchases in a liquidity trap,” *American Economic Review*, 2005, *95*, 110–137.
- Baumeister, Christiane and Luca Benati**, “Unconventional monetary policy and the Great Recession: Estimating the impact of a compression in the yield spread at the zero lower bound,” 2010. Mimeo, European Central Bank.
- Bernanke, Ben S., Vincent R. Reinhart, and Brian P. Sack**, “Monetary policy alternatives at the zero bound: An empirical assessment,” *Brookings Papers on Economic Activity*, 2004, *2*, 1–100.
- Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch**, “Do central bank liquidity facilities affect interbank lending rates?,” 2009. Federal Reserve Bank of San Francisco, Working paper 2009-13.
- Cochrane, John H. and Monika Piazzesi**, “Bond Risk Premia,” *American Economic Review*, 2005, *95*, 138–160.
- D’Amico, Stefania and Thomas B. King**, “Flow and stock effects of large-scale Treasury purchases,” 2010. Working paper, Federal Reserve Board.
- Doh, Taeyoung**, “The efficacy of large-scale asset purchases at the zero lower bound,” *Federal Reserve Bank of Kansas City Economic Review*, 2010, *Q2*, 5–34.

- Duygan-Bump, Burcu, Patrick M. Parkinson, Eric S. Rosengren, Gustavo A. Suarez, and Paul S. Willen**, “How effective were the Federal Reserve emergency liquidity facilities? Evidence from the asset-backed commercial paper money market mutual fund liquidity facility,” 2010. Working paper, Federal Reserve Bank of Boston.
- Eggertsson, Gauti and Michael Woodford**, “The zero bound on interest rates and optimal monetary policy,” *Brookings Papers on Economic Activity*, 2003, 1, 139–211.
- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack**, “Large-scale asset purchases by the Federal Reserve: Did they work?,” 2010. Federal Reserve Bank of New York Staff Reports.
- Greenwood, Robin and Dimitri Vayanos**, “Bond supply and excess bond returns,” 2010. Working Paper, London School of Economics.
- Gürkaynak, Refet S. and Jonathan H. Wright**, “Macroeconomics and the term structure,” 2010. Working Paper, Johns Hopkins University.
- , **Brian Sack, and Jonathan H. Wright**, “The U.S. Treasury Yield Curve: 1961 to the present,” *Journal of Monetary Economics*, November 2007, 54 (8), 2291–2304.
- Hamilton, James D. and Jing Wu**, “Identification and estimation of affine-term-structure models,” 2010. Working paper, University of California, San Diego.
- Hancock, Diana and Wayne Passmore**, “Did the Federal Reserve’s MBS Purchase Program Lower Mortgage Rates?,” 2011. Working paper, Federal Reserve Board.

- Kitchen, John and Menzie D. Chinn**, “Financing U.S. debt: Is there enough money in the world– and at what cost?,” 2010. Working paper, University of Wisconsin-Madison.
- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, “The aggregate demand for treasury debt,” 2010. Working paper, Northwestern University.
- Krugman, Paul R.**, “It’s baaack: Japan’s slump and the return of the liquidity trap,” *Brookings Papers on Economic Activity*, 1998, 2, 137–205.
- Kuttner, Kenneth N.**, “Can central banks target bond prices?,” 2006. NBER Working paper 12454.
- McAndrews, James, Asani Sarkar, and Zhenyu Wang**, “The effect of the term auction facility on the London inter-bank offered rate,” 2008. Working paper, Federal Reserve Bank of New York.
- McCallum, Bennett T.**, “Theoretical analysis regarding a zero lower bound on nominal interest rates,” *Journal of Money, Credit and Banking*, 2000, 32, 870–904.
- Merton, Robert C.**, “Lifetime portfolio selection under uncertainty: The continuous time case,” *Review of Economics and Statistics*, 1969, 51, 247–257.
- Roley, V. Vance**, “The effect of federal debt-management policy on corporate bond and equity yields,” *Quarterly Journal of Economics*, 1982, 97, 645–668.
- Stambaugh, Robert F.**, “Predictive Regressions,” *Journal of Financial Economics*, 1999, 54, 375–421.

Svensson, Lars E.O., “The zero bound in an open economy: A foolproof way of escaping from a liquidity trap,” *Monetary and Economic Studies*, 2001, *19(S-1)*, 277–312.

Swanson, Eric T., “Let’s Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2,” 2011. Working paper, Federal Reserve Bank of San Francisco.

Taylor, John B. and John C. Williams, “A black swan in the money market,” *American Economic Journal: Macroeconomics*, 2009, *1*, 58–83.

Vayanos, Dimitri and Jean-Luc Vila, “A preferred-habitat model of the term structure of interest rates,” 2009. Working paper, London School of Economics.

Woodford, Michael, *Comments on Ashcraft, Garleanu, and Pedersen, ‘Two monetary tools: interest rates and haircuts’*, Vol. 25 of *NBER Macroeconomics Annual 2010 (forthcoming)*, National Bureau of Economic Research, 2010.

Appendix A. Details of the arbitrageurs' portfolio optimization problem.

Let P_{nt} denote the price of a pure-discount n -period bond (with $P_{0t} = 1$), W_t the total wealth of the arbitrageurs, and z_{nt} the portion of their wealth allocated to each bond maturity. Then the arbitrageurs' wealth evolves according to

$$W_{t+1} = \sum_{n=1}^N z_{nt} \frac{P_{n-1,t+1}}{P_{nt}} W_t$$

with associated rate of return

$$r_{t,t+1} = \frac{W_{t+1} - W_t}{W_t} = \sum_{n=1}^N z_{nt} \left[\frac{P_{n-1,t+1}}{P_{nt}} - 1 \right].$$

If the change in prices between t and $t + 1$ is small,²² the portfolio's mean return and variance can be approximated

$$E_t r_{t,t+1} \approx -z_{1t}(\bar{a}_1 + \bar{b}'_1 f_t) + \sum_{n=2}^N z_{nt} \left[\bar{a}_{n-1} + \bar{b}'_{n-1}(c + \rho f_t) + (1/2)\bar{b}'_{n-1} \Sigma \Sigma' \bar{b}_{n-1} - \bar{a}_n - \bar{b}'_n f_t \right] \quad (31)$$

$$\text{Var}_t(r_{t,t+1}) \approx d'_t \Sigma \Sigma' d_t \quad (32)$$

where the $(J \times 1)$ vector d_t summarizes exposures to each of the J factor risks associated with holding the $(N \times 1)$ vector of bonds z_t . The arbitrageurs thus choose z_t so as to maximize (1) subject to (31), (32), (8), and $\sum_{n=1}^N z_{nt} = 1$, for which the first-order condition is given by (6).

²²Suppose that

$$q_{n,t+1} \equiv \frac{(P_{n-1,t+1} - P_{nt})}{P_{nt}} = \exp(\mu_n h + \sqrt{h} \varepsilon_{n,t+1}) - 1$$

where $(\varepsilon_{1,t+1}, \dots, \varepsilon_{N,t+1})' \sim N(0, \Omega)$. Our approximation is derived from the limiting behavior as h becomes small, analogous to those obtained when considering a continuous-time representation of a discrete-time process. Thus as in Merton (1969),

$$E_t \left(\sum_{n=1}^N z_{nt} q_{n,t+1} \right) = \sum_{n=1}^N z_{nt} [\mu_n h + \Omega_{nn} h/2 + o(h)]$$

$$\text{Var}_t \left(\sum_{n=1}^N z_{nt} q_{n,t+1} \right) = z'_t \Omega z_t h + o(h)$$

for Ω_{nn} the row n , column n element of Ω and $z_t = (z_{1t}, \dots, z_{Nt})'$. Equations (31) and (32) are obtained by setting $h = 1$ and $o(h) = 0$. Specifically,

$$\frac{P_{n-1,t+1}}{P_{nt}} = \exp(\bar{a}_{n-1} + \bar{b}'_{n-1} f_{t+1} - \bar{a}_n - \bar{b}'_n f_t)$$

$$\mu_n = \bar{a}_{n-1} + \bar{b}'_{n-1}(c + \rho f_t) - \bar{a}_n - \bar{b}'_n f_t$$

$$\Omega_{nn} = \bar{b}'_{n-1} \Sigma \Sigma' \bar{b}_{n-1}.$$

Appendix B. Arbitrage-free equilibrium.

Note that $y_{nt} = -n^{-1}p_{nt} = -n^{-1}(\bar{a}_n + \bar{b}'_n f_t)$ and suppose that $\zeta_{nt} = \zeta_n + \vartheta'_n f_t$. If we multiply (9) by \bar{b}_{n-1} and sum over $n = 2, \dots, N$, we find using (8) that equilibrium requires

$$d_t = \sum_{n=2}^N \bar{b}_{n-1} \left[\zeta_n + \vartheta'_n f_t + (\alpha_n/n)(\bar{a}_n + \bar{b}'_n f_t) \right].$$

Equation (10) is obtained from (7) with

$$\lambda = \gamma \Sigma' \sum_{n=2}^N \bar{b}_{n-1} [\zeta_n + (\alpha_n/n)\bar{a}_n]$$

$$\Lambda = \gamma \Sigma' \sum_{n=2}^N \bar{b}_{n-1} [\vartheta'_n + (\alpha_n/n)\bar{b}'_n].$$

Appendix C. ATSM estimation for a just-identified model.

We first estimate the parameters of (5) and (16) by OLS:

$$\begin{bmatrix} \hat{c} & \hat{\rho} \end{bmatrix} = \left(\sum_{t=2}^T f_t \begin{bmatrix} 1 & f'_{t-1} \end{bmatrix} \right) \left(\sum_{t=2}^T \begin{bmatrix} 1 \\ f_{t-1} \end{bmatrix} \begin{bmatrix} 1 & f'_{t-1} \end{bmatrix} \right)^{-1}$$

$$\hat{\Sigma} \hat{\Sigma}' = (T-1)^{-1} \sum_{t=2}^T (f_t - \hat{c} - \hat{\rho} f'_{t-1})(f_t - \hat{c} - \hat{\rho} f'_{t-1})'$$

$$\begin{bmatrix} \hat{A} & \hat{B} \end{bmatrix} = \left(\sum_{t=1}^T Y_{2t} \begin{bmatrix} 1 & f'_t \end{bmatrix} \right) \left(\sum_{t=1}^T \begin{bmatrix} 1 \\ f_t \end{bmatrix} \begin{bmatrix} 1 & f'_t \end{bmatrix} \right)^{-1}$$

$$\hat{\Sigma}_e \hat{\Sigma}'_e = T^{-1} \sum_{t=1}^T (Y_{2t} - \hat{A} - \hat{B}f_t)(Y_{2t} - \hat{A} - \hat{B}f_t)'$$

The predicted value for row i of \hat{B} is given by

$$\hat{B}'_i = n_i^{-1} b'_1 \left[I_J + \rho^Q + (\rho^Q)^2 + \dots + (\rho^Q)^{n_i-1} \right] \quad \text{for } i = 1, \dots, M.$$

For the just-identified case with $M = J + 1$, we solve this $[(J + 1) \times J]$ system of equations for the $J(J + 1)$ unknowns ρ^Q and b_1 using numerical search. Taking these values for ρ^Q and b_1 as given, we can then use (11) to solve for \bar{b}_n for any desired n along with

$$\bar{a}_n = n\bar{a}_1 + \sum_{\ell=1}^n \bar{b}'_{\ell-1} c^Q + (1/2) \sum_{\ell=1}^n \bar{b}'_{\ell-1} \Sigma \Sigma' \bar{b}_{\ell-1}.$$

The $J + 1$ values for a_1 and c^Q are then found by numerical solution of the $J + 1$ equations

$$\hat{A}_i = -n_i^{-1} \bar{a}_{n_i} \quad \text{for } i = 1, \dots, M.$$

Appendix D. ATSM estimation for an overidentified model.

We estimated (27) by unconstrained OLS,

$$\begin{bmatrix} \hat{A}_2^\dagger & \hat{B}_2^\dagger \end{bmatrix} = \left(\sum_{t=1}^T Y_{2t} \begin{bmatrix} 1 & Y'_{1t} \end{bmatrix} \right) \left(\sum_{t=1}^T \begin{bmatrix} 1 \\ Y_{1t} \end{bmatrix} \begin{bmatrix} 1 & Y'_{1t} \end{bmatrix} \right)^{-1}$$

for which the inverse of the usual variance matrix for the estimated coefficients is given by

$$\hat{R} = \hat{\Omega}_e^{-1} \otimes T^{-1} \sum_{t=1}^T \begin{bmatrix} 1 \\ Y_{1t} \end{bmatrix} \begin{bmatrix} 1 & Y'_{1t} \end{bmatrix}$$

with $\hat{\Omega}_e$ given by diagonal elements of

$$T^{-1} \sum_{t=1}^T (Y_{2t} - \hat{A}_2^\dagger - \hat{B}_2^\dagger Y_{1t})(Y_{2t} - \hat{A}_2^\dagger - \hat{B}_2^\dagger Y_{1t})'.$$

The minimum-chi-square estimation procedure proposed by Hamilton and Wu (2010) estimates the structural parameters of interest $\theta = (\pi^Q, a_1^*)'$ or $(\pi^Q, a_1^*, a_1)'$ by minimizing

$$T[\hat{\pi} - g(\theta)]' \hat{R}[\hat{\pi} - g(\theta)] \quad (33)$$

where $\hat{\pi} = \text{vec}\left(\begin{bmatrix} \hat{A}_2^\dagger & \hat{B}_2^\dagger \end{bmatrix}'\right)$ and $g(\theta)$ denotes the corresponding predicted value from (28).

Under the null hypothesis that the model is correctly specified, the minimal value achieved for (33) should have an asymptotic $\chi^2(k_1 - k_0)$ distribution, where $k_1 = 14$ is the number of parameters in \hat{A}_2^\dagger and \hat{B}_2^\dagger and $k_0 = 2$ or 3 is the number of elements in θ .

Appendix E. Details of data construction.

Following Greenwood and Vayanos (2010), we started with CRSP data for outstanding Treasury debt by individual CUSIP number to estimate outstanding nominal Treasury debt at the end of each month. We calculated n for each issue by calculating the number of days

between maturity and the last Friday of the month, and converted to weeks by rounding up. The raw source for these data appears to be the *Monthly Statement of the Public Debt of the United States*. We checked these data by summing all the maturities and comparing this sum with the sum of nominal bills, bonds, and notes recorded in the Haver database,²³ which also comes from the same *Monthly Statement*. We found numerous discrepancies, which came from such factors as the CRSP files on occasion missing individual CUSIP series and at other times having incorporated assorted data entry errors. We were able to correct CRSP data errors so as to reduce almost all discrepancies to less than \$200 M by hand comparison of the CRSP numbers with individual copies of the *Monthly Statement* itself.

Although the Federal Reserve currently reports outright Treasury holdings for the System Open Market Account by individual CUSIP, we were unable to secure access to historical archives of these, and settled for rough estimates constructed as follows. The Federal Reserve's weekly H41 release²⁴ reports SOMA each Wednesday by rough maturity breakdowns (less than 15 days, 16-90 days, 91 days to 1 year, over 1 year to 5 years, over 5 years to 10 years, and over 10 years), and we matched up the last Wednesday of each month for SOMA holdings with the last calendar day of the month for Treasury marketable debt. Unfortunately, the reported SOMA maturity categories include both nominal Treasuries as well as TIPS, which we exclude from our analysis. Our solution was to assume that Fed holdings of TIPS as a fraction of the Fed's total holdings of notes and bonds was the same across all maturity categories. Total Fed holdings of notes and bonds are reported on the H41, as are total TIPS holdings (though

²³We thank Christiane Baumeister for sharing these Haver data.

²⁴Available in Table 2 of <http://www.federalreserve.gov/datadownload/Choose.aspx?rel=H41>. Prior to June 2003, we used the end-of-calendar month data compiled by Kuttner (2006) available at <http://econ.williams.edu/people/knk1/research>.

prior to December 2002, we had to read the latter by hand from the notes section of individual reports). We then multiplied each maturity category greater than 1 year by this ratio to get an estimate of total TIPS holdings in those categories. For maturity categories less than 1 year, we multiplied by the product of this ratio with the ratio of the Fed's notes and bonds of maturity less than 1 year to the Fed's total Treasury securities less than one year. We then subtracted the resulting estimates of TIPS holdings within each maturity category from the reported total holdings within each category to get our estimate of nominal Fed holdings for each maturity category. We then allocated this ratio evenly across total outstanding Treasury securities of each weekly maturity falling within that category to arrive at our estimate of how much of those securities were held by the Federal Reserve's SOMA.

Estimated parameters					Implied parameters			
c^Q	0.0116 (0.0002)	-0.0118 (0.0005)	-0.0036 (0.0007)		λ	-0.1378 (0.0717)	0.1604 (0.0727)	-0.0564 (0.0687)
ρ^Q	0.9990 (0.0001)	0.0094 (0.0002)	-0.0140 (0.0005)		Λ	-0.0867 (0.0468)	-0.0480 (0.0594)	-0.0948 (0.1203)
	0.0027 (0.0003)	0.9870 (0.0004)	0.0330 (0.0010)			0.0847 (0.0455)	-0.0266 (0.0825)	0.1773 (0.1200)
	-0.0018 (0.0002)	-0.0028 (0.0002)	0.9867 (0.0008)			-0.0567 (0.0436)	0.0531 (0.0596)	-0.1862 (0.1594)
c	-0.0034 (0.0089)	-0.0003 (0.0074)	0.0006 (0.0066)					
ρ	0.9895 (0.0072)	0.0042 (0.0081)	-0.0244 (0.0157)					
	0.0083 (0.0047)	0.9826 (0.0081)	0.0478 (0.0123)					
	-0.0013 (0.0041)	0.0055 (0.0058)	0.9755 (0.0132)					
$a_1 \times 5200$	4.1158 (0.0074)							
$b_1 \times 5200$	1.0345 (0.0058)	-0.6830 (0.0081)	0.6311 (0.0189)					
Σ	0.1094 (0.0236)	0	0					
	0.0360 (0.0100)	0.1027 (0.0045)	0					
	-0.0670 (0.0188)	0.0025 (0.0130)	0.0968 (0.0149)					
$\Sigma_e \times 5200$	0.0978 (0.0023)	0	0	0				
	0	0.0674 (0.0016)	0	0				
	0	0	0.0531 (0.0013)	0				
	0	0	0	0.1171 (0.0028)				

Table 1: Parameter estimates for the weekly affine-term-structure model, Jan 5, 1990 to July 27, 2007. Small-sample standard errors in parentheses. Sample size: $T = 917$. Bold indicates statistically significantly different from zero at the 5% significance level.

Regressors	6m over 3m	1yr over 6m	2y over 1y	5y over 1y	10y over 1y
c, f_t^*	0.357 (0.000)	0.356 (0.000)	0.331 (0.000)	0.295 (0.000)	0.331 (0.000)
c, f_t, z_t^{A*}	0.410 (0.020)	0.420 (0.119)	0.373 (0.311)	0.300 (0.728)	0.336 (0.665)
c, f_t, z_t^{L*}	0.428 (0.003)	0.501 (0.008)	0.524 (0.006)	0.398 (0.035)	0.357 (0.196)
c, f_t, z_t^{pc*}	0.368 (0.001)	0.361 (0.007)	0.333 (0.062)	0.297 (0.098)	0.334 (0.051)
c, f_t, v_t^*	0.385 (0.016)	0.409 (0.001)	0.388 (0.006)	0.339 (0.008)	0.338 (0.227)
c, f_t, q_t^*	0.444 (0.002)	0.568 (0.000)	0.714 (0.000)	0.617 (0.000)	0.549 (0.001)
c, f_t, z_t^{pc}, q_t^*	0.452 (0.002)	0.571 (0.000)	0.717 (0.000)	0.618 (0.000)	0.550 (0.002)
c, f_t, v_t, q_t^*	0.458 (0.001)	0.595 (0.000)	0.737 (0.000)	0.640 (0.000)	0.552 (0.002)
$c, f_t, z_t^A, z_t^L, q_t^*$	0.476 (0.000)	0.597 (0.001)	0.741 (0.000)	0.670 (0.002)	0.634 (0.054)

Table 2: R^2 and hypothesis tests for holding-return forecasting regressions. Reported numbers are the R^2 for the regressions, with p -values in parentheses, for tests of the null hypothesis that coefficients on starred variables are zero. All regressions also include a constant term (denoted by c) and all hypothesis tests use Newey-West variance matrix with 20 lags. Bold indicates coefficients on starred variables are statistically significantly different from zero at the 5% significance level.

	F test	$\phi_i' \Delta$
level	3.256 (0.023)	0.005 (0.112)
slope	4.415 (0.005)	-0.250 (0.116)
curvature	2.672 (0.049)	-0.073 (0.116)

Table 3: Granger-causality tests and scenario impact estimates for factor vector autoregression. First column reports F test (p -value in parentheses) of null hypothesis that $\phi_i = 0$ in regression $f_{it} = c_i + \rho_i' f_{t-1} + \phi_i' q_{t-1} + \varepsilon_{it}$. Second column reports estimate of $\phi_i' \Delta$ for that regression (with standard error) for Δ the average change in q under the alternative scenario.

	Contemporaneous		Forecast	
	restricted	unrestricted	restricted	unrestricted
3m	0.625	0.668	0.522	0.602
1y	0.891	0.924	0.652	0.767
5y	0.961	0.975	0.753	0.753
30y	0.965	0.972	0.735	0.787

Table 4: R^2 for post-crisis sample (March 3, 2009 to Aug 10, 2010) for unrestricted OLS fit to post-crisis data and for prediction constructed from pre-crisis parameter estimates together with post-crisis estimates of π^Q and a_1^* . Contemporaneous: prediction of y_{nt} given current 6-month, 2-year and 10-year yields. Forecast: predictions of y_{nt} given lagged 6-month, 2-year and 10-year yields.

Study	Measure	Original estimates		Hamilton-Wu estimates	
		Pre-crisis	ZLB	Pre-crisis	ZLB
Gagnon, et. al.	10 yr yield	-20		-14	-13
Greenwood-Vayanos	5yr-1yr spread	-39		-17	-9
	20yr-1yr spread	-74		-25	-18
D'Amico-King	10yr yield		-67	-14	-13
Deutsche Bank	10yr yield		-20	-14	-13

Table 5: Comparison of different estimates of the effect of replacing \$400 billion in long-term debt with short-term debt.

Yield	Factor loadings			Normal effect	ZLB effect
	level	slope	curvature		
10-year Treasury	1.000	0.500	0.167	-14	-13
Aaa Corporate	0.883	0.453	0.379	-14	-15
Baa Corporate	0.888	0.441	0.535	-14	-17
30-year Mortgage	0.933	0.363	0.325	-11	-13

Table 6: Empirical loadings of selected yields on Treasury level, slope and curvature factors, and predicted effect on yield (in basis points) of selling \$400 billion in short-term Treasury debt and buying \$400 billion in long-term Treasury debt.

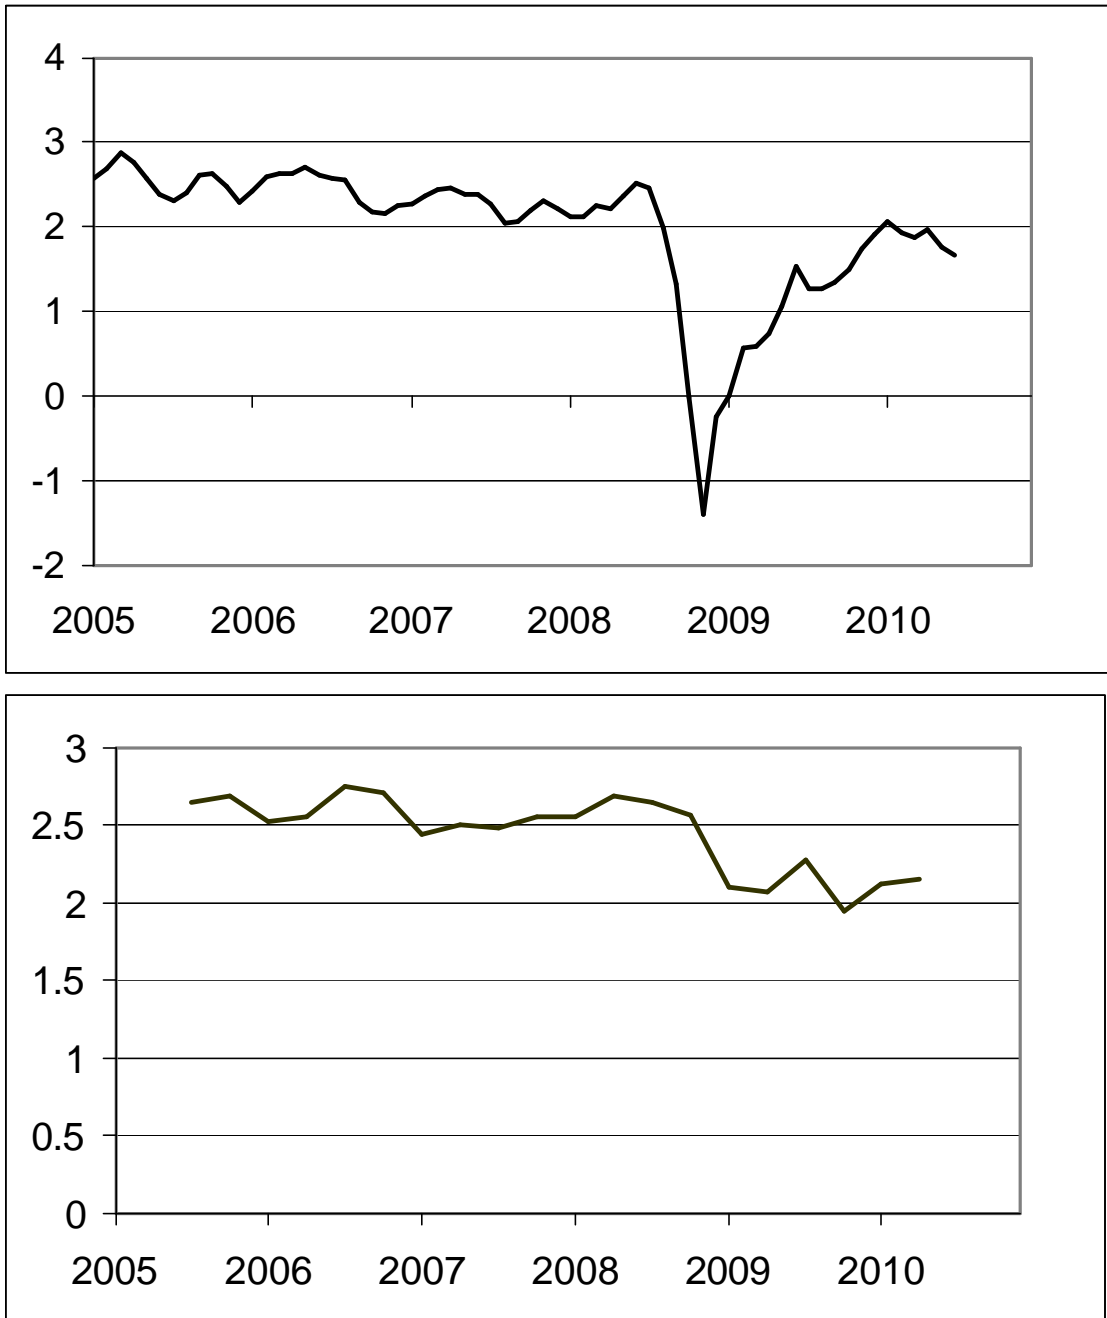


Figure 1. Alternative measures of 5-year expected inflation. Top panel: 5-year break-even inflation rate, calculated as nominal yield of 5-year Treasury bond minus nominal yield on 5-year Treasury Inflation Protected Security, 2005:M1-2010:M6. Data source: FRED database of the Federal Reserve Bank of St. Louis. Bottom panel: 5-year expected CPI inflation, from the average response of the Survey of Professional Forecasters, 2005:Q3 to 2010:Q2. Data source: Federal Reserve Bank of Philadelphia.

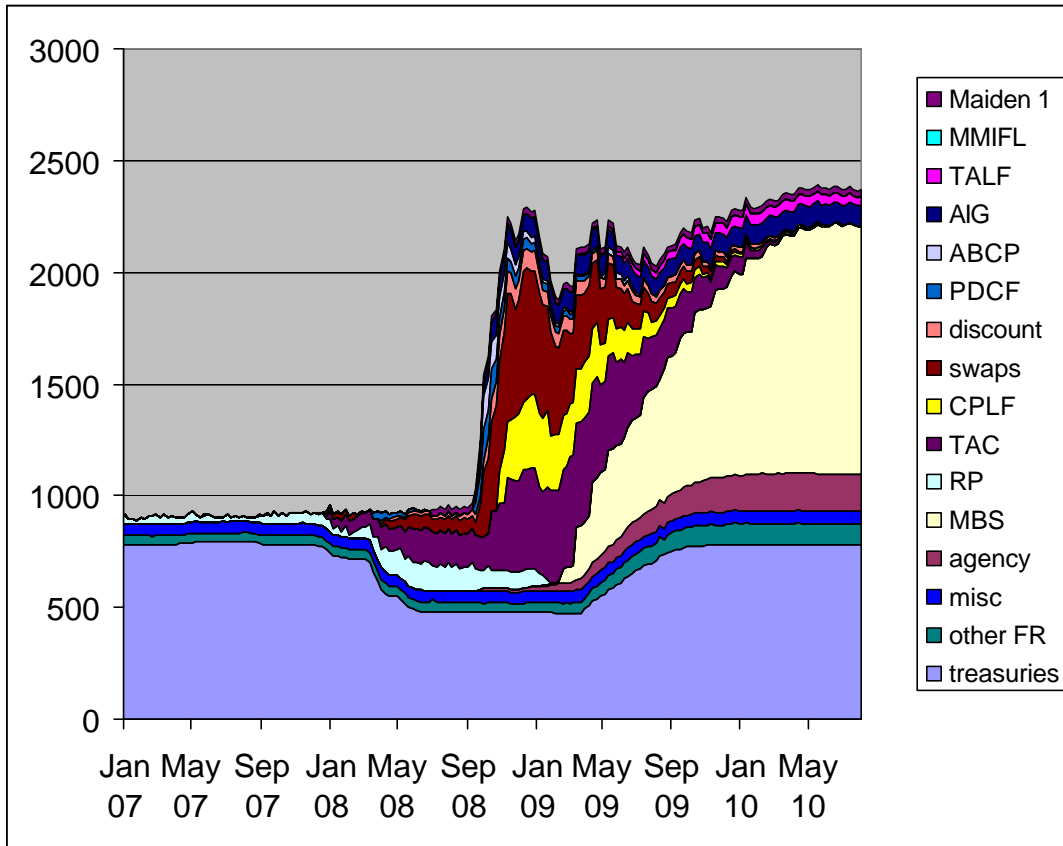


Figure 2. Federal Reserve assets, in billions of dollars, Jan 3, 2007 to Aug 4, 2010, Wednesday values, seasonally unadjusted, from Federal Reserve H41 release. *Maiden 1*: net portfolio holdings of Maiden Lane LLC; *MMIFL*: net portfolio holdings of LLCs funded through the Money Market Investor Funding Facility; *TALF*: loans extended through Term Asset-Backed Securities Loan Facility; *AIG*: sum of credit extended to American International Group, Inc. plus net portfolio holdings of Maiden Lane II and III; *ABCP*: loans extended to Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility; *PDCF*: loans extended to primary dealer and other broker-dealer credit; *discount*: sum of primary credit, secondary credit, and seasonal credit; *swaps*: central bank liquidity swaps; *CPLF*: net portfolio holdings of LLCs funded through the Commercial Paper Funding Facility; *TAC*: term auction credit; *RP*: repurchase agreements; *MBS*: mortgage-backed securities held outright; *agency*: federal agency debt securities held outright; *misc*: sum of float, gold stock, special drawing rights certificate account, and Treasury currency outstanding; *other FR*: Other Federal Reserve assets; *treasuries*: U.S. Treasury securities held outright.

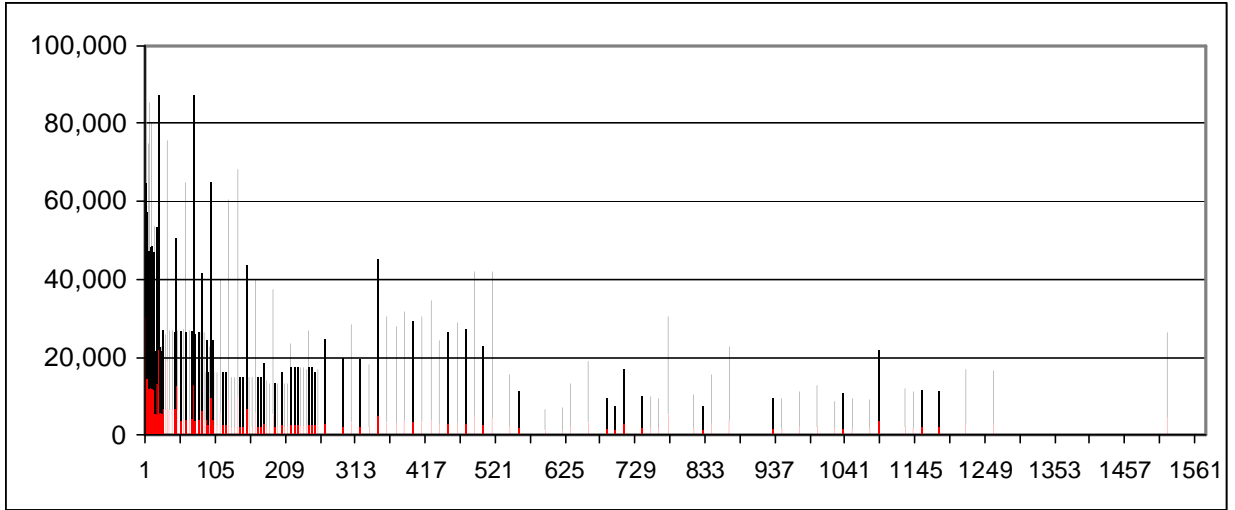


Figure 3. Maturity structure of U.S. federal debt as of December 31, 2006. Horizontal axis: maturity in weeks. Black bars: face value of marketable nominal Treasury securities of that maturity, in millions of dollars. Red bars: imputed holdings of the System Open Market Account of the U.S. Federal Reserve.

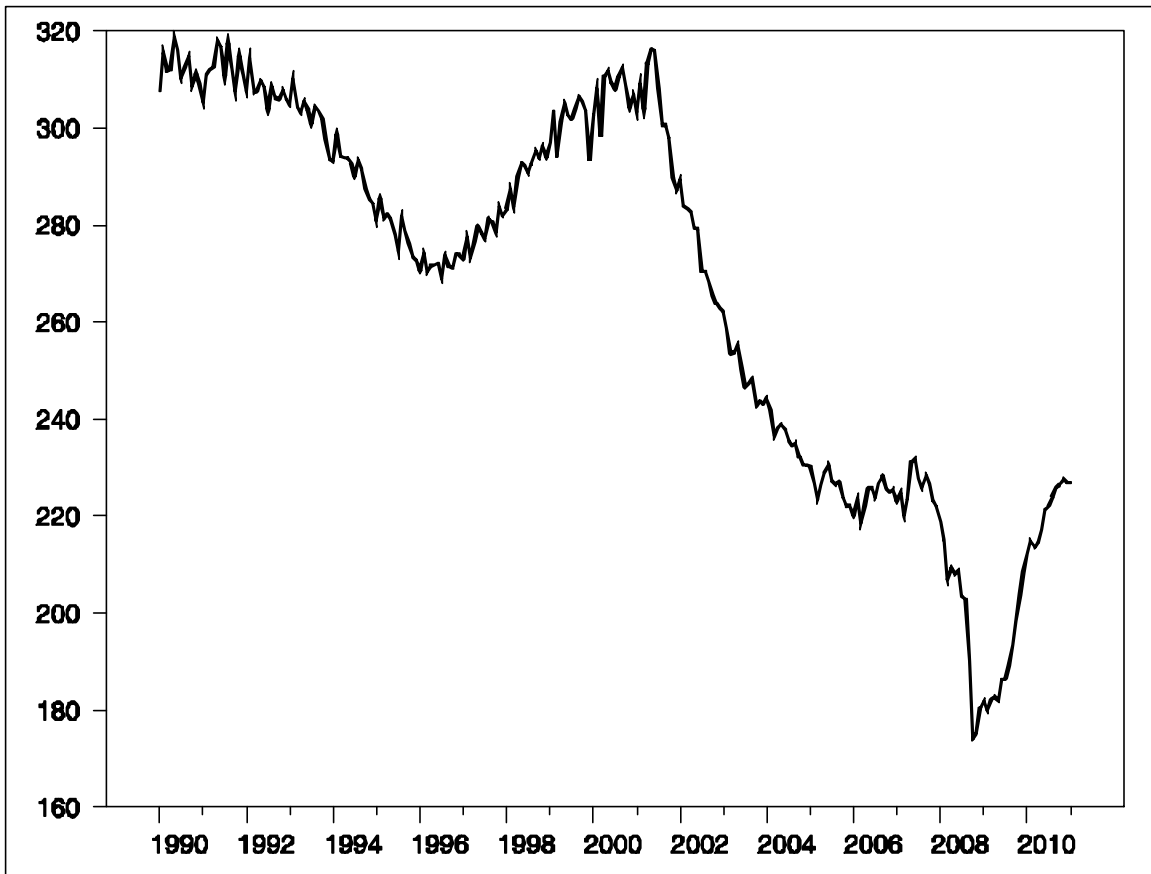


Figure 4. Average maturity in weeks of debt held by the public, plotted monthly from Jan 31, 1990 to Jan 31, 2011.

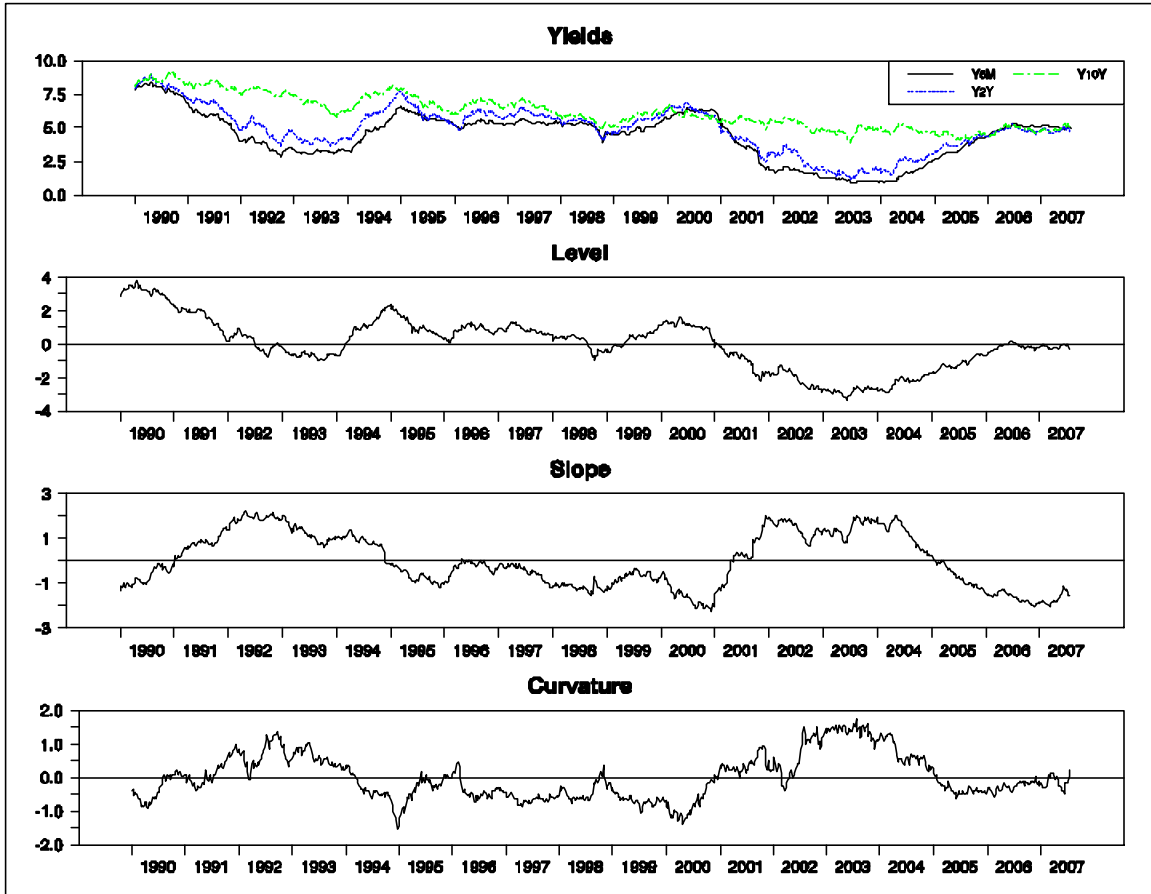


Figure 5. Yields and factors used in baseline estimation, weekly from Jan 5, 1990 to July 27, 2007.

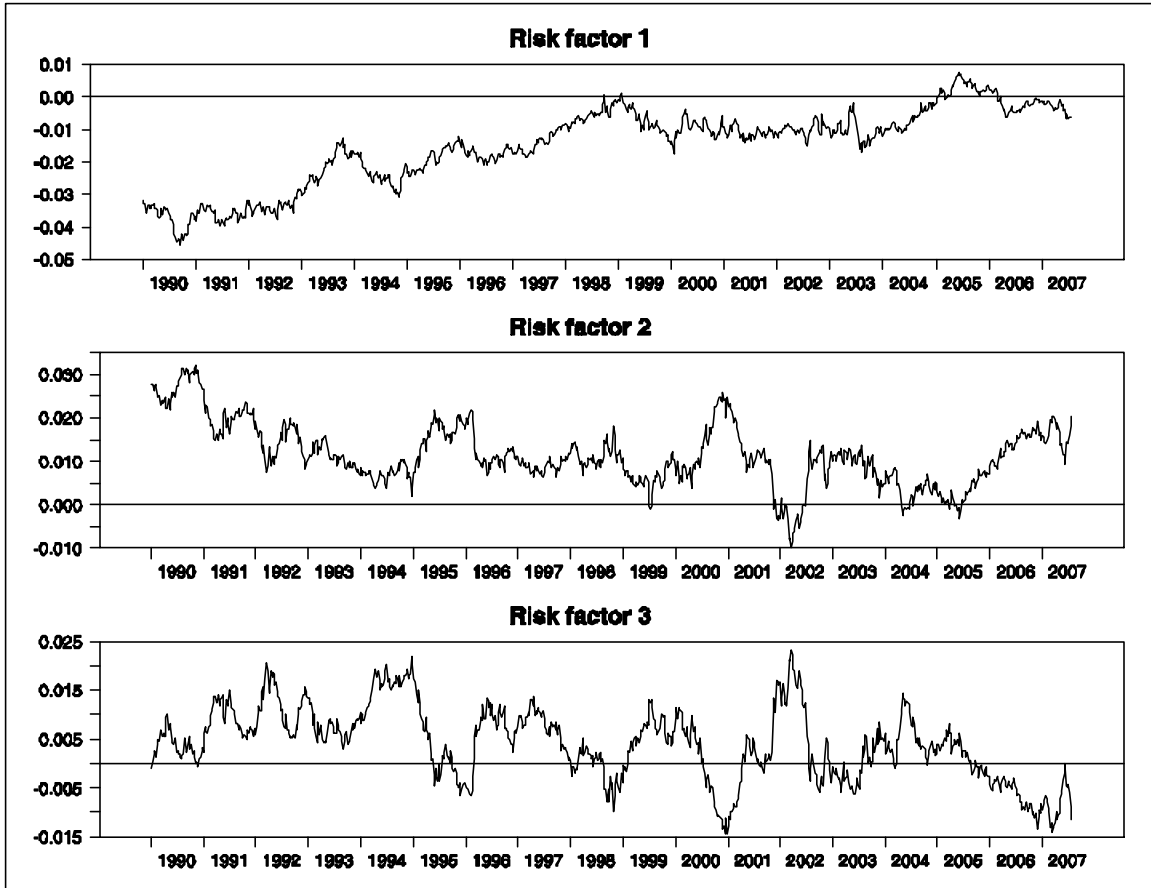


Figure 6. Values of the three elements of ΣI_t , weekly from Jan 5, 1990 to July 27, 2007.

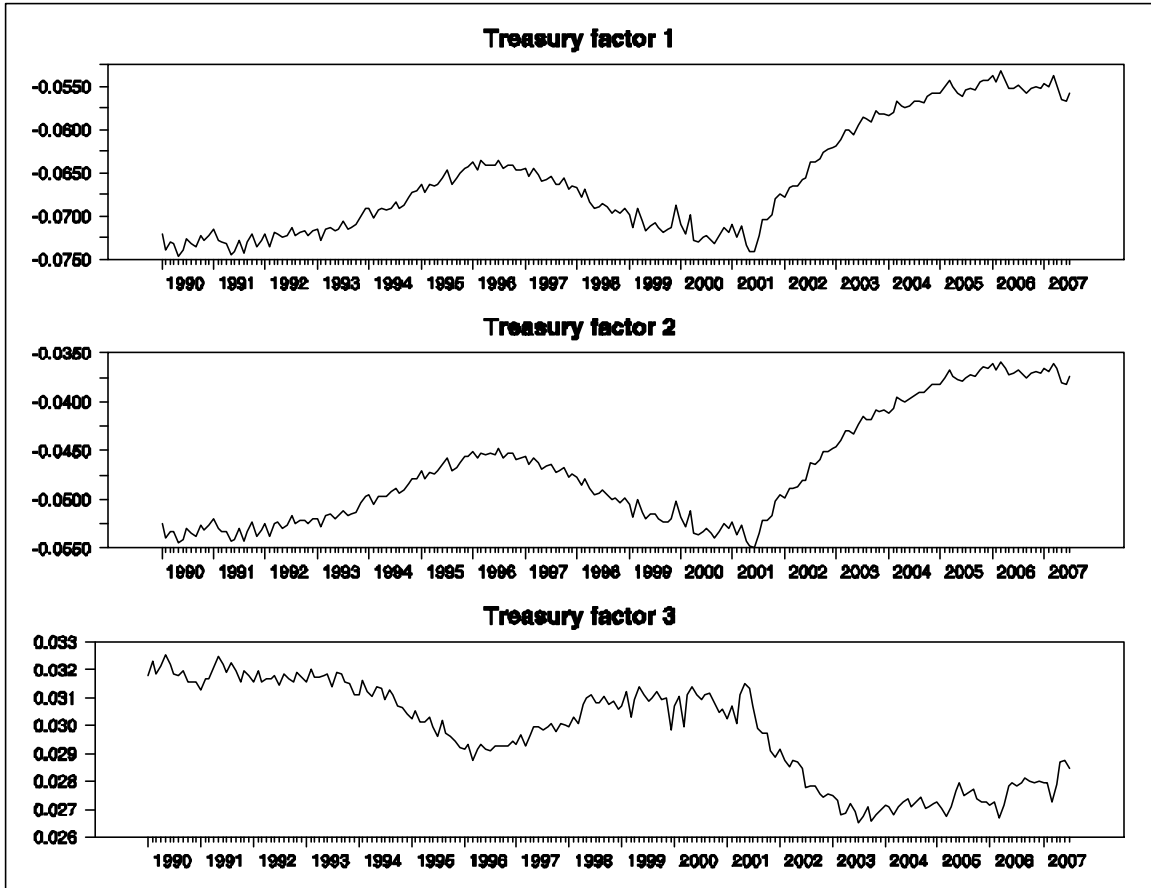


Figure 7. Values of the three elements of $q_t = 100 \Sigma \Sigma' \sum_{n=1}^{N-1} z_{nt} \bar{b}_{n-1}$ monthly from Jan 31, 1990 to July 31, 2007.

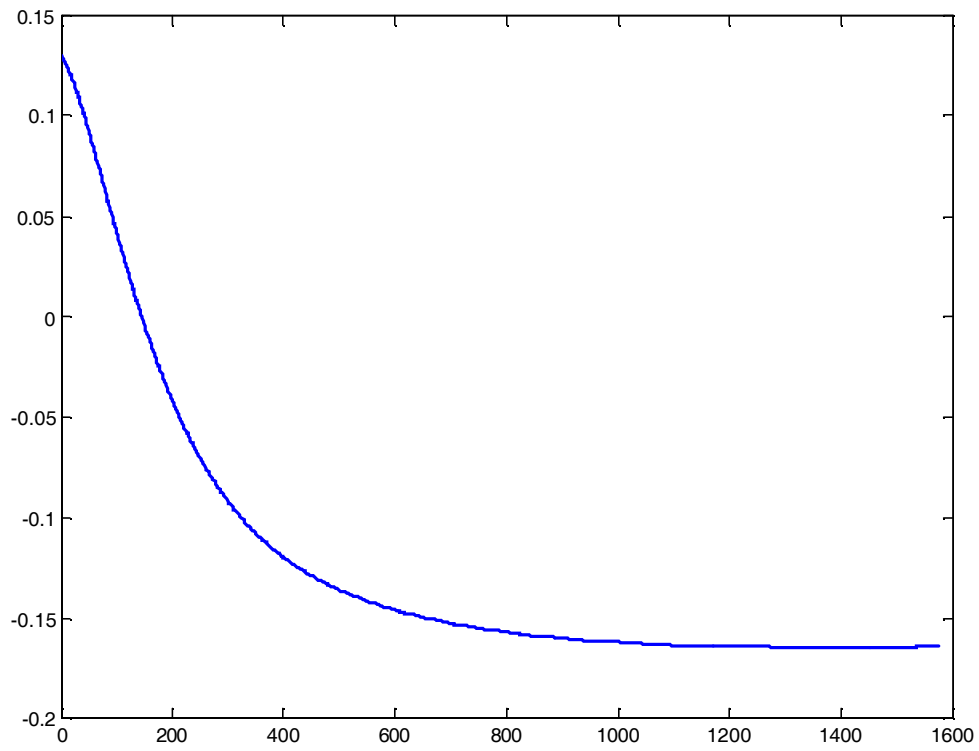


Figure 8. Predicted change in $y_{n,t+1}$ (quoted in annual percentage points) as a function of weeks to maturity n in response to shift in q_t of size ? .

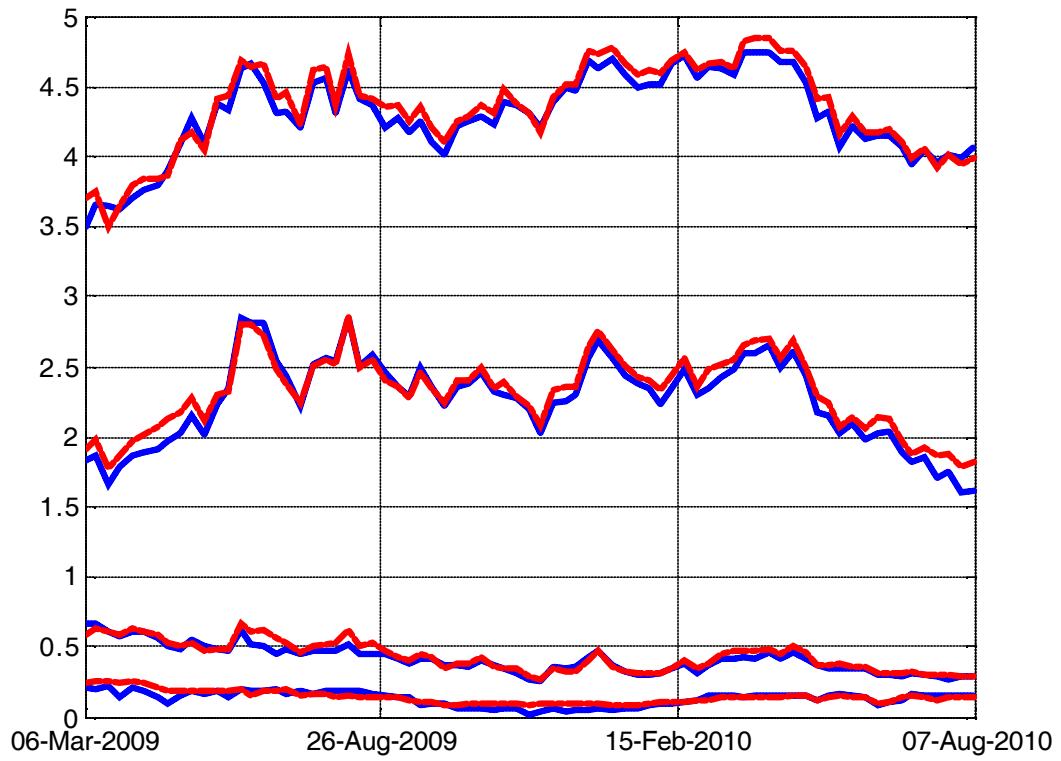


Figure 9. Actual (solid) and predicted (dashed) behavior of selected interest rates, weekly from March 7, 2009 to August 10, 2010. Rates shown (in order from top to bottom) are the 30 year, 5 year, 1 year, and 3 month.

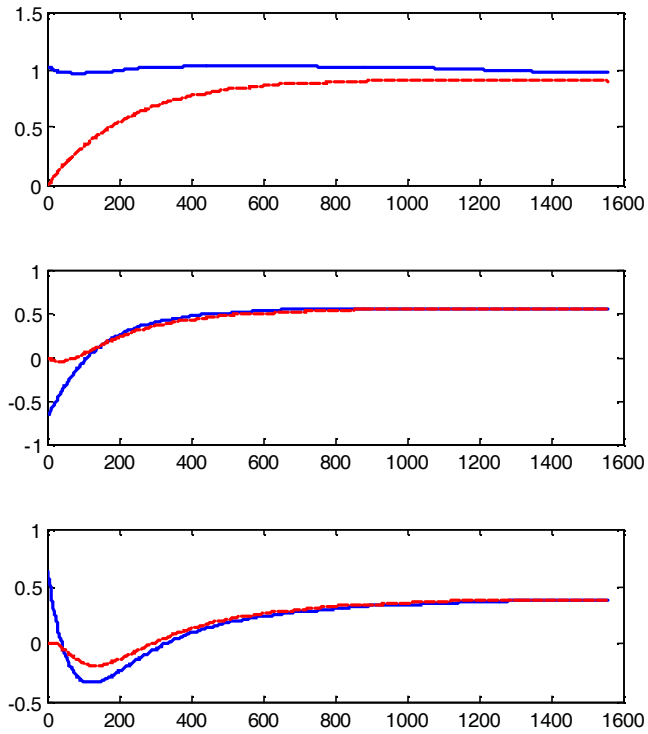


Figure 10. Factor loadings. Solid curves: normal loadings (plots of $5200 b_n$ as function of maturity n in weeks). Dashed curves: zero-lower-bound loadings ($5200 b_n^*$). Top panel: level loadings; middle panel: slope loadings; bottom panel: curvature loadings.

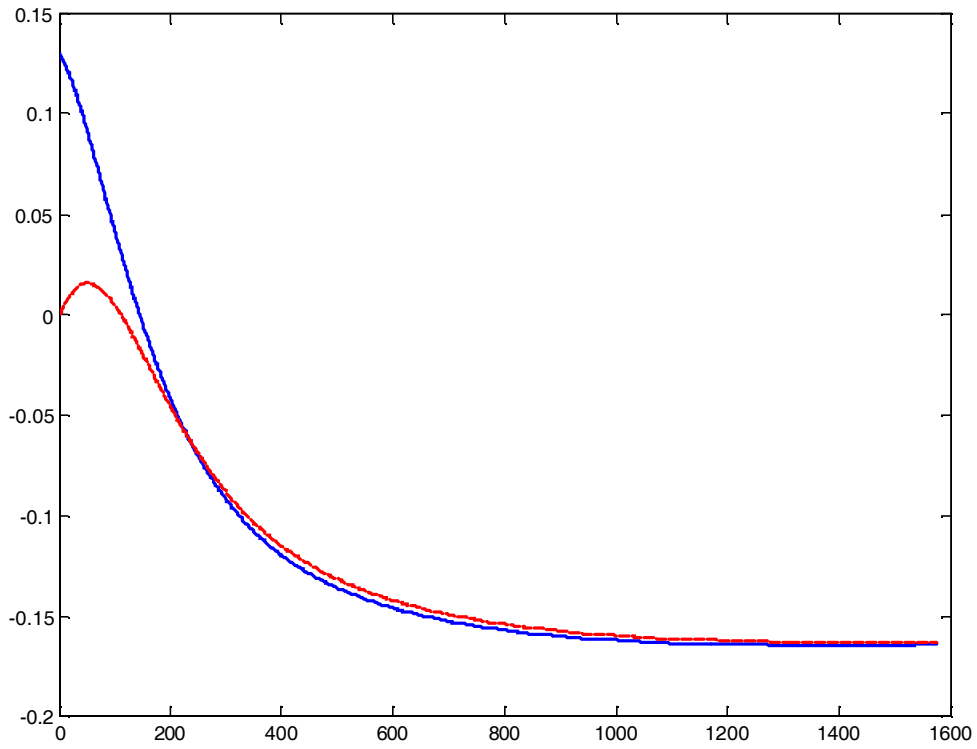


Figure 11. Predicted change in $y_{n,t+1}$ (quoted in annual percentage points) as a function of weeks to maturity n in response to shift in q_t of size δ . Solid: effect in normal times (plot of $5200b_n'f\Delta$ as a function of n); dashed: effect at the zero lower bound (plot of $5200b_n^*f\Delta$).

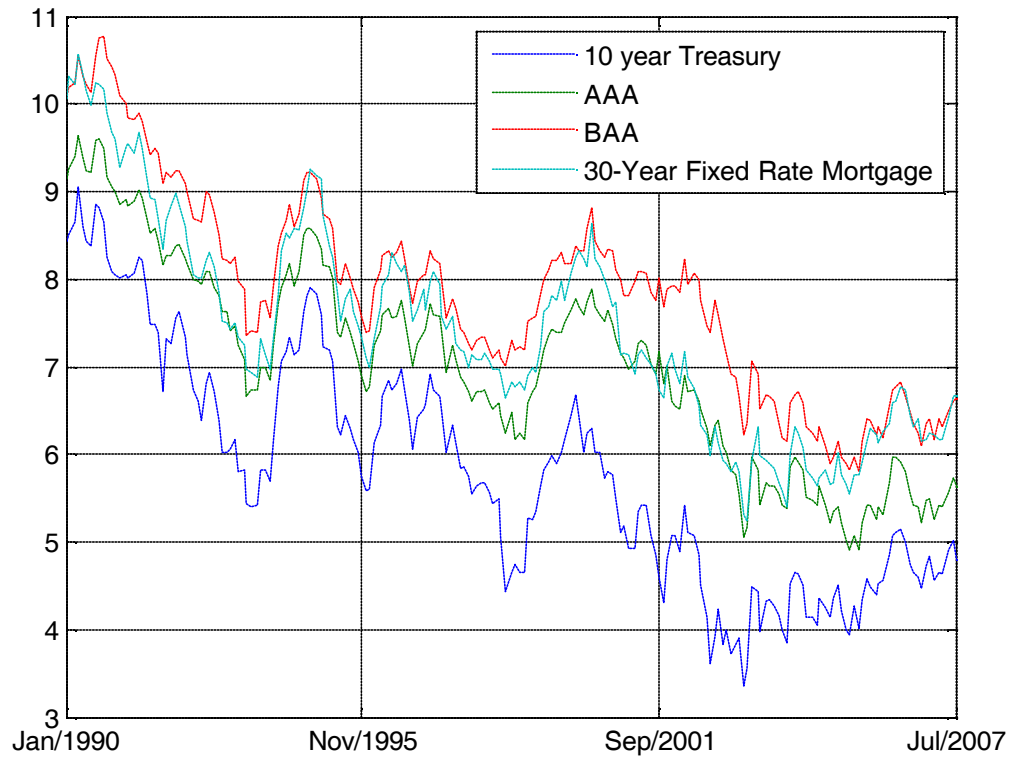


Figure 12. Assorted long-term yields, 1990:M1 to 2007:M7.

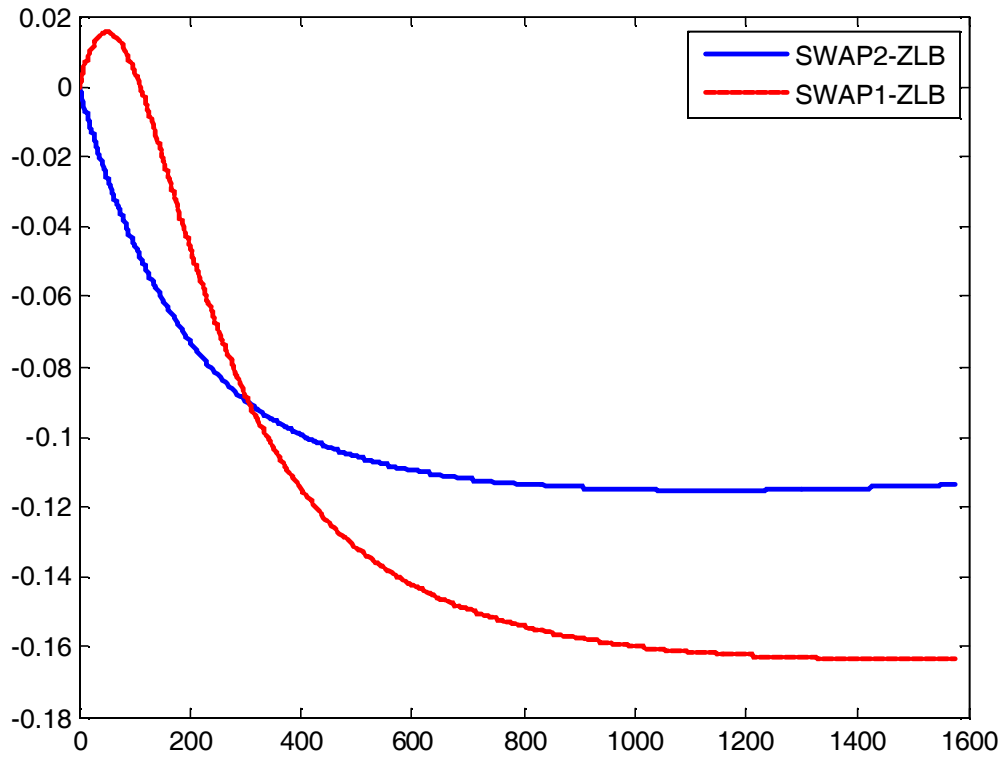


Figure 13. Effects of two different maturity swaps when implemented at the zero lower bound. Dashed curve: Fed sells all its holdings of less than 1-year maturity and retires debt at the longest end of the maturity structure (plot of $5200b_n^* f\Delta$ as a function of n). Solid curve: Fed sells all its holdings of less than 1-year maturity and retire debt evenly across 2-1/2 to 10 year maturities (plot of $5200b_n^* f\Delta_2$).

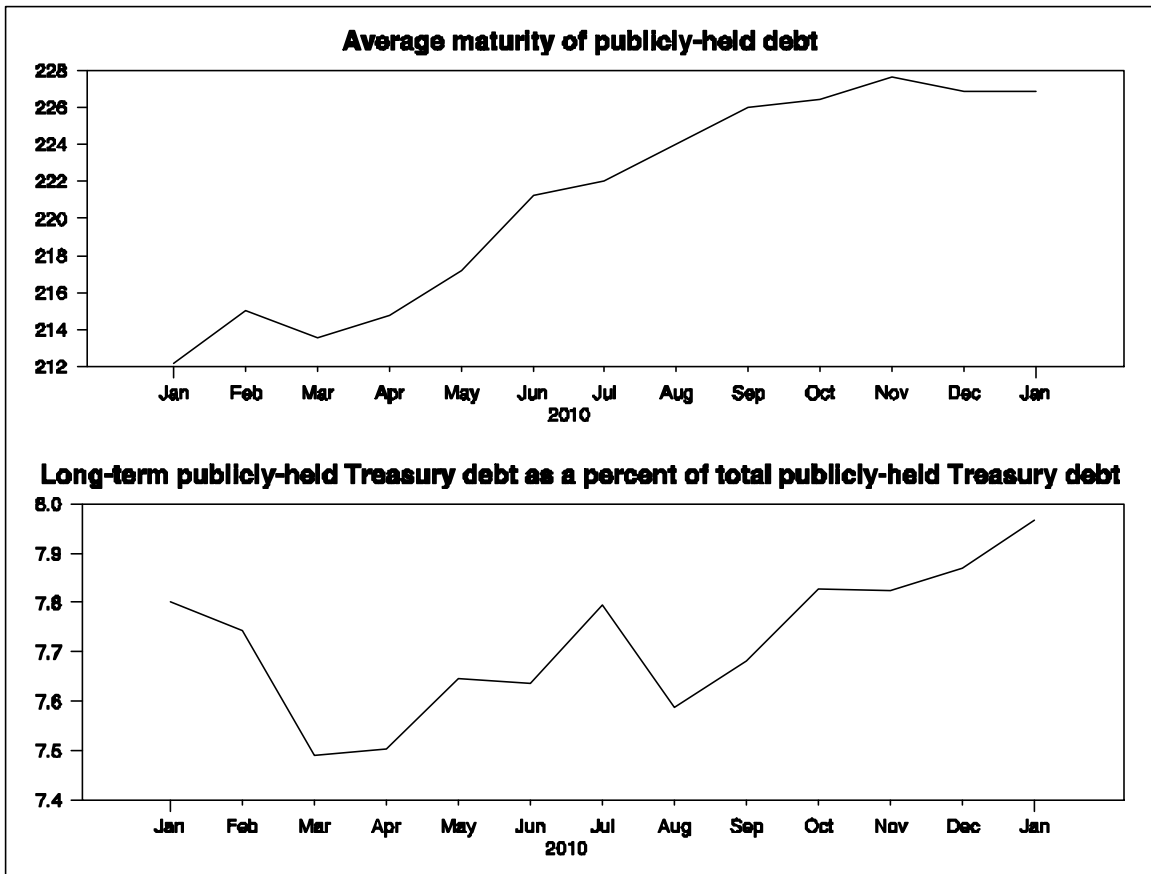


Figure 14. Top panel: average maturity of Treasury debt other than that held by the Federal Reserve (z_t^A), 2010:M1-2011:M1. Bottom panel: fraction of outstanding Treasury debt not held by the Federal Reserve that is of 10 years or longer maturity (z_t^L), 2010:M1-2011:M1.