

Expectation and Duration at the Effective Lower Bound

Thomas B. King^a

^a*Federal Reserve Bank of Chicago, Chicago, Illinois 60604*

Abstract

With risk-averse arbitrageurs and an effective lower bound on nominal rates, nonlinear interactions among short-rate expectations, bond supply, and term premia emerge in equilibrium. These interactions, which are absent from affine models, help explain the observed behavior of the yield curve near the ELB, including evidence about unconventional monetary policy. The impact of both short-rate expectations and bond supply are attenuated at the ELB. However, in simulations of the post-crisis experience in the U.S., shocks to investors' duration-risk exposures have much smaller effects than shocks to the anticipated path of short rates. The latter shocks matter, in part, because of the reduction in interest-rate volatility associated with a longer expected stay at the ELB—a novel channel of unconventional policy.

Keywords: Term structure, portfolio balance, ZLB, quantitative easing, forward guidance

JEL: E43, E52, G12

1. Introduction

Over the last decade, fixed-income markets have witnessed a combination of two extraordinary circumstances: massive changes in the quantity and structure of safe debt—including large-scale purchases of such debt by central banks—and the decline of short-term interest rates to their effective lower bound (ELB). This paper studies how these two phenomena interact in an equilibrium model of arbitrage-free bond pricing.

To see empirically that the ELB may be important for the relationship between interest rates and debt supply, the top panel of Table 1 reports regressions of long-term yields on the weighted-average maturity of outstanding Treasury debt (WAM) and the one-year Treasury yield. The data are monthly, from 1971 through 2015. Using interactive dummies, I allow the coefficients on both variables to change after the ELB was reached in December 2008, but otherwise the regressions are identical to those of Greenwood and Vayanos (2014) (whose sample ended in 2007). Indeed, column 2 replicates the main result of that paper: long-term yields were significantly positively related to the duration risk held by investors in the pre-ELB period. A one-year increase in WAM pushed the 10-year yield up by 22 basis points in this sample, and, consistent with longer-term bonds having greater duration exposure, the coefficients were higher for longer maturities.

As column 3 shows, those coefficients all fall substantially and lose statistical significance during the ELB period. Despite the relatively few observations at the ELB, the t statistics reported in column 4 show that the declines in the coefficient values are significant.¹ The relationships between short- and long-term yields also change at the ELB, as shown in columns 5 through 7. Prior to 2008, the coefficients on the one-year yield were less than 1 and were monotonically decreasing in the maturity of the dependent variable. At the ELB, the coefficients rise above 2, with the 10- and 15-year yields now being more sensitive than the 5-year yield is. Again, t tests show that the differences across the two periods are statistically significant.² The remaining panels of the table show that the shifts in both sets of coefficients are robust to using the maturity-weighted debt-to-GDP ratio in place of WAM and the two-year yield in place of the one-year.

These results suggest important changes in the behavior of the yield curve and its relationship to Treasury supply at the ELB. But what theoretical reasons do we have to expect such changes? I argue that at least three nonlinear mechanisms may be at work:

1. An increase in the quantity of longer-term bonds that investors hold raises the duration risk of their portfolios by an amount that depends directly on interest-rate volatility. Interest-rate volatility is lower when the short-rate distribution is truncated. Thus, if term premia are increasing in the amount of duration risk held by investors, the effects of bond supply will be damped at the ELB.
2. Interest-rate volatility moves together with short-rate expectations at the ELB because an increase in the length of time that the ELB is expected to bind reduces near-term uncertainty about short rates.³ If term premia depend on this uncertainty, changes in rate expectations will induce changes in term premia at the ELB.

¹The standard errors are calculated using the Newey and West (1987) procedure, with 36 lags, again following Greenwood and Vayanos (2014).

²Gilchrist et al. (2015) document similar changes in the relationship between shorter- and longer-term yields at the ELB.

³Hattori et al. (2016) show that accommodative monetary-policy announcements during the ELB period caused declines in implied interest-rate volatility across the term structure.

3. At the ELB, near-term rate expectations are constrained and are unlikely to move much in response to shocks. Consequently, changes in expectations will have relatively larger effects on medium- and long-term yields.

To formalize and quantify these possibilities, I incorporate the ELB into a structural model of bond pricing in the style of Vayanos and Vila (2009). In this type of model, the marginal investors are arbitrageurs with limited risk-bearing capacity. When they are given more long-term bonds to hold, the duration risk of their portfolios rises and they demand higher risk premia—a phenomenon sometimes known as the “duration channel” of bond supply. Many recent empirical studies on the effects of duration shocks have explicitly pointed to this framework for motivation and interpretation,⁴ and other papers have extended and applied it in various ways.⁵ Although those models have been useful for understanding the relationships between bond supply and bond yields, they have almost exclusively been developed under the assumption that the short-term interest rate follows a linear process. The three arguments just mentioned, together with the evidence in Table 1, suggest that the nonlinearity associated with the ELB may induce first-order changes in the yield curve’s behavior. The situation in which the ELB binds is particularly important to consider because central banks have typically sought to exploit the duration channel through long-term bond purchases only after they have cut their traditional policy rates close to zero.

To model the ELB, I generalize the standard affine process for the short-term interest rate in Vayanos and Vila (2009) and its followers to a “shadow rate” process, following the use of that device in the recent empirical term-structure literature, such as Kim and Singleton (2012), Krippner (2012), and Wu and Xia (2016).⁶ When the shadow rate is below the ELB, shocks to its value correspond to changes in investor beliefs about the length of time the ELB is expected to bind. Thus, they may capture both explicit forward guidance about the short-term interest rate and the “signaling channel” of asset purchases, through which expansions of the central bank’s balance sheet might be viewed as a commitment to keep rates near zero for a longer time.⁷

I show analytically that allowing for the ELB in this way endogenously gives rise to the three nonlinear phenomena discussed above. Furthermore, when parameterized to match the unconditional moments of Treasury yields since 1971, the model delivers a number of quantitatively accurate results: it replicates the basic features of the yield curve when the short rate is close to zero; it produces conditional elasticities similar to those reported in Table 1; and, when subjected to shocks that approximate the Federal Reserve’s unconventional monetary policy over the ELB period, it implies both a cumulative yield-curve impact similar to what event studies suggest and a hump shape in the forward curve that matches the pattern observed in those studies. All of these empirical successes depend crucially on the nonlinear mechanisms that emerge endogenously in the

⁴E.g., Gagnon et al. (2011), Joyce et al. (2011), Swanson (2011), Krishnamurthy and Vissing-Jorgensen (2011), D’Amico and King (2013).

⁵E.g., Hamilton and Wu (2012), Greenwood and Vayanos (2014), King (2015), Altavilla et al. (2015), Greenwood et al. (2015b), Haddad and Sraer (2015), Hayashi (2016), Malkhozov et al. (2016).

⁶Bauer and Rudebusch (2014) argue that the shadow-rate specification does a good job of capturing yield-curve dynamics near the ELB, greatly outperforming traditional affine models. Notably, however, this literature has so far been dominated by atheoretical term structure models. This paper is among the first to incorporate a shadow-rate process into a structural model of the yield curve.

⁷Woodford (2012), Bauer and Rudebusch (2014), and Bhattarai et al. (2015) argue for the importance of the signaling channel. As noted by Swanson (2017), because many announcements of asset purchases were accompanied by changes in the FOMC’s communications about future short rates, it is impossible to distinguish empirically between the effects of the signaling channel and those of forward guidance.

model. Consequently, comparable affine models, which ignore the ELB, cannot generally match these features of the data.

Having thus validated the model, I use it to address two quantitative questions. First, I ask through which channel unconventional monetary-policy shocks had their largest effects. The answer is that, given the magnitude of the observed shocks, changes in the expected path of the short rate were responsible for at least half of the cumulative downward shift in the yield curve during the ELB period. In addition, at longer maturities, about one-third of the total decline in yields is explained by the term-premium effects associated with reduced short-rate uncertainty at the ELB—a channel of unconventional policy that has been overlooked by previous literature. The bond-supply shocks account for less than 25% of the total decline in the ten-year yield, and at shorter maturities their contribution is even smaller. In other words, the model suggests that the duration channel of asset purchases was considerably less important than implicit or explicit forward guidance about the path of short-term rates.

Second, I ask the model whether the relative effectiveness of the two shocks changes in different environments. I find that bond-supply shocks are most powerful, relative to shadow-rate shocks, when the shadow rate is deeply negative and the amount of duration held by the market is high. In this situation, the efficacy of both types of shocks is attenuated because of the damping effects associated with the ELB, but the attenuation is greater for the shadow-rate shocks. A negative shadow rate and a high quantity of market duration are precisely the conditions under which most Federal Reserve asset purchases were conducted. Thus, even though those purchases appear to have had only modest effects through the duration channel, their use could have been consistent with the Fed optimizing across its policy tools in the ELB environment.

This paper is related to several others in the recent literature. As noted above, a number of studies have used variants of the Vayanos-Vila (2009) framework to analyze the effects of fluctuations in bond supply in a linear environment. Hamilton and Wu (2012) briefly considered a version in which, once the short rate reached the ELB, investors believed that it would stay there with an exogenously given probability. However, because that probability was assumed to be constant, their model did not contain a mechanism for signaling or forward guidance. In addition, away from the ELB it priced bonds as if the ELB did not exist. Thus, their model lacked the key nonlinearities and interactions that drive most of my results.

Greenwood et al. (2015b) note the hump-shaped pattern in forward rates in response to unconventional policy announcements and argue that expectations of future changes in bond supply likely account for that pattern. While my model does not rule out their type of mechanism, it implies the effects of bond-supply shocks are relatively modest and that the empirical hump in the forward curve can alternatively be explained by the non-monotonic effects of changes in short-rate expectations that arise at the ELB.

A final set of related papers are the empirical studies that have attempted to decompose the effects of unconventional policy into various channels. Krishnamurthy and Vissing-Jorgensen (2011, 2013) argue, based on event studies, that the evidence for the duration channel is weak, consistent with what my model implies. Swanson (2017) conducts event studies on unconventional policy to isolate a component reflecting short-rate expectations and a residual component that he essentially interprets as reflecting the duration channel. Although he concludes that the latter is important for long-term yields, his approach requires that factor loadings for interest rates in the ELB period were similar to those in the pre-ELB period. My model effectively allows for endogenous changes

in loadings at the ELB and suggests that those changes could be quite substantial.⁸

2. Theoretical Framework

2.1. Investor behavior and equilibrium bond prices

I begin with the same portfolio-choice problem that forms the basis of the models in Vayanos and Vila (2009) and the several theoretical papers that have followed it. Investors have access to a continuum of zero-coupon bonds with maturities 0 to T . At each point in time t , they choose to hold a market-value quantity $x_t(\tau)$ of each maturity τ . Let $P_t^{(\tau)}$ represent the time- t price of a bond with remaining maturity τ . In addition, investors have access to a risk-free security that pays the instantaneous rate r_t . Investors' time- t wealth W_t is the sum of the market-value of the bond portfolio and the risk-free asset, and it thus evolves according to

$$dW_t = \int_0^T x_t(\tau) \frac{dP_t^{(\tau)}}{P_t^{(\tau)}} d\tau + \left(W_t - \int_0^T x_t(\tau) d\tau \right) r_t dt \quad (1)$$

Investors have mean-variance preferences, and thus, taking W_t as given, they choose quantities $x_t(\tau)$ to solve the problem

$$\max_{x_t(\tau) \forall \tau} E_t[dW_t] - \frac{a}{2} \text{var}_t[dW_t] \quad (2)$$

subject to (1), where a is absolute risk aversion and E_t and var_t represent expectation and variance conditional on the time- t state.

The first-order conditions for this problem can be written as

$$E_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}} \right] = r_t dt + a \int_0^T x_t(s) \text{cov}_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}}, \frac{dP_t^{(s)}}{P_t^{(s)}} \right] ds \quad (3)$$

for all τ . Note that, under risk-neutrality ($a = 0$), all bonds have the same expected return, equal to the risk-free rate. Otherwise, the risk premium demanded for each bond is proportional to the covariance of that bond's price with the return on the whole portfolio of bonds.

The model is closed by assuming that the government exogenously supplies a time-varying quantity of bonds $z_t(\tau)$ at each maturity. A solution to the model is a set of state-contingent bond prices that clear the market. Specifically, market clearing requires

$$z_t(\tau) = x_t(\tau) \quad (4)$$

at each maturity τ and at each point in time t . Prices adjust to make (3) and (4) hold jointly in all states of the world. Since investors optimize without constraints on their portfolio weights, the equilibrium is arbitrage free.

Denote log bond prices as $p_t^{(\tau)} = \log P_t^{(\tau)}$. By Itô's Lemma,

⁸D'Amico and King (2013) and Cahill et al. (2013) present event-study evidence that asset purchases may also operate through a scarcity or "local supply" channel, whereby imperfect substitutability causes yields to fall by more for maturities where more purchases occurred. My model is silent about this type of phenomenon.

$$\mathbb{E}_t \left[dp_t^{(\tau)} \right] = \mathbb{E}_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}} \right] - \frac{1}{2} \text{var}_t \left[dp_t^{(\tau)} \right] \quad (5)$$

and

$$\text{cov}_t \left[dp_t^{(\tau)}, dp_t^{(s)} \right] = \text{cov}_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}}, \frac{dP_t^{(s)}}{P_t^{(s)}} \right] \quad (6)$$

for any two bond maturities τ and s . Expected log returns in equilibrium are thus

$$\mathbb{E}_t \left[dp_t^{(\tau)} \right] = r_t dt + a \int_0^T z_t(s) \text{cov}_t \left[dp_t^{(\tau)}, dp_t^{(s)} \right] ds - \frac{1}{2} \text{var}_t \left[dp_t^{(\tau)} \right] \quad (7)$$

Because all bonds payoff a face value of 1 with certainty at maturity—that is, they satisfy the boundary condition $p_t^{(0)} = 0$ at all t —bonds with positive maturities are given by the sum of expected future returns:

$$p_t^{(\tau)} = - \int_0^T \mathbb{E}_t \left[\frac{dp_{t+s}^{(\tau-s)}}{ds} \right] ds \quad (8)$$

The following standard relationships then determine bond yields $y_t^{(\tau)}$ and τ -period instantaneous forward rates $f_t^{(\tau)}$:

$$y_t^{(\tau)} \equiv -p_t^{(\tau)} / \tau \quad (9)$$

$$f_t^{(\tau)} \equiv - \frac{\partial p_t^{(\tau)}}{\partial \tau} \quad (10)$$

The exogenous variables in this model are the short rate r_t and the maturity-specific bond supplies $z_t(\tau)$. Throughout the paper, I assume that r_t is determined by a shadow-rate process, described in the following subsection. I consider a number of possibilities for $z_t(\tau)$. In Section 3, to illustrate the properties of the model most transparently, I simply assume that $z_t(\tau)$ is constant across maturities and over time. In Sections 4 and 5, where the focus is on quantitative estimates, I consider more-realistic stochastic processes for bond supply.

2.2. The short rate

I assume that the short rate r_t is determined by

$$r_t = \max[\hat{r}_t, b] \quad (11)$$

where the shadow rate \hat{r}_t follows the Ornstein-Uhlenbeck process

$$d\hat{r}_t = \kappa(\mu - \hat{r}_t)dt + \sigma dB_t \quad (12)$$

for mean, persistence, and variance parameters μ , κ , and σ , with B_t being a Brownian motion. I assume that $0 < \kappa < 1$. This is the same process used in the empirical shadow-rate literature

mentioned in the introduction. As noted there, that literature generally shows that the shadow-rate specification performs well in describing the reduced-form dynamics of the yield curve at the ELB. Obviously, a special case that produces an affine specification for the short rate is $b = -\infty$. This will be a useful case for comparison, because it is the specification used in the previous theoretical literature on the duration channel.⁹

Given (12), the conditional distribution of future shadow rates s periods ahead is normal, with mean $E_t[\hat{r}_{t+h}]$ and variance $\text{var}_t[\hat{r}_{t+h}]$ given by the standard affine prediction equations. In a model with no lower bound, where $\hat{r}_t = r_t$ in all states of the world, those equations also describe the conditional distribution of future short rates. Once the ELB is imposed, however, the short rate at any period in the future is distributed *truncated* normal if $\hat{r}_{t+s} > b$ and is simply equal to b otherwise. Therefore, the mean of r_{t+s} , conditional on information at time t , is given by

$$E_t[r_{t+s}] = (1 - \Phi_t^{(s)})b + \Phi_t^{(s)}E_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] \quad (13)$$

where $\Phi_t^{(s)}$ is the probability that \hat{r}_{t+s} exceeds b in period $t + s$ —that is,

$$\Phi_t^{(s)} = \Phi\left(\frac{E_t[\hat{r}_{t+s}] - b}{\sqrt{\text{var}_t[\hat{r}_{t+s}]}}\right) \quad (14)$$

where $\Phi(\cdot)$ denotes the standard-normal CDF. It follows that $E_t[r_{t+h}] > E_t[\hat{r}_{t+h}]$; thus, expectations of r_{t+s} are always higher in the presence of a finite b than they would be in the affine case. Similarly, since the variance of r_t is zero whenever $\hat{r}_t < b$, the unconditional variance of $\text{var}_t[r_{t+s}]$ is always lower in a model with an ELB.¹⁰

Since \hat{r}_{t+s} is distributed normally at all horizons s , the moments of r_{t+s} can be calculated analytically. The solid lines in Figure 1 depict the conditional moments of the forward short rate

⁹Although the ELB is imposed *a priori* here, it is trivial to extend the model to endogenize it by allowing investors to hold an elastic supply of cash (paying zero nominal return) in addition to the risk-free asset. Alternative short-rate processes that impose the ELB also exist. For example, in Monfort et al. (forthcoming), once the ELB is reached, the short rate stays there with some time-varying probability. At least qualitatively, such differences in specification are unimportant. The crucial features are that short-rate volatility is low at the ELB and that the ELB is “sticky,” in the sense that the short rate tends to stay there for some time once it reaches it. Any model that generates these properties (which are amply evident in the data) will produce results along the lines of those presented below.

¹⁰By the law of total variance, the variance of the shadow rate can be decomposed as

$$\begin{aligned} \text{var}_t[\hat{r}_{t+s}] &= \Phi_t^{(s)}\text{var}_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] + (1 - \Phi_t^{(s)})\text{var}_t[\hat{r}_{t+s}|\hat{r}_{t+s} \leq b] \\ &\quad + \Phi_t^{(s)}(E_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] - E_t[r_{t+s}])^2 + (1 - \Phi_t^{(s)})(E_t[\hat{r}_{t+s}|\hat{r}_{t+s} \leq b] - E_t[r_{t+s}])^2 \end{aligned}$$

which, from (13), simplifies to

$$\begin{aligned} \text{var}_t[\hat{r}_{t+s}] &= \Phi_t^{(s)}\text{var}_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] + (1 - \Phi_t^{(s)})\text{var}_t[\hat{r}_{t+s}|\hat{r}_{t+s} \leq b] \\ &\quad + \Phi_t^{(s)}(1 - \Phi_t^{(s)})(E_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] - E_t[\hat{r}_{t+s}|\hat{r}_{t+s} \leq b])^2 \end{aligned}$$

An analogous calculation shows

$$\begin{aligned} \text{var}_t[r_{t+s}] &= \Phi_t^{(s)}\text{var}_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] \\ &\quad + \Phi_t^{(s)}(1 - \Phi_t^{(s)})(E_t[\hat{r}_{t+s}|\hat{r}_{t+s} > b] - b)^2 \end{aligned}$$

Direct comparison of these last two equations reveals that $\text{var}_t[\hat{r}_{t+s}]$ is greater than $\text{var}_t[r_{t+s}]$.

using $\mu = 4.9\%$, $\kappa = 0.019$, $\sigma = 0.77\%$, and $b = 0.17\%$, which are values that are consistent with the data.¹¹ The moments are plotted for horizons of 2, 5, 10, and 15 years across different values of the current shadow rate \hat{r}_t , shown on the horizontal axes. Vertical lines indicate the location of the ELB. For comparison, the dashed lines show the corresponding conditional moments in the affine case, $b = -\infty$, keeping all other parameters the same.

Panel A shows how expected future short rates in the shadow-rate model asymptote to b as $\hat{r}_t \rightarrow -\infty$. Note that the derivatives of $E_t[r_{t+s}]$ with respect to \hat{r}_t decrease and eventually go to zero as the shadow rate falls below the ELB. Thus, shocks to \hat{r}_t will have generally weaker effects on the expectations component of yields when $\hat{r}_t < b$. Moreover, these shocks may have larger effects on the expectations component of medium or long-term yields than on shorter-term yields. This contrasts to an environment far above the ELB, where the effects of shocks to \hat{r}_t are always largest at the short end of the curve.

Similarly, in Panel B, as \hat{r}_t falls below b , the conditional variance of the future short rate drops notably. The reason for this is intuitive—when the shadow rate is far below the ELB, the *actual* short rate will almost certainly be *equal* to the ELB for a long time and therefore will display little variation. All else equal, the lower short-rate volatility near the ELB will mean that the volatilities of all yields are lower near the ELB. Thus the covariance terms that represent the multipliers on z_t in equation (7) will generally be smaller. Heuristically, this implies that both the levels of risk premia and their sensitivity with respect to bond holdings will be smaller at the ELB than they are away from the ELB (or in an affine model). The following section explores these assertions more rigorously.

3. Consequences of the ELB for bond pricing

In this section I demonstrate how the presence of the ELB affects the equilibrium determination of bond prices. To isolate the key mechanisms, I consider a simplified version of the model in which bond supply is flat across maturities and constant over time. That is, $z_t(\tau) = \zeta$, for some constant $\zeta > 0$. Of course, for quantitative purposes, it will ultimately be necessary to relax this assumption, and the subsequent sections of the paper will consider stochastic-supply versions of the model in specific numerical simulations. For the purposes of illustrating the general properties of the model, however, holding bond supply fixed greatly reduces analytical complexity without meaningfully altering the conclusions.

3.1. Bond returns

In general, bond prices are a nonlinear function of the shadow rate. Define the time and maturity-specific coefficients $A_t^{(\tau)} \equiv \partial p_t^{(\tau)} / \partial \hat{r}_t$. Since \hat{r}_t is the only source of variation in the model, the covariance between the returns on two bonds of arbitrary maturities τ and s can be written as

$$\text{cov}_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}}, \frac{dP_t^{(s)}}{P_t^{(s)}} \right] = A_t^{(\tau)} A_t^{(s)} \sigma^2 \quad (15)$$

With $x(\tau) = \zeta$ for all τ in equilibrium, equation (3) gives the expected excess return (i.e., risk premium) on a τ -maturity bond:

¹¹The details of the parameterization are discussed in Section 4.2.

$$\mathbb{E}_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}} \right] - r_t dt = a\zeta\sigma^2 A_t^{(\tau)} \int_0^T A_t^{(s)} ds \quad (16)$$

Solving the model is tantamount to solving for the coefficients $A_t^{(\tau)}$. In the special case where the ELB never binds ($b = -\infty$), one can show that the solution is

$$A_t^{(\tau)}|_{(b=-\infty)} = \int_0^\tau e^{-\kappa s} ds = \frac{1 - e^{-\kappa\tau}}{\kappa} \quad (17)$$

Indeed, this is precisely the affine model studied by Vayanos and Vila (2009) with constant bond supply. Note that $A_t^{(\tau)}|_{(b=-\infty)}$ is non-negative and increasing in τ , implying that risk premia are strictly positive and larger for longer-maturity bonds. Because $A_t^{(\tau)}|_{(b=-\infty)}$ is constant for any τ , at any given maturity the model produces a return volatility and a risk premium that do not vary with \hat{r}_t . Finally, since (17) does not involve ζ , the effect of a hypothetical change in bond supply on the maturity- τ expected return is given by

$$\frac{\partial \mathbb{E}_t[dP_t^{(\tau)}/P_t^{(\tau)}]}{\partial \zeta} = a\sigma^2 A_t^{(\tau)} \int_0^T A_t^{(s)} ds \quad (18)$$

Thus, expected returns at any given maturity have a (positive) sensitivity to the level of bond supply that is independent of the value of the shadow rate.

Once we allow the ELB to bind occasionally, the model no longer admits an analytical solution. However, Appendix A shows that, to a first-order approximation, $A_t^{(\tau)}$ in this more general case is given by

$$A_t^{(\tau)} \approx \int_0^\tau e^{-\kappa s} \Phi_t^{(s)} ds \quad (19)$$

where, as in the previous section, $\Phi_t^{(s)} = \Pr_t[\hat{r}_{t+s} > b]$.¹² In words, $A_t^{(\tau)}$ represents the accumulated stream of probabilities that the short rate will be unconstrained by the ELB in each of the next τ periods, discounted at rate κ . The more likely \hat{r}_t is to spend time below b between periods t and $t + \tau$, the lower $A_t^{(\tau)}$ will be. The Vayanos-Vila model in (17) is a limiting case that holds when the ELB never binds so that $\Phi_t^{(s)} = 1$ for all t and s .

The key result is that, outside of the affine case, $A_t^{(\tau)}$ is strictly increasing in \hat{r}_t . Specifically,

$$\frac{\partial A_t^{(\tau)}}{\partial \hat{r}_t} \approx \int_0^\tau e^{-2\kappa s} \varphi_t^{(s)} ds > 0 \quad (20)$$

where $\varphi_t^{(s)}$ is the PDF associated with $\Phi_t^{(s)}$. Since this derivative is positive for all bonds when $b > -\infty$, the covariance between any two bonds' returns, given in equation (15), is also increasing in \hat{r}_t . In particular, bond-return volatilities (and covariances) become damped as the short rate

¹²Equation (19) holds exactly if either $a = 0$ or $b = -\infty$. The absence of a closed-form expression for bond prices is a general property of Gaussian no-arbitrage shadow-rate models, even when no structural equilibrium conditions are imposed. See Kim and Singleton (2012), for example.

approaches the ELB from above. Effectively, the reduced variance of r_t that was evident in Figure 1 gets passed through to the variances of longer-term bonds.

This state-dependence of the return covariances has implications for risk premia through equation (16). Risk premia are uniformly higher when the shadow rate is higher. Even when r_t is constrained by the ELB, shocks to the shadow rate—such as those induced by forward guidance about monetary policy—have effects on risk premia through this channel. An increase in the length of time the ELB is expected to bind, for example, is associated with a reduction in $A_t^{(\tau)}$ and therefore leads to lower expected excess returns on bonds, in addition to a lower expected path of the short rate. Furthermore, the derivative in (18) is directly proportional to the $A_t^{(\tau)}$ terms and is therefore also increasing in the shadow rate. Thus, the effects of changes in bond supply will be lower when r_t is at or near the ELB than they are when it is far above the ELB.¹³

These results can also be viewed in terms of the price and quantity of risk. In particular, one can rearrange equation (16) to express expected excess returns as

$$\mathbb{E}_t \left[\frac{dP_t^{(\tau)}}{P_t^{(\tau)}} \right] - r_t = A_t^{(\tau)} \lambda_t \quad (21)$$

where the market price of risk λ_t is the same for all bonds in period t and is given by

$$\lambda_t = a\zeta\sigma^2 \int_0^T A_t^{(s)} ds \quad (22)$$

Risk prices depend positively on all of the $A_t^{(\tau)}$ terms, which means that they are increasing in the shadow rate—that is, risk prices are lower at the ELB.¹⁴ In contrast, in the affine version of the model risk prices are constant. Furthermore, (22) makes it clear that bond supply affects returns by determining the price of risk. In the affine model, the sensitivity of the risk price to bond supply $\partial\lambda_t/\partial\zeta$ is constant. In contrast, with $b > -\infty$, this sensitivity is strictly increasing in \hat{r}_t so that risk prices are less responsive to bond supply at the ELB.

3.2. Bond yields

The discussion to this point has been cast entirely in terms of expected bond returns, but in the remainder of the paper the primary interest will be in bond *yields*. Although they are somewhat more cumbersome to work with, all of the qualitative results just discussed for expected returns also hold for yields. In particular, by substituting equation (16) into (5), (8), and (9), it is straightforward to show the following:

- The sensitivity of yields to bond supply ($\partial y_t^{(\tau)}/\partial\zeta$) is strictly increasing in the shadow rate.
- Term premia, $y_t^{(\tau)} - \mathbb{E}_t \int_0^\tau r_{t+s} ds$, are strictly increasing in the shadow rate.

¹³Equation (18) still holds as an approximation to the effect of bond supply in the presence of the ELB, because (19) does not involve ζ .

¹⁴Indeed, risk prices are nonlinear in \hat{r}_t . This result presents a challenge to empirical shadow-rate term-structure models, which typically assume that risk prices are affine in the states (e.g., Kim and Singleton, 2010; Wu and Xia, 2016). The model here instead suggests that risk prices endogenously decrease in a nonlinear way around the ELB. I thank Don Kim for pointing this out to me.

- At the ELB, shocks to the shadow rate have their largest impact on intermediate-maturity yields.

The first two of these propositions follow immediately from the discussion of bond returns above. The third proposition arises from the following argument. Because the state variable in the model is stationary, expectations for returns in the far future must be invariant to today's shocks. Because yields are averages of expected returns, this implies $\partial y_t^{(\tau)} / \partial \hat{r}_t \rightarrow 0$ as $\tau \rightarrow \infty$. But, at the ELB, short-term yields are constrained at b . Therefore, $\partial y_t^{(0)} / \partial \hat{r}_t = 0$ when $\hat{r}_t < b$. Since yields do not move at the very short or very long end, they must move the most in the middle of the curve.

Figures 2 and 3 demonstrate these propositions quantitatively, solving the model numerically. The numerical solution, described in Appendix B, relies on an iterative projection method and is more precise than the first-order approximation described above. For the purposes of illustration, I use the same short-rate parameters as in Figure 1 and set risk aversion to $a = 0.15$, which is the value used in the two-factor version of the model discussed in the next section.

Figure 2 shows comparative statics over the supply parameter ζ at two different levels of the shadow rate. In the left-hand graph, the shadow rate is at $\hat{r}_t = 5.2\%$, the unconditional mean of the short rate in the data. In the right-hand graph, it is at $\hat{r}_t = -2.7\%$, which is the average value obtained by the shadow rate during the ELB period, according to Krippner's (2012) estimates. In both graphs, the black line depicts the yield curve when the model is solved using $\zeta = 0.42$, a value which allows the model to match the average level of the ten-year yield in the data. The red line shows the yield curve using $\zeta = 0.21$. Thus, loosely speaking, the comparison shows what happens to the yield curve when we remove half of the bond supply from the market, conditional on different values of the shadow rate. Since expected future short rates are unaffected by the value of ζ , the entirety of the difference between the black and red lines reflects a difference in term premia.

The reduction in bond supply has a notable effect on yields in both panels. However, when the short rate is above the ELB, the effect is substantially larger. For example, the ten-year yield falls by 95 basis points when $\hat{r}_t = 5.2\%$ but only by 53 basis points when $\hat{r}_t = -2.7\%$ in response to the same decrease in supply. This illustrates the attenuation of supply effects when the ELB is binding.

Figure 3 shows the effect of a one-standard-deviation negative shock to the shadow rate, starting from the same two illustrative values of \hat{r}_t . The top panel shows the responses of the spot-rate yield curve, while the bottom panel shows the same information in terms of forward rates, where the patterns are somewhat easier to see. In each graph, the pink region depicts the change, on impact, resulting from the shock through the expectations component, while the blue region depicts the change in the term premium. Far above the ELB, the shadow rate shock has a monotonic effect on yields that arises almost entirely through the expectations component. This is similar to the outcome in the affine model, where the term premium is unaffected by shocks to the short rate. (Indeed, it asymptotes to that case as the starting value for \hat{r}_t is moved farther above the ELB.) Below the ELB, the situation is quite different. The overall effect of the shock is smaller at all maturities; the reaction across forward rates is hump-shaped rather than monotonic; and the response is due to both a decline in the expectations component and a decline in the term premium. The reduction in the term premium in response to the shock accounts for about a third of the overall yield decline at longer maturities in this example (e.g., 10 of the 28 basis points on the ten-year yield). Its largest effect on forward rates is at the 8-year maturity, contributing to the hump shape of the response. Thus, these results illustrate both the term premium effects of shadow-rate shocks at the ELB and the relatively large reaction of the middle of the forward curve to such shocks in these circumstances.

4. Quantitative evaluation of the model with stochastic bond supply

I now turn to the quantitative implications of the model when there are shocks to the supply of bonds available to investors. I first show that, when suitably parameterized, the model matches a number of key features of the data, particularly at the ELB. In Section 5, I then use the model to simulate the effects of the Federal Reserve’s unconventional monetary policies.

4.1. Bond supply process

To lend further realism to the model, I relax the assumption that bond supply is flat across maturities and constant over time. Since bond supply is continuous across maturities, the object $z_t(\tau)$ is an infinite-dimensional vector. Previous literature specifies it as an affine function of a finite state vector β_t :

$$z_t(\tau) = \zeta(\tau) + \theta(\tau)\beta_t \quad (23)$$

where $\zeta(\tau)$ and $\theta(\tau)$ are maturity-specific intercepts and coefficients. Meanwhile, β_t is assumed to follow the Ornstein-Uhlenbeck process

$$d\beta_t = -\kappa_\beta \beta_t dt + \sigma_\beta dB_{\beta,t} \quad (24)$$

where κ_β and σ_β are parameters governing the persistence and variance of β_t , and $B_{\beta,t}$ is a Brownian motion with innovations independent of those in B_t . I take the dimension of β_t to be 1.

I further assume that the intercept in (23) is constant across maturities: $\zeta(\tau) = \zeta$. This involves only a small loss of generality, since $\zeta(\tau)$ is integrated out in equation (7) and is thus only a level shifter. Similarly, the individual factor loadings $\theta(\tau)$ do not matter for yields; only the weighted sum $\int_0^T \theta_t(s) \text{cov}_t \left[dp_t^{(\tau)}, dp_t^{(s)} \right] ds$ does. This suggests that the exact specification of the function $\theta(\tau)$ is not of first-order importance, so long as it can generate realistic behavior for overall portfolio duration.¹⁵ For simplicity, I follow Greenwood et al. (2015b) by assuming that this function is linear in τ :

$$\theta(\tau) = \left(1 - \frac{2\tau}{T} \right) \quad (25)$$

This specification implies that the bond distribution behaves with a see-saw motion across maturities. Positive supply shocks reduce the amount of long-term bonds and increase the amount of short-term bonds in equal measure, with the fulcrum at $\tau = T/2$.

Two helpful summary measures of bond supply that are frequently used in the literature have direct counterparts in the model. The first measure is the weighted-average maturity (WAM) of the outstanding debt, which is given by

$$WAM_t = v \frac{\int_0^T \tau z_t(\tau) d\tau}{\int_0^T z_t(\tau) d\tau} \quad (26)$$

where v is the length of one period, expressed in years. The second measure is the amount of “ten-year-equivalent” bonds outstanding. This variable is defined as the dollar value of ten-year bonds that would produce the same duration-weighted value that the actual portfolio of outstanding

¹⁵Malkhozov et al. (2016) make a similar point.

bonds has. (Thus, for example, a portfolio of 5-year bonds with a value of \$100 is worth \$50 in terms of ten-year equivalents.) Mathematically, the amount of ten-year equivalents is defined as

$$10YE_t = \frac{v}{10} \int_0^T \tau z_t(\tau) d\tau \quad (27)$$

The integrals in both of these equations can be evaluated analytically, given the assumed process for $z_t(\tau)$, providing convenient ways of translating real-world changes in the outstanding bond distribution into the bond-supply shocks of the model.

4.2. Parameters and solution

There are nine parameters in the model, which I set to match empirical moments of Treasury supply and the yield curve. Specifically, I use the Gurkaynak et al. (2007) zero-coupon yields available on the Federal Reserve Board's website and the Treasury security data available in CRSP.¹⁶ I start the sample in August 1971 because at that time 10-year yields become available. The sample ends in December 2015. The specific moments that I match are discussed below. The parameters are summarized in the top line of Table 2.

A period is normalized to one calendar quarter (i.e., $v = 1/4$), and I take T , the maximum-maturity bond available to investors, to be 60 quarters. This matches the longest maturity bond that was continuously available over the period under consideration. It also happens to be close to the typical duration of a 30-year coupon bond, which is the longest bond issued by the Treasury at any point during this period. Alternative values for T make little difference, however.

I calibrate the autoregressive coefficient on the supply factor κ_β to match the persistence of the weighted-average maturity of outstanding Treasury debt (the same series used in the regressions in Table 1). In the data, this variable is calculated as the value-weighted timing of all cash flows on all Treasury instruments held by the public. In the model, the WAM of the debt held by investors, expressed in years, is given by equation (26):

$$WAM_t = v \frac{\int_0^T \tau [\zeta + (1 - \frac{2\tau}{T}) \beta_t] d\tau}{\int_0^T \zeta + (1 - \frac{2\tau}{T}) \beta_t d\tau} = vT(\frac{1}{2} - \frac{1}{6\zeta} \beta_t) \quad (28)$$

Since WAM_t is linear in β_t , it has the same persistence. Thus, I match the persistence of WAM in the data, using the four-quarter autocorrelation (0.92) to abstract from seasonal patterns in Treasury issuance. This gives $\kappa_\beta = 0.021$.¹⁷ The parameter σ_β determines the scale of the bond-supply factor. Since β_t is unitless, this parameter has no economic content and indeed is not separately identified. Without loss of generality, therefore, I set it such that the unconditional variance of β_t is normalized to 1.

I determine the remaining parameters by matching the long-run empirical features of the yield curve. In general, all of the moments of yields are affected by all of the model parameters, and I

¹⁶Center for Research in Security Prices, Booth School of Business, The University of Chicago. Used with permission. All rights reserved. crsp.uchicago.edu.

¹⁷Of course, equating investors' bond holdings in the model with Treasury debt in the data might be taking the model too literally given that real investors may also have duration exposure through other instruments. However, other calibrating κ_β to other autocorrelations in the 0.8 to 1 range, holding the rest of the parameters constant, produces similar outcomes to those reported below.

therefore search for vectors of parameters that jointly match the data using simulated method of moments. Specifically, I (1) solve the model numerically for each candidate parameter vector, (2) simulate 1 million periods of data using that solution, and (3) compute the model-implied moments that result from those simulations. I iterate this procedure until the moments all are matched to within three significant digits. There are six remaining parameters and I match six moments in the data, so the parameter values are just-identified.

The specific moments I match are the unconditional mean and standard deviation of the three-month Treasury yield (5.2% and 3.6% respectively), the unconditional mean and standard deviation of the ten-year yield (6.7% and 2.9%), and the correlation between the three-month and ten-year Treasury yield (0.91). In addition, I match the average value of the three-month Treasury yield at the ELB. Specifically, between December 2008 and December 2015, the three-month yield averaged 0.22%, with a maximum value of 0.68; I thus require parameters that achieve a mean $y_t^{(1)}$ of 0.0022 conditional on $y_t^{(1)} < 0.0068$.¹⁸ The values for κ , μ , σ , a , b , and ζ that produce these same moments in the simulations, given the values of the other parameters, are those reported in the top row of Table 2.

For comparison, I consider two alternative models in which $b = -\infty$, i.e., models with an affine process for the short rate. In the first such model, shown in row 2 of Table 2, I set all parameters other than b equal to the same values as in the shadow-rate model in order to isolate the effects of imposing the ELB. In the second affine model, shown in row 3, I recalibrate the parameters to match the same set of unconditional yield-curve moments that the shadow-rate model matches. The parameters turn out to be fairly similar to those in the baseline model, with the primary differences being that $\mu_{\hat{r}}$ is a bit higher and $\sigma_{\hat{r}}$ is a bit lower due to the truncation effects noted earlier.

I solve the model globally by discretizing the state space and iteratively (a) calculating state-contingent bond prices in equation (8) given conditional expectations, and (b) calculating conditional expectations given state-contingent prices using the transition densities implied by equations (12) and (24). Cubic interpolation between the discretized nodes is used for situations, such as model simulation, in which state values are required to be continuous. The details of the solution method are described in Appendix B.

4.3. Model fit

Table 3 summarizes the properties of bond yields produced by the calibrations in the shadow-rate and affine models and compares these results to the data. The model-implied moments are calculated by drawing 1,000,000 times from the distributions of e_t^r and e_t^β and simulating the resulting paths of the state variables \hat{r}_t and β_t . To illustrate the importance of the ELB, I report the results conditional on the one-quarter yield being both below and above the value 0.68%. Again, the reason for choosing this threshold is that it was the maximum attained by the three-month Treasury yield in the data during the time that the Federal Reserve kept its policy rate in the 0 – 25 basis point range.

The shadow-rate model matches the data quite well when the short rate is at its lower bound, coming within a few basis points of the means and standard deviations of all but the longest yields.

¹⁸Note that the three-month yield used in these calculations is the fitted value of the Gurkaynak et al. curves, which are based on Treasury coupon-security data. It is not a Treasury bill rate. It thus avoids any premium associated with very liquid, “money-like” assets.

In contrast, the affine model with the same calibration predicts a one-quarter yield that averages -1.3% (negative Treasury rates never actually appear in the data) and an average yield-curve slope that is dramatically steeper than what was observed. It also predicts slope volatilities that are somewhat farther from the data than those given by the shadow-rate model. The affine model that is recalibrated to match the unconditional yield curve moments does slightly better at the ELB, but it still significantly underperforms the shadow-rate model. It also underestimates the frequency with which the ELB binds by about 40%. The shadow-rate model achieves its successes near the ELB without sacrificing performance relative to the affine models in other regions of the state space. As the bottom panel shows, all three models differ from each other by only a few basis points for all of the reported statistics when the one-quarter yield is greater than 0.68%.

4.4. Yield sensitivities to state variables

I now consider how bond supply and the shadow rate affect the yield curve in different regions of the state space. To facilitate this discussion, define the state-dependent factor loadings $a_{\hat{r},t}^{(\tau)} \equiv \partial y_t^{(\tau)} / \partial \hat{r}_t$ and $a_{\beta,t}^{(\tau)} \equiv \partial y_t^{(\tau)} / \partial \beta_t$.¹⁹ The solid lines in Figure 4 depict these loadings in the shadow-rate model across a range of values for \hat{r}_t , holding β_t fixed at its mean value of zero. The dashed lines depict the corresponding loadings in the affine model under the baseline calibration (line 2 of Table 2). Since yields are affine functions of the states in the affine model, the factor loadings are constant and the dashed lines are always flat. The factor loadings in the shadow-rate model asymptote to those of the affine model as \hat{r}_t rises farther above the ELB.

For reasons discussed in Section 3, $a_{\hat{r},t}^{(\tau)}$ is monotonically increasing in \hat{r}_t , as shown in panel A of Figure 4. Furthermore, for \hat{r} low enough, $a_{\hat{r},t}^{(\tau)} < a_{\hat{r},t}^{(s)}$ when $\tau < s$. Consequently, longer-term yields respond to shadow-rate shocks by more than shorter-term yields do. This pattern is the opposite of what we observe when $\hat{r}_t > b$, and it is the opposite of what the affine model predicts. Meanwhile, as shown in panel B, in the shadow-rate model, $a_{\beta,t}^{(\tau)}$ is monotonically decreasing (i.e., becoming more negative) in \hat{r}_t . (Recall that increases in β_t reduce the duration exposure of investors and therefore have negative effects on term premia.) This result replicates in the stochastic-supply version of the model the damping effects of the ELB on bond the effects of bond supply that were demonstrated in Section 3. As discussed there, it occurs because $\text{var}_t[r_t]$ falls as \hat{r}_t moves below the ELB, causing the covariance terms in (7) to become smaller. Consequently, at the ELB—and particularly when \hat{r}_t is deeply negative—bond-supply shocks have smaller effects on yields than they do in the affine model.

Figure 6 shows the factor loadings plotted across a range of values for β_t , holding \hat{r}_t fixed at either the unconditional mean of the short rate (Panel A) or at a value of -2.7% (Panel B). As before, I choose the latter value for illustration of the ELB environment because it is the average of the shadow rate estimated by Krippner (2012) during the ELB period in the U.S..²⁰ From this perspective, the differences between the shadow-rate and the affine model are evident even when the short rate is at its unconditional mean, and they are of first-order importance when the shadow rate is negative. In that region, two particularly noteworthy results stand out. First, $a_{\hat{r},t}^{(\tau)}$ is not

¹⁹Note that $a_{\hat{r},t}^{(\tau)} = -A_t^{(\tau)} / \tau$ in the one-factor model of the previous section. That is, it is the loading in yield space, rather than return space.

²⁰To interpret the meaning of $\hat{r}_t = -.027$, simulations starting from this value produce a modal time of 6 quarters until the shadow rate moves above the ELB. This is roughly consistent with survey evidence on market participants' expectations and other evidence collected during much of the ELB period (see Femia et al. (2013)).

only strictly lower than it is when \hat{r}_t is positive, it is also decreasing in β_t . The reason is that, when β_t is positive, investors have relatively little exposure to long-term bonds. Consequently, when the shadow rate rises, the resulting increase in short-rate risk has a relatively small effect on term premia. When β_t is negative, in contrast, investors' bond exposures are greater, and increases in the shadow rate have a larger impact on term premia through their effects on short-rate volatility. Second, $A_{\beta,t}^{(\tau)}$ is increasing (becoming less negative) in β_t . Intuitively, higher levels of β_t reduce exposure to long-term bonds, making long-term yields less sensitive to the changes in short-rate risk induced by the shadow-rate. Consequently, positive shocks to β_t reduce the volatility of yields, making *further* shocks to β_t less potent. This result implies, for example, that the marginal effects of asset purchases decline as the central bank does more of them. It will also be important for analyzing the relative effectiveness of alternative policies in different environments in Section 5.

The state-dependent factor loadings explain the empirical patterns that were illustrated in Table 1 in the introduction. Recall that the regressions reported there showed that Treasury supply had smaller effects during the ELB period than in the pre-ELB sample. This is exactly the result predicted by panel B of Figure 4. Table 1 also showed that regression coefficients of long-term on short-term yields became larger at the ELB, in particular rising from less than 1 to greater than 1. That the model reproduces this result can be seen by examining the factor loadings in the panel 2.A, and in particular how they cross in the sub-ELB region.

Table 4 makes these patterns clearer by computing the linearized relationships between long-term yields, shorter-term yields, and WAM_t in the model. Given equation (28), the response of the τ -period yield to a change in WAM_t is $-6\zeta a_{\beta,t}^{(\tau)}/vT$. Panel A reports these responses at various values of the shadow rate, for $\tau = 5, 10$, and 15 years, in the first group of columns. To make the comparison to Table 1 clearer, the second and third groups of columns report the sensitivity to WAM_t holding fixed the one-year yield or the two-year yield.²¹ The sensitivities are similar in magnitude to the estimates in Table 1. They show a clear decline at the ELB, just as the regression results do.²² Similarly, panel B shows state-contingent linearized coefficients of the τ -period yield on the one- and two-year yields, holding the supply factor fixed ($a_{\hat{r},t}^{(\tau)}/a_{\hat{r},t}^{(4)}$ and $a_{\hat{r},t}^{(\tau)}/a_{\hat{r},t}^{(8)}$). As in Table 1, the coefficients rise from less than 1 to greater than 1 at the ELB and switch from being most-sensitive to least-sensitive at the five-year maturity. These results further support the ability of the model to explain the empirical behavior of the yield curve and its relationship to bond supply in the ELB environment.

5. Assessing unconventional policy

5.1. Simulating policy paths

I now use the model to study the effects of unconventional monetary policy. The Federal Reserve implemented two main types of such policy: asset purchases (also known as “quantitative easing”) and forward guidance about the future course of the short-term interest rate. Jointly, these policies can be mapped into the shadow-rate and bond-supply shocks of the model. However, as a number of authors have noted, QE may have worked in part through a “signaling channel,”

²¹These coefficients are given by $-6\zeta(a_{\beta,t}^{(\tau)} - a_{\beta,t}^{(s)}/a_{\hat{r},t}^{(s)})/vT$ for $s = 4$ or 8 quarters.

²²Since there is no concept of GDP in the model, it is not possible to mimic the specifications in Table 1 that use the maturity-weighted debt-to-GDP ratio. However, if the model results are computed replacing WAM_t with the model-implied measure of ten-year-equivalent bonds the same patterns hold.

serving as a commitment by the Fed to keep the short rate at the ELB for a longer time. (E.g., Woodford (2012); Bauer and Rudebusch (2014); Bhattarai et al. (2015).) If so, then such policies involve shocks to both bond supply and the shadow rate. For this reason, I do not attempt to distinguish the effects of forward guidance and QE *per se* but rather model the joint effects of changes in the anticipated short rate and bond supply.

The strategy is to feed the model a set of shocks that approximate those associated with unconventional policy during the ELB period and calculate the yield-curve effects of those shocks. To conduct this exercise, one must translate the actions taken by the Federal Reserve into shocks that can be input into the model. Cumulatively, we know fairly precisely how large these shocks were. r_t remained at the ELB for seven years, and so the shadow-rate shocks in the simulation must keep r_t at b for exactly 28 periods. Over the same time, Greenwood et al. (2015a) report that the Fed removed approximately \$2.7 trillion of ten-year-equivalent bonds from the market, including Treasuries, agency debt, and MBS. This was approximately 21% of the total 10-year equivalents outstanding in these markets as of December 2015, which from equation (27) is sufficient to pin down the cumulative size of the shocks to bond supply.²³

While it is tempting to interpret the bond-supply shocks associated with QE events simply as realizations of e_t^β , Federal Reserve asset purchases likely differed in important ways from the other types of bond-supply fluctuations that dominate the long span of data. In particular, in the baseline model above the parameter κ_β was calibrated to a value of 0.021, implying a half-life of 8.5 years, to match the persistence of Treasury debt since 1971. But Fed asset purchases were almost certainly interpreted as less persistent than that. Carpenter et al. (2015) inferred from surveys of market participants, conducted while the QE programs were taking place, that the size of the Fed’s balance sheet was expected to normalize by August 2020. By that reckoning, the expansion of the Fed’s balance sheet, which occurred between December 2008 and December 2014, had a perceived half-life of less than 4.5 years on average, substantially less than that of bond-supply shocks under the baseline calibration.²⁴

To account for these differences, I extend equation (23) to allow for an additional supply factor Q_t representing bond-supply shocks due to changes in the Federal Reserve’s balance sheet:

$$z_t(\tau) = \zeta + \theta(\tau)(\beta_t + Q_t) \quad (29)$$

where

$$dQ_t = -\kappa_Q Q_t dt + \sigma_Q dB_{Q,t} \quad (30)$$

for “Fed balance-sheet” shocks $dB_{Q,t}$. I set $\kappa_Q = 0.04$, giving Q_t shocks a half life of 4.25 years, and I continue to specify $\theta(\tau)$ as in (25). When calculating bond yields in this model, I set σ_Q to zero, so that the perceived risk associated with total bond supply is the same as in the baseline

²³In December 2015, the CRSP Treasury data show ten-year-equivalent Treasury bonds of \$9.5 trillion, while SIFMA data show \$7.4 trillion of agency-backed MBS and CMOs and \$1.3 trillion of long-term agency debt outstanding (<http://www.sifma.org/research/statistics.aspx>). Hanson (2014) shows that the average duration of a 30-year MBS is about 3.5 years, and I assume that the duration of long-term agency debt is 5 years. Under these approximations, ten-year equivalents outstanding totaled \$12.7 trillion.

²⁴Other evidence on the persistence of QE is mixed, but it does not suggest mean reversion as low as 0.021. Wright (2012) estimates a half-life of less than a year for the effects of unconventional monetary-policy shocks on yields in a VAR. Altavilla and Giannone (2016) show that markets expected most of the effects of unconventional policy to persist for at least a year, but the survey data they use do not extend beyond that horizon. Similarly, Swanson (2017) finds that the persistence of most QE shocks was large, but he limits the estimation horizon to 180 business days.

model.²⁵ All other parameters continue to take the values shown in the top row of Table 2. Since Q_t does not add additional risk to the model, the conditional moments of yields and the factor loadings for \hat{r}_t and β_t also remain the same as above. I note that, although the addition of the balance-sheet factor adds realism to the model, the results presented below are largely unchanged if asset purchases are simply treated as ordinary shocks to β_t .

The nonlinearities induced by the ELB mean that both starting values and the trajectory of the shocks matter. I use starting values based on the configuration of the yield curve on the eve of unconventional policy. Letting $t = 0$ denote the period immediately before unconventional policies were enacted, I set $\hat{r}_0 = 0.0017$, just at the ELB. I set $\beta_0 = -0.42$, which produces a ten-year yield-curve slope of 3.0%, the observed slope as of the FOMC meeting prior reaching the ELB. I initialize Q_0 to zero, since QE did not exist prior to the ELB.²⁶

Since we cannot directly observe the trajectories of the shadow rate and the Fed balance-sheet factor in the data, I simulate a range of possible trajectories, with each trajectory being consistent with the observed outcomes of (1) a short rate that stays at zero for exactly 28 periods and (2) a cumulative net reduction in 10-year-equivalent bonds of 21%.²⁷ The details of these simulations are discussed in Appendix C. Each simulation i consists of a set of 28 shocks to both the shadow rate and the Fed's balance sheet $\{(e_{i,1}^{\hat{r}}, e_{i,1}^Q), \dots, (e_{i,28}^{\hat{r}}, e_{i,28}^Q)\}$, which accumulate into the state trajectories $\{(\hat{r}_{i,0}, Q_{i,0}), \dots, (\hat{r}_{i,28}, Q_{i,28})\}$ via equations (12) and (30). (By construction, $e_{i,t}^{\hat{r}} = B_{\hat{r},t} - E_{t-1}[B_{\hat{r},t}]$ and similarly for $e_{i,t}^Q$.) The initial values $(\hat{r}_{i,0}, Q_{i,0})$, which were just discussed, are the same in all simulations. The shocks to β_t are set to zero, so that variable simply decays back toward its mean over the period, following a path that is identical across all simulations.

Panel A of Figure 6 shows the resulting distribution of the simulated trajectories for \hat{r}_t . This distribution spans empirical estimates of the shadow-rate path during the ELB period, including those of Krippner (2012) and Wu and Xia (2016). Panel B shows the distribution of the Q_t trajectories, converted to cumulative percentage changes in ten-year-equivalent bonds outstanding for ease of interpretation, using equation (27). It is more difficult to know what the “right” path of this variable ought to be (see footnote 27), but the distribution covers a fairly wide range of possibilities.

5.2. Yield curve responses and their decomposition

With the simulated distributions of the state-variable trajectories in hand, I use the model to extract the yield-curve responses. To report the results, for each period in each simulation I calculate how the yield curve changes, relative to how it would have changed if there had been no shock in

²⁵This assumption is justified because QE purchases account for very little of the unconditional variation in the duration risk of investors' portfolios. Allowing for a positive σ_Q does not substantively change the results below, as long as it is less than the unconditional variance of σ_β .

²⁶The ELB was officially reached on December 16, 2008, when the FOMC cut the target federal funds rate from 1% to a range of 0 to 25 basis points. However, from the Treasury market's perspective the effective date may have been slightly earlier. The three-month yield declined 102 basis points over the intermeeting period leading up to December 16, in anticipation of the cut. In addition, the first announcement of asset purchases came on November 25. Using a starting value for bond supply based on the situation as of October 29, 2008, ensures that it does not include these pre-ELB influences of unconventional policy.

²⁷One might use data on the Fed's holdings as an observed measure of Q_t , and such a path is in fact spanned by the set of trajectories I simulate. I allow for more ambiguity, however, because the term premium depends on expectations of future asset purchases, not just the amount the Fed currently holds. In the model here, $E_t[Q_{t+h}]$ depends only on Q_t , making the distinction between current balances and expected future balances fuzzy. Greenwood et al. (2015b) explore this distinction in detail.

that period. I then sum these differences across periods within each simulation. This procedure is analogous to empirical event studies that attempt to isolate and accumulate the immediate impact of policy shocks without accounting for their dynamics. (Here, there is an “event” in every period.) Specifically, letting $y^{(\tau)}(\hat{r}_t, \beta_t, Q_t)$ denote the τ -maturity yield as a function of the state variables, I calculate

$$D_i^{(\tau)} = \sum_{t=1}^{28} \left[y^{(\tau)}(\hat{r}_{i,t}, \beta_t, Q_{i,t}) - y^{(\tau)}(\hat{r}_{i,t} - e_{i,t}^r, \beta_t, Q_{i,t} - e_{i,t}^Q) \right] \quad (31)$$

The distribution of $D_i^{(\tau)}$ across simulations is shown in Figure 7.A, with the corresponding calculation for forward rates shown in 7.B.

The median decline in the ten-year yield produced by the simulated unconventional policy shocks is 207 basis points. This estimate does not differ much across simulations, with the middle 90% of the distribution spanning only the range of -211 to -199 basis points. It is worth noting that these are similar magnitudes to the effects that have been estimated in event studies of unconventional policy. For example, looking at 23 important policy announcements during the ELB period, and controlling for macroeconomic news, Altavilla and Giannone (2016) find a net effect on the ten-year yield of -176 basis points. Comparisons between these kinds of results and the model cannot be made precise because empirical event studies necessarily capture only a subset of the relevant shocks, but they suggest that the model’s results are quantitatively realistic.

The model also reproduces another key stylized fact from the event-study literature. Rogers et al. (2014) and Greenwood et al. (2015b) show that unconventional policy announcements typically resulted in a hump-shaped reaction across the forward curve, with forward rates in the 5- to 10-year range moving the most. The model generates exactly this pattern (see Figure 7.B). As was evident in Figure 5, it does so because of the non-monotonic effects of shadow-rate shocks induced by the ELB, which operate through both the expectations and term-premium components of yields.

I decompose the cumulative contemporaneous yield-curve reaction shown in Figure 7 into various channels of unconventional policy. While this calculation captures the sources of the changes in yields in the periods when shocks occurred, it does not account for the dynamic effects of those shocks. Therefore, I also calculate a decomposition of the total model-implied variance in yields during the ELB period. In both exercises, the breakdown is calculated by computing what the change in yields would have been if only the shadow-rate or the Fed balance-sheet shocks had occurred (again, relative to a baseline case in which there are no shocks at all). In the case of the shadow-rate shocks, the response can be further decomposed into the expectations and term-premium components. Finally, because of the nonlinearities, the responses to the individual shocks do not sum exactly to the total response when both types of shocks occur simultaneously; an “interaction” term captures the residual.

Table 5 shows the results of these decompositions according to the medians across simulations, with the 5% and 95% quantiles reported in parentheses. Panel A presents the decomposition of the cumulative contemporaneous effects, and panel B presents the total-variance decomposition. Looked at either way, the shadow-rate shocks are responsible for considerably more of the change in yields than the Fed balance-sheet shocks are. For example, as shown in column 4 of panel A, the model implies that the duration effects associated with QE lowered the ten-year yield by a total of just 47 basis points in the periods when they occurred. The shadow-rate shocks explain over 75% of the contemporaneous declines in yields at this maturity, and at maturities of less than five years they explain even more of the change (columns 2 and 3). The expectations component constitutes the bulk of the effects of the shadow-rate shocks, but the term-premium effects of such shocks are

also significant. They account for 70 basis points of the decline in yields at the ten-year maturity.

In panel B, the nonlinear interactions in column 6 loom larger, so that the precise contributions of each factor to the overall variance of yields are somewhat less clear. Nonetheless, in the case with Fed balance-sheet shocks alone (column 4), yields beyond the two-year horizon have less than 3% of the variance that they have when both shocks are present. In contrast, in the case with shadow-rate shocks alone (columns 2 and 3), the variance of yields is similar to the variance when both shocks are present. At intermediate and long maturities, the term-premium effects of the shadow-rate shocks explain about the same fraction of the total variance in yields that the expectations component explains.

The above exercise has assumed that shadow-rate shocks and Fed balance-sheet shocks are independent. In practice, it is likely that they were positively correlated at the ELB—for example, several FOMC announcements during this period contained information about both the future short rate and QE policies. While it is difficult to know exactly what the correlation was, I check the sensitivity of the results to the extreme assumption that the two shocks were perfectly correlated. While this still (by construction) results in the same length of time at the ELB and the same amount of cumulative bond purchases as shown in Figure 6, it changes the trajectories and their correlations, which could affect the results due to nonlinearities. The details of the simulations are the same as described above and in Appendix C, except that $e_{i,t}^{\hat{r}}$ and $e_{i,t}^{\beta}$ are drawn from perfectly correlated distributions for each i and t .

The results are shown in Table 6. Comparison with Table 5 demonstrates that the correlation assumption has very little quantitative effect on the results. In particular, it remains the case that Fed balance-sheet shocks have much smaller overall effects than shadow-rate shocks do and that a significant portion of the effect of shadow-rate shocks operates through their endogenous effects on term premia.

5.3. Discussion

The above results imply that the duration removal that occurred through Federal Reserve asset purchases may have had significantly smaller effects on yields than has sometimes been suggested. For example, summarizing the empirical literature, Williams (2014) concludes that each \$600 billion of asset purchases likely reduced the ten-year yield by 15 to 25 basis points. The Fed bought nearly \$4 trillion of securities in total, so if the Williams estimates are interpreted as entirely reflecting duration removal they suggest an effect of the duration channel that is about three times larger than the one produced by the model.

Yet the empirical literature is far from conclusive that the duration channel is the only—or even the primary—channel through which QE operates. Evidence in Krishnamurthy and Vissing-Jorgensen (2011) and Bauer and Rudebusch (2015), for example, indicates that the signaling channel may have been the most important aspect of asset purchases. The simulations above are consistent with those results. Furthermore, it may be that previous empirical work on QE has confused the duration channel with the link between policy expectations and term premia that emerges in the model of this paper. For example, some empirical studies assume that declines in term premia around QE announcements must be due to duration effects (e.g., Gagnon et al., 2011; Christensen and Rudebusch, 2012). The model shows that this is not necessarily true. The signaling channel of bond purchases (or explicit forward guidance that is issued concurrently with QE announcements) can also lower the term premium through its effect on interest-rate volatility.²⁸

²⁸It is also worth repeating that the bond-supply effects generated by the model are consistent with evidence that

These observations show that the effects of shadow-rate and Fed balance-sheet shocks produced in the model simulations are plausible with respect to the data. Yet a question still remains about why those effects have the magnitudes that they do. In particular, one may wonder why \$4 trillion of asset purchases “only” buys a 47-basis-point reduction in the ten-year yield. This is a difficult question to answer, because the quantitative outcomes of the simulations depend on all of the model parameters, starting values, shock processes, and other simulation choices. However, two features of the model seem particularly noteworthy.

First, as discussed throughout the paper, duration effects are attenuated by the ELB. The same amount of asset purchases would have had a larger effect in the simulations if those purchases had taken place in an environment with the short rate closer to normal levels. Depending on the details, the factor loadings shown in Figure 4 suggest that the duration effects could have been over 50% larger in such a situation.

Second, the yield data themselves place some limits on how large the duration channel can be within the context of the type of structural model used here. This can be seen most easily by returning to the one-factor model of Section 3. Equation (16) shows that excess returns are linear in bond supply. Thus, any change in parameters that increases the sensitivity of excess returns to bond supply must increase the average *level* of excess returns in equal proportion. To triple the effects of QE (for example, by tripling the value of the parameter a), one would also have to accept a tripling of the average risk premium. In the model of the paper, the parameters are set to match the 6.7% average level of the ten-year yield observed in the data. A tripling of the risk premium would raise this average to 9.7%—a value that would easily be rejected empirically. (The parameter change would also increase the volatility of yields significantly beyond the observed levels.) Thus, while one could generate larger effects of QE in the model through alternative choices of parameters, nearly any such modification would put the model at odds with other features of the data.²⁹

5.4. Policy options in alternative environments

Because they take different units, it is not meaningful to ask whether the shadow-rate shocks or the bond-supply shocks are “more powerful” in general. However, one can compare their *relative* effectiveness in different states of the world. One reason that such a comparison may be interesting is that policymakers, who presumably have some notion of the implicit cost of implementing each type of policy, may favor one over the other depending on the circumstances.

To measure relative efficacy, I compute the size of the bond-supply shock that would be required to generate the same effect on the τ -period yield that a 25-basis-point decline in the shadow rate has. Specifically, again letting $y^{(\tau)}(\hat{r}_t, \beta_t, Q_t)$ denote the τ -maturity yield as a function of the state variables, I solve for $\Delta\beta$ such that

$$y^{(\tau)}(\hat{r}_t - .0025, \beta_t, Q_t) = y^{(\tau)}(\hat{r}_t, \beta_t + \Delta\beta, Q_t) \quad (32)$$

does not stem from QE episodes. As noted in Section 4.4, the effects of bond supply in the model largely match those estimated in regressions of the Greenwood-Vayanos (2014) type. That evidence arguably gives a cleaner read on duration effects than the QE event studies do, because it avoids some of the identification problems of event studies, is not contaminated by a monetary-policy signaling channel, and is largely free of distortions associated with the ELB.

²⁹One might worry that the parameters matched to the full-sample moments are not appropriate for the period when QE was conducted. However, Table 3 and other results presented above show that the model continues to do a good job of describing the behavior of yields during this time.

using a range of initial values (\hat{r}_t, β_t) . I repeat an analogous exercise to solve for the relative efficacy of Fed balance-sheet shocks, ΔQ .³⁰ In an affine model, the values $\Delta\beta$ and ΔQ that solve these equations are constant across the state space. In the shadow-rate model, as was evident in Figures 4 and 5, the elasticities of yields with respect to shadow-rate and bond-supply shocks differ in different areas of the state space, and therefore their relative efficacy also differs.

Figure 8 presents contour maps of relative efficacy for 10- and 15-year yields, with darker colors indicating bigger values—i.e., areas of the space in which bond-supply shocks have relatively large effects compared to those of shadow-rate shocks. Both β_t and Q_t achieve their greatest relative efficacy in the southwest quadrant of the maps, where both \hat{r}_t and β_t are deeply negative. As noted earlier, both bond-supply and shadow-rate shocks are attenuated when the shadow rate is below the ELB. However, when β_t is negative (i.e., more duration in the market), the attenuation of the shadow-rate shocks is greater than the attenuation of the bond-supply shocks.³¹ Thus, for example, a shock to β_t of about 0.10 or a shock to Q_t of about 0.14 in this region is sufficient to lower the ten-year yield by the same amount that a 25-basis-point shock to \hat{r}_t would achieve. In contrast, at the unconditional means of the states, the respective sizes of the β_t and Q_t shocks required are closer to 0.14 and 0.20, respectively.

Interestingly, this high-relative-efficacy region for the bond-supply shocks is approximately the region of the space in which the Fed asset purchases were conducted in practice. The greatest removal of duration from the market occurred during the QE and maturity extension programs that mostly operated between 2011 and 2013. During that time, empirical shadow-rate term-structure models show \hat{r}_t near its nadir, with the Krippner (2012) estimate, for example, averaging -4.5% over those three years. Meanwhile, the Treasury was lengthening the maturities of its issuance, so that the average duration outstanding stood near the upper end of its historical range. Moreover, fiscal expansion increased the total quantity of Treasury debt outstanding, further boosting the amount of interest-rate risk held by investors.³² Thus, one possible interpretation of the Fed’s actions during this time is that it saw the cost-benefit calculations around its policy options changing. During normal times, the Fed has a revealed preference for not engaging in asset purchases. This preference may have shifted during the ELB period if the FOMC perceived that the marginal benefits of forward guidance declined sufficiently relative to those of asset purchases.

6. Conclusion

This paper has augmented a model of risk-averse arbitrage in the bond market to account for the effective lower bound on nominal interest rates. At and near the ELB, the effects of bond supply are damped, the expectations and term-premium components of yields become endogenously correlated, and shocks to short-rate expectations have their largest effect on intermediate-maturities forward rates. When calibrated to the long-run features of yields, the model successfully reproduces the conditional moments of the yield curve, particularly near the ELB, as well as empirical evidence on the effects of bond supply on yields. When considering shocks that approximate the experience of unconventional monetary policy in the U.S., the main finding is that the duration effects of shocks

³⁰Note that the solution for $\Delta\beta$ is the same regardless of whether we use the baseline process (23) or include Q_t in the model as in (29). In the latter case, the initial value of Q_t is set to zero.

³¹This result can be seen to some extent in Figure 5, panel B.

³²Over the 2011 - 2013 period, the maturity-weighted debt-to-GDP ratio averaged 4.4, compared to a pre-ELB average of 2.4, according to the CRSP data.

to the Federal Reserve’s balance sheet are relatively weak, accounting for less than one-quarter of the overall change in the ten-year yield and even less at shorter maturities. Instead, the majority of the effects of such policies come through the expectations component of yields. The term premium effects of changes in anticipated short rates—a channel that does not exist in affine models and has been ignored by previous literature—also plays a significant role.

In addition to their implications for policy choices during future ELB episodes, the results are relevant for empirical work in this area. For example, economists have sometimes studied questions about the yield curve at the ELB (including questions about unconventional monetary policy) by extrapolating the results of linear models estimated on pre-ELB data. The analysis here suggests that that approach could be quite dangerous. There are reasons to expect factor loadings to change both quantitatively and qualitatively at the ELB, with some transmission channels becoming diminished and others enhanced. Because bond prices are forward-looking and depend on both expectations and uncertainty, nonlinearities matter and can operate on the term structure in subtle ways.

Acknowledgements

For helpful comments and discussions I thank Stefania D’Amico, Sam Hanson, Don Kim, Dimitri Vayanos, Andrea Vedolin, and seminar participants at the Federal Reserve Bank of Chicago, the 2017 FRB San Francisco/Bank of Canada/Simon Fraser University conference on Advances in Fixed Income and Macro-Finance Research, the 2018 American Finance Association meetings, and the 2018 Banque de France Workshop on Monetary Policy and Asset Prices. Zachry Wang provided excellent research assistance. The views expressed here do not reflect official positions of the Federal Reserve.

References

- Altavilla, C., Carboni, G., Motto, R., 2015. Asset purchase programmes and financial markets: Lessons from the euro area. European Central Bank Working Papers, no 1864.
- Altavilla, C., Giannone, D., 2016. The effectiveness of non-standard monetary policy measures: Evidence from survey data. *Journal of Applied Econometrics* 32, 952–964.
- Bauer, M. D., Rudebusch, G. D., 2014. The signaling channel for federal reserve bond purchases. *International Journal of Central Banking* 10 (3), 233–289.
- Bhattarai, S., Eggertsson, G. B., Gafarov, B., 2015. Time consistency and the duration of government debt: A signalling theory of quantitative easing. NBER Working Paper 21336.
- Cahill, M. E., D’Amico, S., Li, C., Sears, J. S., 2013. Duration risk versus local supply channel in treasury yields: Evidence from the federal reserve’s asset purchase announcements. Federal Reserve Board FEDS working paper 2013-35.
- Carpenter, S. B., Ihrig, J. E., Klee, E. C., Quinn, D. W., Boote, A. H., 2015. The federal reserve’s balance sheet and earnings: A primer and projections. *International Journal of Central Banking* 11 (2), 237–283.

- Christensen, J. H. E., Rudebusch, G. D., 2012. The response of interest rates to u.s. and u.k. quantitative easing. *Economic Journal* 122, F237–414.
- D’Amico, S., King, T. B., 2013. Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics* 108 (2), 425–448.
- Femia, K., Friedman, S., Sack, B., 2013. The effects of policy guidance on perceptions of the fed’s reaction function. Federal Reserve Bank of New York, Staff Report no.652.
- Gagnon, J. E., Raskin, M., Remache, J., Sack, B. P., 2011. The financial market effects of the federal reserve’s large-scale asset purchases. *International Journal of Central Banking* 7 (1), 3–43.
- Gilchrist, S., Lopez-Salido, D., Zakrajsek, E., 2015. Monetary policy and real borrowing costs at the zero lower bound. *American Economic Journal: Macroeconomics* 7 (1), 77–109.
- Greenwood, R., Hanson, S., Rudolph, J. S., Summers, L. H., 2015a. Chapter 1: The optimal maturity of government debt; chapter 2: Debt management conflicts between the u.s. treasury and the federal reserve. Book chapters in in *The \$13 Trillion Question: How America Manages Its Debt*, edited by David Wessel, 43-89. Brookings Institution Press.
- Greenwood, R., Hanson, S., Vayanos, D., December 2015b. Forward guidance in the yield curve: Short rates versus bond supply. NBER Working Paper No. 21750.
- Greenwood, R., Vayanos, D., 2014. Bond supply and excess bond returns. *Review of Financial Studies* 27 (3), 663–713.
- Gurkaynak, R. S., Sack, B., Wright, J. H., 2007. The u.s. treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54 (8), 2291–304.
- Haddad, V., Sraer, D. A., July 2015. The banking view of bond risk premia. Working Paper.
- Hamilton, J. D., Wu, J. C., 2012. The effectiveness of alternative monetary policy tools in a zero lower bound environment. *Journal of Money, Credit and Banking* 44 (s1), 3–46.
- Hanson, S. G., 2014. Mortgage convexity. *Journal of Financial Economics* 113 (2), 270–299.
- Hattori, M., Schrimpf, A., Sushko, V., 2016. The response of tail-risk perceptions to unconventional monetary policy. *American Economic Journal: Macroeconomics* 8 (2), 111–36.
- Hayashi, F., April 2016. Affine term structure pricing with bond supply as factors. FRB Atlanta CQER Working Paper 2016-1.
- Joyce, M. A. S., Lasasosa, A., Stevens, I., Tong, M., 2011. The financial market impact of quantitative easing in the united kingdom. *International Journal of Central Banking* 7 (3), 32–49.
- Kim, D. H., Singleton, K. J., 2012. Term structure models and the zero bound: An empirical investigation of japanese yields. *Journal of Econometrics* 170 (1), 32–49.
- King, T. B., 2015. A portfolio-balance approach to the nominal term structure. FRB Chicago Working Paper 2013-18.
- Krippner, L., March 2012. Modifying gaussian term structure models when interest rates are near the zero lower bound. CAMA Working Paper 2012-05.

- Krishnamurthy, A., Vissing-Jorgensen, A., 2011. The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity* (2), 215–265.
- Malkhozov, A., Mueller, P., Vedolin, A., Venter, G., 2016. Mortgage risk and the yield curve. *Review of Financial Studies* 29 (5), 1220–53.
- Monfort, A., Pegoraro, F., Renne, J.-P., Roussellet, G., forthcoming. Staying at zero with affine processes: An application to term structure modeling. *Journal of Econometrics*.
- Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55 (3), 703–8.
- Rogers, J. H., Scotti, C., Wright, J. H., 2014. Evaluating asset-market effects of unconventional monetary policy: A cross-country comparison. *Economic Policy* 29 (80), 749–799.
- Swanson, E. T., 2011. Let’s twist again: A high-frequency event study analysis of operation twist and its implications for qe2. *Brookings Papers on Economic Activity* (1), 151–88.
- Swanson, E. T., 2017. Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. NBER Working Paper 21816.
- Vayanos, D., Vila, J.-L., November 2009. A preferred-habitat model of the term structure of interest rates. NBER Working Paper 15487.
- Williams, J. C., January 2014. Monetary policy at the zero lower bound: Putting theory into practice. Hutchins Center on Fiscal and Monetary Policy at Brookings, January 16.
- Woodford, M., 2012. Methods of policy accommodation at the interest-rate lower bound. Federal Reserve Bank of Kansas City, Jackson Hole symposium, September 16.
- Wright, J., 2012. What does monetary policy do to long-term interest rates at the zero lower bound? *Economic Journal* 122, F447–66.
- Wu, J. C., Xia, F. D., 2016. Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit, and Banking* 48 (2-3), 253–291.

Appendices

Appendix A. Approximate solution to the one-factor model

I look for coefficients $A_t^{(\tau)}$ that jointly satisfy (7) and (15) for all maturities τ and all values of the shadow rate \hat{r}_t .

Log bond returns can be decomposed as,

$$dp_t^{(\tau)} = E_t \left[dp_t^{(\tau)} \right] + \frac{\partial p_t^{(\tau)}}{\partial \hat{r}_t} \sigma dB_t \quad (\text{A.1})$$

so,

$$\text{cov}[dp_t^{(\tau)}, dp_t^{(s)}] = \frac{\partial p_t^{(\tau)}}{\partial \hat{r}_t} \frac{\partial p_t^{(s)}}{\partial \hat{r}_t} \sigma^2 \quad (\text{A.2})$$

Expanding (7) and (8), time- t log bond prices are

$$p_t^{(\tau)} = - \int_0^\tau E_t[r_{t+s}] ds - a\zeta \int_0^\tau E_t \left[\int_0^T \text{cov}_{t+s} \left[dp_{t+s}^{(\tau)}, dp_{t+s}^{(j)} \right] dj - \frac{1}{2} \text{var}_{t+s} \left[dp_{t+s}^{(\tau)} \right] \right] ds \quad (\text{A.3})$$

Substituting (A.2),

$$p_t^{(\tau)} = - \int_0^\tau E_t[r_{t+s}] ds - a\zeta \sigma^2 \int_0^\tau E_t \left[\int_0^T \frac{\partial p_{t+s}^{(\tau)}}{\partial \hat{r}_t} \frac{\partial p_{t+s}^{(j)}}{\partial \hat{r}_t} dj - \frac{1}{2} \left(\frac{\partial p_{t+s}^{(\tau)}}{\partial \hat{r}_t} \right)^2 \right] ds \quad (\text{A.4})$$

Taking the derivative,

$$\begin{aligned} \frac{\partial p_t^{(\tau)}}{\partial \hat{r}_t} &= - \int_0^\tau \frac{\partial E_t[r_{t+s}]}{\partial \hat{r}_t} ds \\ &\quad - a\zeta \sigma^2 \int_0^\tau E_t \left[\int_0^T \left(\frac{\partial^2 p_{t+s}^{(\tau)}}{\partial \hat{r}_t^2} \frac{\partial p_{t+s}^{(j)}}{\partial \hat{r}_t} + \frac{\partial p_{t+s}^{(\tau)}}{\partial \hat{r}_t} \frac{\partial^2 p_{t+s}^{(j)}}{\partial \hat{r}_t^2} \right) dj - \frac{\partial p_{t+s}^{(\tau)}}{\partial \hat{r}_t} \frac{\partial^2 p_{t+s}^{(\tau)}}{\partial \hat{r}_t^2} \right] ds \end{aligned} \quad (\text{A.5})$$

Approximating $A_t^{(\tau)}$ by setting second derivatives in this expression to zero,

$$\begin{aligned} A_t^{(\tau)} &= - \frac{\partial p_t^{(\tau)}}{\partial \hat{r}_t} \approx \int_0^\tau \frac{\partial E_t[r_{t+s}]}{\partial \hat{r}_t} ds \\ &= \int_0^\tau e^{-\kappa s} E_t \left[\frac{\partial r_{t+s}}{\partial \hat{r}_{t+s}} \right] ds \end{aligned} \quad (\text{A.6})$$

where the last equality follows because \hat{r}_t is a standard Ornstein-Uhlenbeck process. The term in brackets is equal to 1 if \hat{r}_{t+s} is above the ELB and 0 otherwise. Therefore, its expectation is just the probability that $\hat{r}_{t+s} > b$, and we have

$$A_t^{(\tau)} \approx \int_0^\tau e^{-\kappa s} \Phi_t^{(s)} ds \quad (\text{A.7})$$

which is the result given in the text. (One can also show this last step by directly differentiating (13) with respect to \hat{r}_t .)

Appendix B. Solution Algorithm

Consider a small but finite time interval Δ . For $\tau > \Delta$ we have

$$\begin{aligned} p_t^{(\tau)} &= \mathbb{E}_t \left[p_{t+\Delta}^{(\tau-\Delta)} \right] - \int_0^\Delta \mathbb{E}_t \left[dp_{t+\delta}^{(\tau-\delta)} \right] d\delta \\ &\approx \mathbb{E}_t \left[p_{t+\Delta}^{(\tau-\Delta)} \right] - \frac{\Delta}{2} \left(\mathbb{E}_t \left[dp_{t+\Delta}^{(\tau-\Delta)} \right] + \mathbb{E}_t \left[dp_t^{(\tau)} \right] \right) \end{aligned} \quad (\text{B.1})$$

From (7), (23), and (25), the expectations terms on the right-hand side are given by

$$\mathbb{E}_t \left[dp_t^{(\tau)} \right] \approx r_t + \frac{1}{\Delta} \left(a \int_0^T \left[\zeta + \left(1 - \frac{2s}{T} \right) \beta_t \right] \text{cov}_t \left[p_{t+\Delta}^{(\tau-\Delta)}, p_{t+\Delta}^{(s-\Delta)} \right] ds - \frac{1}{2} \text{var}_t \left[p_{t+\Delta}^{(\tau-\Delta)} \right] \right) \quad (\text{B.2})$$

and similarly for $\mathbb{E}_t \left[dp_{t+\Delta}^{(\tau-\Delta)} \right]$, where we have made use of the approximation

$$\text{cov}_t \left[dp_t^{(\tau)}, dp_t^{(s)} \right] \approx \frac{1}{\Delta} \text{cov}_t \left[p_{t+\Delta}^{(\tau-\Delta)}, p_{t+\Delta}^{(s-\Delta)} \right] \quad (\text{B.3})$$

Note that the approximation errors in (B.1) through (B.3) vanish as Δ goes to zero.

Discretize the state and maturity space into $N_\tau \times N_{\hat{r}} \times N_\beta$ nodes. Let τ^n , \hat{r}^n , and β^n denote the values of the maturity and state variables at node n . Set $p(0, \hat{r}_t, \beta_t) = -\max[\hat{r}_t, b]$ for all β^n . Let $\pi(\hat{r}_{t+\Delta}, \beta_{t+\Delta} | \hat{r}_t, \beta_t)$ be the conditional PDF of the states associated with the laws of motion (12) and (24). Finally, transform log prices from the time/maturity domain to the state-space domain by writing them as a function $p(\cdot)$:

$$p_t^{(\tau)} = p(\tau, \hat{r}_t, \beta_t) \quad (\text{B.4})$$

Our objective is to solve for this function by calculating state-contingent values of all of the objects on the right-hand side of (B.1).

The algorithm proceeds as follows:

Step 0. Set $i = 0$. Begin with an initial guess of the pricing function $p^0(\cdot)$. For example, choose $p^0(\tau, \hat{r}_t, \beta_t) = -\max[\hat{r}_t, b]$ for all τ, \hat{r}_t, β_t .

Step 1. At each node n , evaluate the functions

$$\begin{aligned} F_1^i(\tau^n, \hat{r}^n, \beta^n) &\equiv \int \int \pi(\hat{r}, \beta | \hat{r}^n, \beta^n) p^i(\tau^n - \Delta, \hat{r}, \beta) d\hat{r} d\beta \\ &= \mathbb{E}_t \left[p_{t+\Delta}^{(\tau^n - \Delta)} \mid (\hat{r}_t, \beta_t) = (\hat{r}^n, \beta^n) \right] \end{aligned} \quad (\text{B.5})$$

$$\begin{aligned}
F_2^i(\tau^n, \hat{r}^n, \beta^n) &\equiv \int \int \pi(\hat{r}, \beta | \hat{r}^n, \beta^n) [p^i(\tau^n - \Delta, \hat{r}, \beta) - F_1^i(\tau^n, \hat{r}^n, \beta^n)]^2 d\hat{r} d\beta \\
&= \text{var}_t \left[p_{t+\Delta}^{(\tau^n - \Delta)} \mid (\hat{r}_t, \beta_t) = (\hat{r}^n, \beta^n) \right]
\end{aligned} \tag{B.6}$$

$$\begin{aligned}
F_3^i(\tau^n, \hat{r}^n, \beta^n) &\equiv \int \int \pi(\hat{r}, \beta | \hat{r}^n, \beta^n) F_2^i(\tau^n - \Delta, \hat{r}, \beta) d\hat{r} d\beta \\
&= \text{E}_t \left[\text{var}_{t+\Delta} \left[p_{t+2\Delta}^{(\tau^n - 2\Delta)} \right] \mid (\hat{r}_t, \beta_t) = (\hat{r}^n, \beta^n) \right]
\end{aligned} \tag{B.7}$$

$$\begin{aligned}
F_4^i(\tau^n, s, \hat{r}^n, \beta^n) &\equiv \left[\zeta + \left(1 - \frac{2s}{T} \right) \beta^n \right] \int \int \{ \pi(\hat{r}, \beta | \hat{r}^n, \beta^n) [p^i(\tau^n - \Delta, \hat{r}, \beta) - F_1^i(\tau^n, \hat{r}^n, \beta^n)] \\
&\quad \times [p^i(s, \hat{r}, \beta) - F_1^i(s, \hat{r}^n, \beta^n)] \} d\hat{r} d\beta \\
&= \left[\zeta + \left(1 - \frac{2s}{T} \right) \beta_t \right] \text{cov}_t \left[p_{t+\Delta}^{(\tau^n - \Delta)}, p_{t+\Delta}^{(s - \Delta)} \mid (\hat{r}_t, \beta_t) = (\hat{r}^n, \beta^n) \right]
\end{aligned} \tag{B.8}$$

and

$$\begin{aligned}
F_5^i(\tau^n, s, \hat{r}^n, \beta^n) &\equiv \int \int \{ \pi(\hat{r}, \beta | \hat{r}^n, \beta^n) \left[\zeta + \left(1 - \frac{2s}{T} \right) \beta \right] F_2^i(\tau^n - \Delta, s, \hat{r}, \beta) d\hat{r} d\beta \\
&= \text{E}_t \left[\left[\zeta + \left(1 - \frac{2s}{T} \right) \beta_{t+\Delta} \right] \text{cov}_{t+\Delta} \left[p_{t+2\Delta}^{(\tau^n - 2\Delta)}, p_{t+2\Delta}^{(s - \Delta)} \right] \mid (\hat{r}_t, \beta_t) = (\hat{r}^n, \beta^n) \right]
\end{aligned} \tag{B.9}$$

with all five functions set to zero when $\tau = 0$.

Step 2. Update the pricing function by substituting the functions above for the moments in (B.1) and evaluating them at each node. That is, calculate

$$\begin{aligned}
p^{i+1}(\tau^n, \hat{r}^n, \beta^n) &= F_1^i(\tau^n, \hat{r}^n, \beta^n) - \max[\hat{r}^n, b] \\
&\quad - \frac{a\Delta}{2} \sum_{s=1}^T [F_4^i(\tau^n, s, \hat{r}^n, \beta^n) + F_5^i(\tau^n, s, \hat{r}^n, \beta^n)] \\
&\quad + \frac{\Delta}{4} [F_2^i(\tau^n, \hat{r}^n, \beta^n) + F_3^i(\tau^n, \hat{r}^n, \beta^n)]
\end{aligned} \tag{B.10}$$

Set $i = i + 1$.

Repeat steps (1) and (2) to convergence.

The expectations in Step 1 are computed numerically using the probability function $\pi(\cdot)$ and the pricing function $p^i(\cdot)$. The integration is performed by quadrature and, to ensure accuracy, relies on a much finer grid than the price computation in Step 2 does. To obtain bond prices over this refinement of the space, the values of $p^i(\cdot)$ are interpolated between each pair of nodes,

at each iteration, using a cubic spline. At the edges of the discretized space, to avoid explosive behavior, prices are log-linearly extrapolated for the purposes of computing expectations. (So long as the edges are far away from the region of the space that is being considered, the conditional expectations used there have little influence on the results.)

In the baseline model of the paper, I use $N_\tau = 60$, $N_{\hat{r}} = 101$, and $N_\beta = 25$, for a total of 151,500 nodes, distributed uniformly in each dimension over the intervals $\tau = [1, 60]$, $\hat{r} = [-0.25, 0.35]$, and $\beta = [-6.0, 6.0]$. Expanding the density of the nodes or their range beyond this point had no noticeable effect on the results reported in the paper. The algorithm converges to three significant digits in approximately 400 iterations.

Appendix C. Simulation details

This appendix describes the construction of the simulated distributions of the state variables used in the pseudo event study, as depicted in Figure 6.

Within each simulation i , I draw a discrete quarterly series of shadow-rate shocks $\{e_{i,1}^{\hat{r}}, \dots, e_{i,28}^{\hat{r}}\}$, each of which represents the change in \hat{r}_t over the course of each quarter that was unanticipated as of the beginning of that quarter. I.e., $e_{i,t} = \hat{r}_t - E_{t-1}[\hat{r}_t]$. Without further restrictions each shock has mean zero and standard deviation $\tilde{\sigma}$ given by

$$\tilde{\sigma} = \frac{1 - e^{-\kappa}}{\kappa} \sigma \quad (\text{C.1})$$

However, *ex post* it clearly had to be the case that the average shadow-rate shock during the ELB period was negative. To ensure that this condition is satisfied, for each simulation i , I draw the series of shadow-rate shocks from the distribution $N[\mu_i, \tilde{\sigma}]$, where $\mu_i < 0$ is chosen to make the terminal value of the shadow rate $\hat{r}_{i,28}$ exactly equal b . I reject any draw in which the simulated value of \hat{r}_t ever rises above b .

To determine the size of the Fed balance-sheet shocks $\{e_{i,1}^Q, \dots, e_{i,28}^Q\}$, note that equation (27) implies that changes in the supply factors translate into percentage changes in 10-year-equivalent bonds as follows:

$$\begin{aligned} \% \Delta 10YE_{t+s} &= \frac{\int_0^T \tau \left[\left(1 - \frac{2\tau}{T}\right) (\beta_{t+s} + Q_{t+s} - \beta_t - Q_t) \right] d\tau}{\int_0^T \tau \left[\zeta + \left(1 - \frac{2\tau}{T}\right) (\beta_t + Q_t) \right] d\tau} \\ &= -\frac{\Delta \beta_{t+s} + \Delta Q_{t+s}}{3\zeta - \beta_t - Q_t} \end{aligned} \quad (\text{C.2})$$

Since the e_t^β are taken to be zero in this exercise, plugging in the actual percentage change in 10-year equivalents that resulted from QE allows one to uniquely solve for ΔQ_{t+s} , given initial values. In particular, at the end of the simulation we must have $Q_{28} = 0.32$ in order to achieve a reduction in ten-year equivalents of 21% relative to a case in which $Q_{28} = 0$.³³

This calculation provides a value for the cumulative effect of the Fed balance-sheet shocks on bond supply, but it does not tell us about the individual values of those shocks. I take a conservative approach by considering the widest possible distribution for the shocks, while respecting

³³The difference between the path of ten-year equivalents in the simulation and the counterfactual case in which no QE occurs is $-\frac{\Delta Q_{t+s}}{3\zeta - \beta_t - Q_t} = -\frac{Q_{28}}{3(0.37) + 0.42}$. Setting this equal to -0.21 gives $Q_{28} = 0.32$.

the empirical facts that (1) QE balances never fell below their starting value of zero, and (2) QE attained its maximum value at the end of the ELB period. Specifically, for each simulation i , I take draws $\{\tilde{e}_{i,1}, \dots, \tilde{e}_{i,28}\}$ from $N[0, 1]$ and compute the balance-sheet shocks $\{e_{i,1}^Q, \dots, e_{i,28}^Q\} = \sigma_i\{\tilde{e}_{i,1}, \dots, \tilde{e}_{i,28}\}$, where σ_i is the the largest value that is consistent with $\min[\{Q_{i,1}, \dots, Q_{i,28}\}] > 0$ and $\max[\{Q_{i,1}, \dots, Q_{i,28}\}] = Q_{i,28}$.

Table 1. Regressions of long-term yields on Treasury duration and short-term yields

Independent variables							Adj. R ²
Dep. Var.	WAM of Treas. debt			1y yield			
	Pre-ELB	ELB	Break <i>t-stat</i>	Pre-ELB	ELB	Break <i>t-stat</i>	
5y yield	0.140 (0.095)	0.002 (0.101)	-2.17	0.842*** (0.050)	2.271*** (0.785)	1.84	0.951
10y yield	0.221* (0.121)	0.058 (0.116)	-2.25	0.736*** (0.060)	3.028** (1.203)	1.92	0.901
15y yield	0.261* (0.133)	0.110 (0.126)	-2.05	0.688*** (0.065)	2.966** (1.276)	1.80	0.870

Independent variables							Adj. R ²
Dep. Var.	WAM of Treas. debt			2y yield			
	Pre-ELB	ELB	Break <i>t-stat</i>	Pre-ELB	ELB	Break <i>t-stat</i>	
5y yield	0.102* (0.373)	-0.002 (0.060)	-2.57	0.901*** (0.032)	1.910*** (0.217)	4.74	0.981
10y yield	0.187** (0.094)	0.053 (0.088)	-2.25	0.794*** (0.048)	2.328*** (0.429)	3.61	0.942
15y yield	0.227** (0.109)	0.113 (0.108)	-1.62	0.746*** (0.056)	2.167*** (0.537)	2.68	0.915

Independent variables							Adj. R ²
Dep. Var.	Maturity-weighted debt/GDP			1y yield			
	Pre-ELB	ELB	Break <i>t-stat</i>	Pre-ELB	ELB	Break <i>t-stat</i>	
5y yield	0.179* (0.010)	-0.058 (0.083)	-2.65	0.850*** (0.049)	2.097*** (0.679)	1.89	0.952
10y yield	0.250* (0.129)	-0.056 (0.093)	-2.81	0.743*** (0.062)	2.940*** (1.043)	2.15	0.902
15y yield	0.282** (0.140)	-0.027 (0.101)	-2.74	0.696*** (0.070)	2.962*** (1.128)	2.06	0.871

Independent variables							Adj. R ²
Dep. Var.	Maturity-weighted debt/GDP			2y yield			
	Pre-ELB	ELB	Break <i>t-stat</i>	Pre-ELB	ELB	Break <i>t-stat</i>	
5y yield	0.126* (0.067)	-0.044 (0.051)	-2.86	0.906*** (0.032)	1.818*** (0.205)	4.77	0.981
10y yield	0.207** (0.104)	-0.058 (0.076)	-2.80	0.800*** (0.050)	2.301*** (0.698)	4.32	0.943
15y yield	0.239** (0.118)	-0.013 (0.095)	-2.43	0.751*** (0.600)	2.218*** (0.463)	3.34	0.915

Notes: Each row in each table reports the estimates of a single regression, where the dependent variable is a longer-term Treasury yield, as indicated in the first column. Each regression uses two independent variables: either the weighted-average maturity of Treasury debt in public hands or the maturity-weighted Treasury-debt-to-GDP ratio and either the one- or two-year zero-coupon Treasury yield. In each regression, the coefficient on each variable is allowed to differ between the period when the ELB was not binding (prior to December 2008) and the period when it was binding (December 2008 through December 2015), with the break accomplished using interactive dummy variables. The samples begin in August 1971 for the 5- and 10-year maturities and in December 1971 for the 15-year maturity. Yield data are Gurkaynak et al. (2007) zero-coupon yields. Treasury debt variables are constructed from CRSP data, following Greenwood and Vayanos (2014). All data are monthly. Newey-West standard errors, using 36 lags, are reported in parentheses, and statistical significance at the 10% (*), 5% (**), and 1% (***) levels is indicated by asterisks. The *t* statistics, reported in italics, test the significance of the break in each of the two coefficients in each regression.

Table 2. Model parameters

	Bond supply				Short rate				Risk aversion
	T	κ_β	σ_β	ζ	μ	κ	σ	b	a
[1] Shadow-rate model	60	0.021	0.20	0.37	4.9%	0.019	0.77%	0.17%	0.15
[2] Affine model – base	60	0.021	0.20	0.37	4.9%	0.019	0.77%	$-\infty$	0.15
[3] Affine model – recalibrated	60	0.021	0.20	0.39	5.2%	0.019	0.70%	$-\infty$	0.17

Notes: The table shows the calibrated values of the parameters in the baseline two-factor shadow-rate model, as well as in two models with affine short-rate processes. In the first affine model, all parameters (except the ELB) are the same as in the shadow-rate model. In the second affine model, the parameters are recalibrated to match the same set of unconditional yield-curve moments that the shadow-rate model matches. Details of the calibration strategy are provided in the text.

Table 3. Conditional moments of yield curve in data vs. two-factor models

Short rate below 0.68%

	% of obs.	3m rate	Slopes (to 3m)			
			2Y	5Y	10Y	15Y
Conditional means						
Data	16%	0.2%	0.3%	1.3%	2.5%	3.1%
Shadow-rate model	15%	0.2%	0.4%	1.2%	2.4%	3.5%
Affine Model – base calibration	15%	-1.3%	0.7%	1.8%	3.4%	4.5%
Affine Model – recalibrated	10%	-0.9%	0.7%	1.8%	3.3%	4.5%
Conditional standard deviations						
Data		0.1%	0.3%	0.6%	0.8%	0.8%
Shadow-rate model		0.2%	0.3%	0.7%	1.1%	1.4%
Affine Model – base calibration		1.7%	0.3%	0.7%	1.3%	1.7%
Affine Model – recalibrated		1.5%	0.3%	0.7%	1.2%	1.5%

Short rate above 0.68%

	% of obs.	3m rate	Slopes (to 3m)			
			2Y	5Y	10Y	15Y
Conditional means						
Data	84%	6.1%	0.5%	0.9%	1.3%	1.5%
Shadow-rate model	85%	6.0%	0.3%	0.7%	1.3%	1.8%
Affine Model – base calibration	85%	6.0%	0.3%	0.7%	1.2%	1.7%
Affine Model – recalibrated	90%	5.9%	0.3%	0.7%	1.3%	1.8%
Conditional standard deviations						
Data		3.1%	0.9%	1.3%	1.6%	1.7%
Shadow-rate model		3.2%	0.3%	0.8%	1.5%	1.9%
Affine Model – base calibration		3.2%	0.3%	0.8%	1.5%	2.0%
Affine Model – recalibrated		3.0%	0.3%	0.8%	1.4%	1.8%

Notes: The table shows conditional moments of zero-coupon yields simulated from the shadow-rate and affine models, based on the model parameters shown in Table 2, together with the corresponding moments from the data. Model results are based on 1 million simulations of the state variables. Yield data are from the Gurkaynak et al. (2007) dataset and cover the period August 1971 – December 2015, except for the 15-year yields, which begin in December 1971.

Table 4. Model-implied relationships among long-term yields, shorter-term yields, and Treasury supply

A. Sensitivity of long-term yields to WAM

	Shadow rate	Sensitivity to WAM			Sensitivity to WAM, holding 1Y yield fixed			Sensitivity to WAM, holding 2Y yield fixed		
		5Y	10Y	15Y	5Y	10Y	15Y	5Y	10Y	15Y
<i>ELB</i>	8%	0.10	0.17	0.22	0.08	0.16	0.21	0.06	0.14	0.20
	4%	0.10	0.17	0.22	0.08	0.16	0.21	0.06	0.14	0.19
	2%	0.09	0.17	0.21	0.08	0.15	0.20	0.06	0.14	0.19
	1%	0.09	0.16	0.21	0.08	0.15	0.20	0.06	0.13	0.19
	0%	0.08	0.15	0.20	0.06	0.14	0.19	0.05	0.12	0.18
	-1%	0.07	0.14	0.19	0.05	0.12	0.18	0.04	0.11	0.16
	-2%	0.05	0.13	0.18	0.04	0.11	0.16	0.03	0.10	0.15
	-4%	0.03	0.10	0.15	0.02	0.08	0.13	0.02	0.07	0.13

Notes: The first group of columns reports the model-implied sensitivity of 5-, 10-, and 15-year yields to the weighted-average maturity of Treasury debt at various values of the shadow rate. The second and third groups of columns report these sensitivities, holding fixed the level of the one-year or the two-year yield. The calculations, which are given in the text, rely on the derivatives of yields with respect to the shadow-rate and the bond-supply factors ($a_{\hat{r}}^{(\tau)}$ and $a_{\hat{\beta}}^{(\tau)}$), evaluated under the baseline parameter values (line 1 of Table 2).

B. Sensitivity of long-term yields to shorter-term yields

	Shadow rate	Sensitivity to 1Y			Sensitivity to 2Y		
		5Y	10Y	15Y	5Y	10Y	15Y
<i>ELB</i>	8%	0.9	0.7	0.6	0.9	0.7	0.6
	4%	0.9	0.7	0.6	0.9	0.7	0.6
	2%	0.9	0.7	0.6	0.9	0.8	0.6
	1%	0.9	0.8	0.7	0.9	0.8	0.7
	0%	1.2	1.2	1.0	1.1	1.0	0.9
	-1%	2.3	2.5	2.3	1.4	1.5	1.4
	-2%	5.8	7.1	7.0	2.0	2.5	2.4
	-4%	78.7	135.1	146.6	5.5	9.4	10.2

Notes: The table reports the model-implied sensitivity of 5-, 10-, and 15-year yields to 1- and 2-year yields, holding the bond-supply factor fixed, at various values of the shadow rate. The coefficients are calculated as the ratio $a_{\hat{r}}^{(\tau)}/a_{\hat{r}}^{(4)}$ or $a_{\hat{r}}^{(\tau)}/a_{\hat{r}}^{(8)}$ where $a_{\hat{r}}^{(\tau)}$ is the derivative of the yield at maturity τ with respect to \hat{r} under the baseline parameter values (line 1 of Table 2).

Table 5. Decompositions of yield responses to unconventional policy shocks (zero correlation)

A. Contemporaneous responses (bps)

Maturity [1]	Shadow-rate shocks		Fed balance-sheet shocks	Interaction [5]	Total [6]
	Expectations component [2]	Term premium component [3]	Term premium component [4]		
2 years	-59 (-82, -39)	-22 (-25, -16)	-13 (-14, -12)	7 (5, 8)	-90 (-116, -63)
5 years	-90 (-106, -69)	-51 (-52, -47)	-30 (-31, -26)	12 (9, 14)	-160 (-177, -135)
10 years	-102 (-109, 91)	-70 (-76, -62)	-47 (-50, -41)	12 (8, 16)	-207 (-211, -199)
15 years	-98 (-100, -92)	-72 (-82, -63)	-57 (-60, -49)	10 (7, 14)	-215 (-219, -210)

B. Total variance (bps²/100)

Maturity [1]	Shadow-rate shocks		Fed balance-sheet shocks	Interaction [5]	Total [6]
	Expectations component [2]	Term premium component [3]	Term premium component [4]		
2 years	19 (13, 26)	5 (2, 9)	0.0 (0.0, 0.2)	0.2 (0, 0.3)	24 (15, 34)
5 years	18 (10, 30)	13 (5, 24)	0.2 (0.2, 0.4)	1.0 (-0.8, 2.3)	32 (17, 54)
10 years	16 (8, 29)	19 (8, 40)	0.7 (0.5, 0.9)	2.4 (-1.8, 5.3)	39 (19, 71)
15 years	14 (6, 25)	19 (8, 41)	1.0 (0.7, 1.3)	3.1 (-2.0, 6.8)	37 (17, 70)

Notes: The table summarizes the results of simulations of unconventional monetary policy in the shadow-rate model. Panel A reports the cumulative response of the spot zero-coupon yield curve in model simulations based on the distribution of state-variable trajectories shown in Figure 6, summing the responses to the shocks in each period. Panel B reports the total variance in yields in the simulations, relative to a baseline scenario in which no shocks occur. In both cases, for each maturity, the median response is reported, with the 5% and 95% quantiles in parentheses below. The total effect on the yield of each maturity is shown in the last column. The bond-supply and shadow-rate shocks are simulated both separately and together to obtain the decomposition reported in the other columns. The “interaction” column represents the effect of nonlinearities that cause the sum of the two individual simulations to differ from that of the joint simulation. For the shadow-rate shocks, the change in the expectations component is calculated from equation (13), while the change in the term-premium component is calculated as the difference between the total change in yields and the change in the expectations component. By construction, the individual components sum to the totals in each simulation, but the median values in columns [2] through [5] may not sum to the values in column [6] because of the asymmetry in the distributions across simulations.

Table 6. Decompositions of yield responses to unconventional policy shocks (perfect correlation)

A. Contemporaneous responses (bps)

Maturity [1]	Shadow-rate shocks		Fed balance-sheet shocks	Interaction [5]	Total [6]
	Expectations component [2]	Term premium component [3]	Term premium component [4]		
2 years	-58 (-80, -36)	-22 (-25, -15)	-13 (-13, -12)	6 (5, 7)	-86 (-111, -57)
5 years	-87 (-101, -68)	-49 (-50, -46)	-27 (-28, -25)	9 (7, 11)	-154 (-169, -128)
10 years	-96 (-97, -95)	-66 (-68, -64)	-41 (-43, -38)	9 (7, 11)	-196 (-199, -189)
15 years	-93 (-95, -91)	-68 (-75, -63)	-48 (-51, -45)	7 (6, 10)	-201 (-206, -196)

B. Total variance (bps²/100)

Maturity [1]	Shadow-rate shocks		Fed balance-sheet shocks	Interaction [5]	Total [6]
	Expectations component [2]	Term premium component [3]	Term premium component [4]		
2 years	18 (12, 25)	5 (2, 9)	0.0 (0.0, 0.1)	0.5 (0.1, 0.7)	24 (19, 30)
5 years	17 (9, 29)	13 (5, 23)	0.3 (0.2, 0.4)	0.8 (-0.5, 1.9)	32 (15, 54)
10 years	16 (8, 27)	19 (8, 38)	0.7 (0.6, 0.8)	1.4 (-2.1, 3.5)	37 (18, 68)
15 years	14 (6, 24)	18 (9, 39)	0.9 (0.8, 1.2)	1.7 (-2.7, 4.2)	35 (16, 64)

Notes: The table summarizes the results of simulations of unconventional monetary policy in the shadow-rate model, under the assumption that the realized shocks to the shadow rate and the Fed balance sheet were perfectly correlated. Panel A reports the cumulative response of the spot zero-coupon yield curve in model simulations based on the distribution of state-variable trajectories described in the text, summing the responses to the shocks in each period. Panel B reports the total variance in yields in the simulations, relative to a baseline scenario in which no shocks occur. In both cases, for each maturity, the median response is reported, with the 5% and 95% quantiles in parentheses below. The total effect on the yield of each maturity is shown in the last column. The bond-supply and shadow-rate shocks are simulated both separately and together to obtain the decomposition reported in the other columns. The “interaction” column represents the effect of nonlinearities that cause the sum of the two individual simulations to differ from that of the joint simulation. For the shadow-rate shocks, the change in the expectations component is calculated from equation (13), while the change in the term-premium component is calculated as the difference between the total change in yields and the change in the expectations component. By construction, the individual components sum to the totals in each simulation, but the median values in columns [2] through [5] may not sum to the values in column [6] because of the asymmetry in the distributions across simulations.

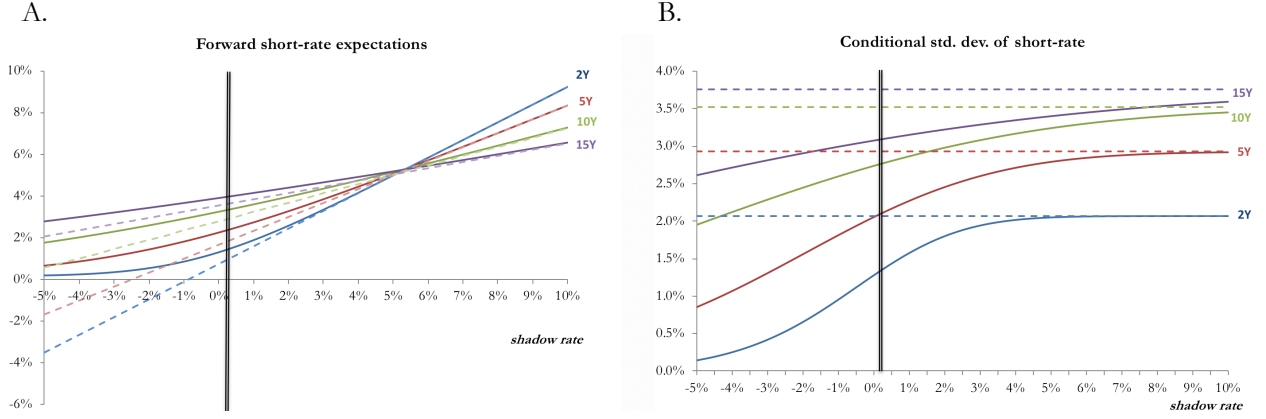


Figure 1. Conditional moments of short rate at various horizons. The figure shows the conditional mean (panel A) and standard deviation (panel B) of the time $t+s$ short rate, conditional on the value of the shadow rate in time t , where $s = 2, 5, 10$, and 15 years. The solid lines show these moments in the shadow-rate model, using the parameters shown in the top line of Table 2. The dashed lines show the moments in an affine model with the same parameter values but with the ELB removed.


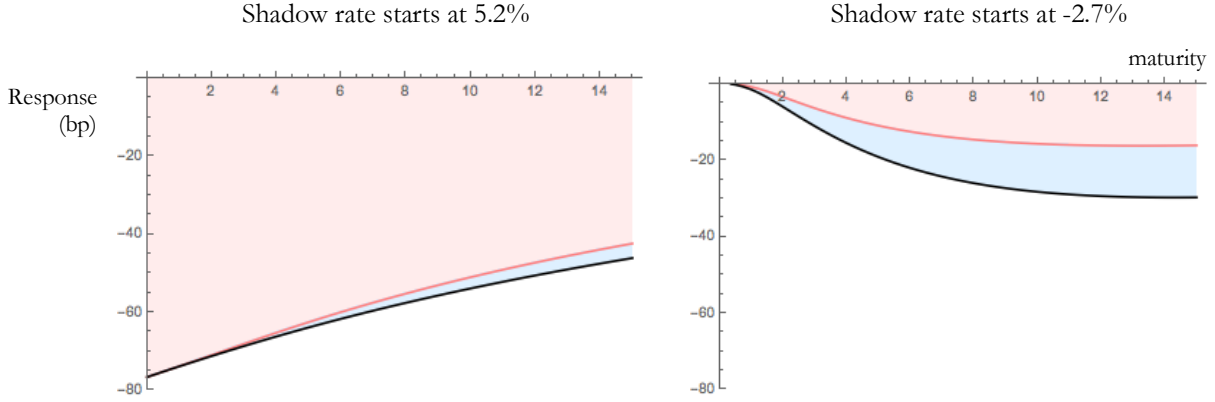


Figure 2. Comparative statics of bond supply in the one-factor model. The figure shows the yield curves generated by the one-factor model at two different values of the shadow rate \hat{r}_t and using two different values of the parameter ζ , which represents the total quantity of bonds in the economy. The baseline value of $\zeta = 0.42$ (shown in black) is calibrated to allow the model to match the average value of the ten-year yield. The value of $\zeta = 0.21$ (shown in red) considers a hypothetical 50% reduction in bond supply relative to the baseline case. The left-hand panel compares the two yield curves resulting from these parameter values when the shadow rate is at the mean value of the short rate (5.2%), while the right-hand panel shows the curves when the shadow rate is at -2.7%, its average during the ELB period according to the Krippner (2012) estimates. The model is solved numerically using the other parameters values described in the text.

A. Spot yield curve



B. Forward-rate curve

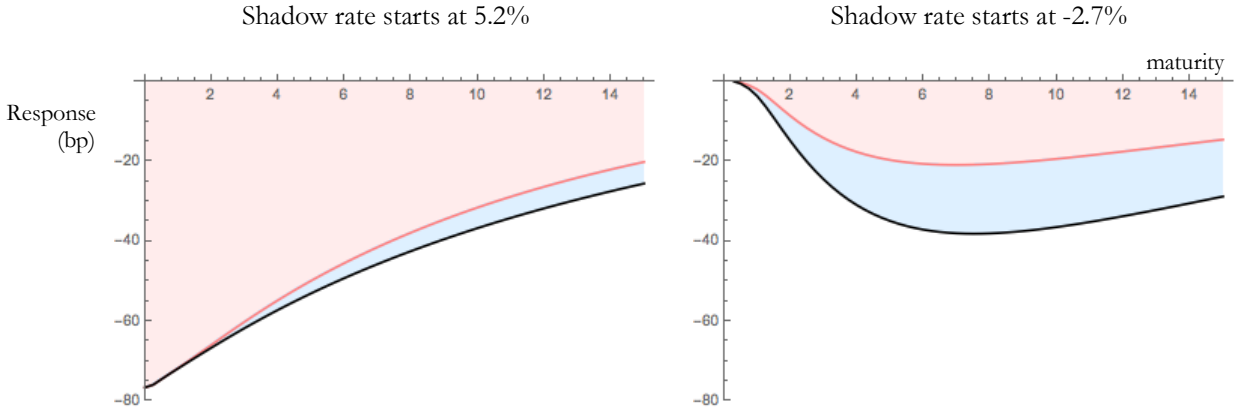


Figure 3. Response to shadow-rate shocks in the one-factor model. The figure shows the model-implied response of the yield curve (panel A) and the forward-rate curve (panel B) to one-standard-deviation shock to the shadow rate in the period when the shock occurs. The pink region shows the change in the expectations component of yields, while the blue region shows the change in the term premium. The change in the expectations component is calculated from equation (13), while the change in the term-premium component is calculated as the difference between the total change in yields and the change in the expectations component. Responses are evaluated starting both from a shadow rate at the mean value of the short rate (5.2%) and a value of -2.7%, its average during the ELB period according to the Krippner (2012) estimates. The model is solved numerical using the parameter values described in the text.

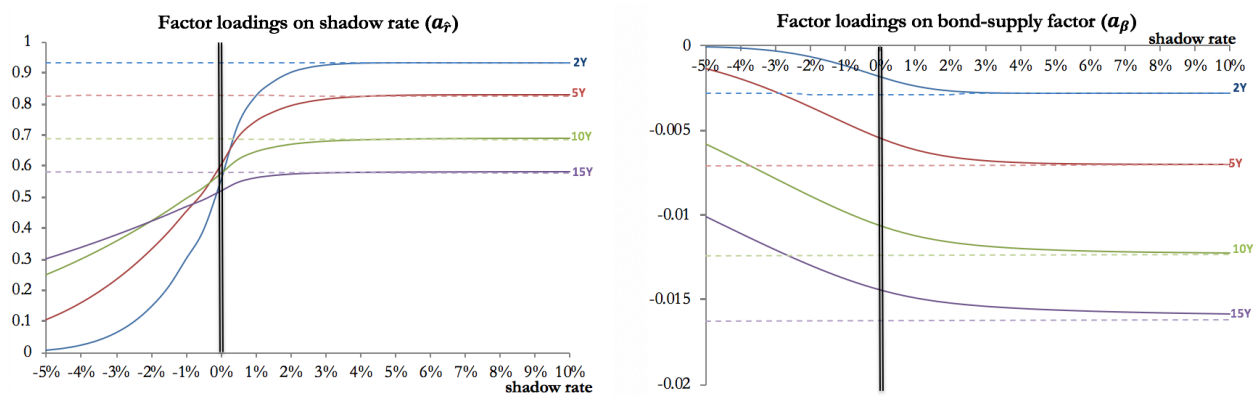
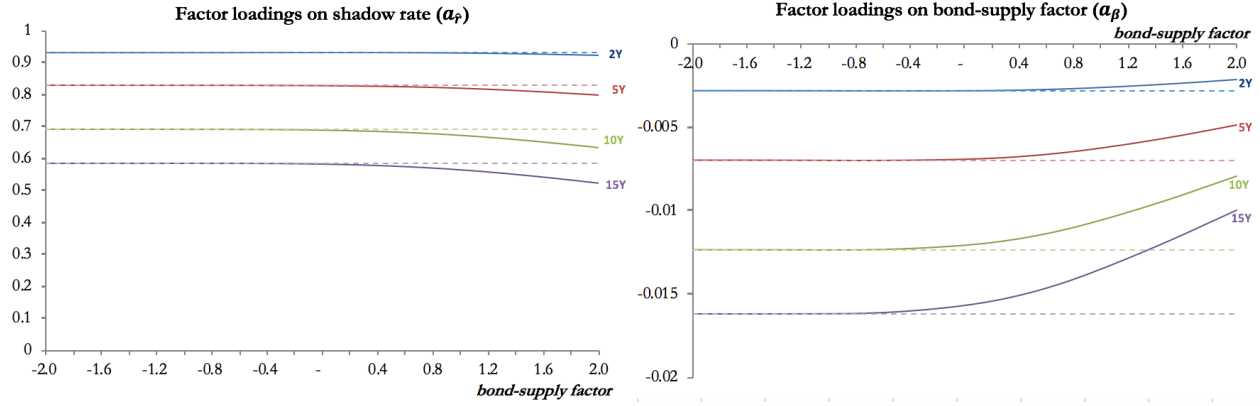


Figure 4. Factor loadings across values of the shadow rate at $\beta = 0$. The figure shows the model-implied factor loadings for the τ -period yield, conditional on the time- t value of the shadow rate, where $\tau = 2, 5, 10$, and 15 years. The solid lines show the loadings in the shadow-rate model, under the parameters shown in line 1 of Table 2. The dashed lines show the loadings in an affine model with the same parameter values but with the ELB removed (line 2 of Table 2). The bond-supply factor β is held fixed at its mean of zero.

A. $\hat{r}_t = 5.2\%$



B. $\hat{r}_t = -2.7\%$





Figure 5. Factor loadings across values of the supply factor. The figure shows the model-implied factor loadings for the τ -period yield, across time- t values of the bond-supply factor, where $\tau = 2, 5, 10$, and 15 years. The solid lines show the loadings in the shadow-rate model, under the parameters shown in line 1 of Table 2. The dashed lines show the loadings in an affine model with the same parameter values but with the ELB removed (line 2 of Table 2). In panel A, the shadow rate is held fixed at the mean value of the short rate (5.2%), while in panel B it is held fixed at -2.7%, its average value during the ELB period according to the Krippner (2012) estimates.

A. Shadow rate



B. %Change in 10-year equivalents





Figure 6. Distributions of state-variable trajectories in model simulations. The figures show the distributions of 100,000 simulated paths of the state variables during the ELB period, in terms of the pointwise medians (solid lines) and 5% and 95% quantiles (dotted lines). The simulations are constructed to exactly match the observed values of the short-term interest rate and the percentage of ten-year equivalent bonds held by the Fed as of December 2008 and December 2015, with the intermediate values simulated from the model as described in Appendix C. For the purposes of presentation, in panel B the balance-sheet state variable (Q) is converted to a cumulative percentage change in 10-year equivalent bonds held by investors, relative to the amount that would have been outstanding in the absence of shocks, using equation (27).

A. Spot yield curve



B. Forward-rate curve


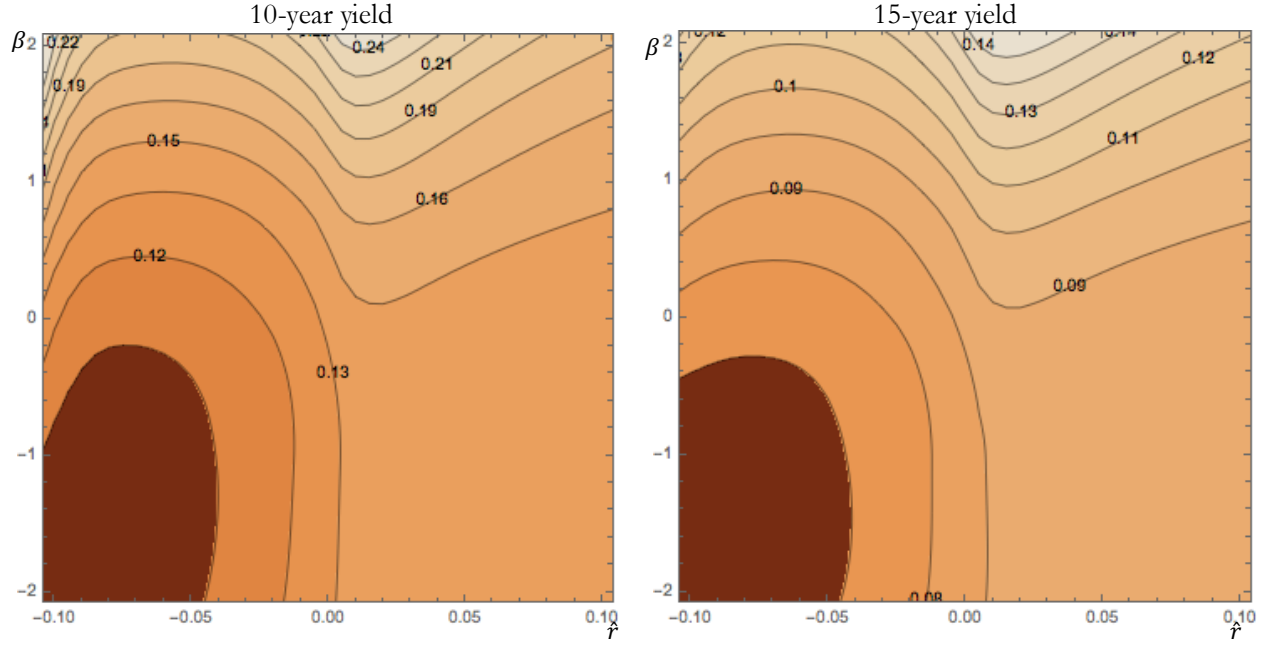


Figure 7. Cumulative yield-curve responses in model simulations. The figure shows the cumulative response of bond yields (panel A) and forward rates (panel B) in model simulations based on the distribution of state-variable trajectories shown in Figure 6. The figures sum the contemporaneous responses to the shocks in each of the 28 simulated periods, without accounting for dynamics. In each panel, the solid line shows the pointwise median and dashed lines show 5% and 95% quantiles.

A. Bond-supply factor (β_t)



B. Fed balance-sheet factor (Q_t)

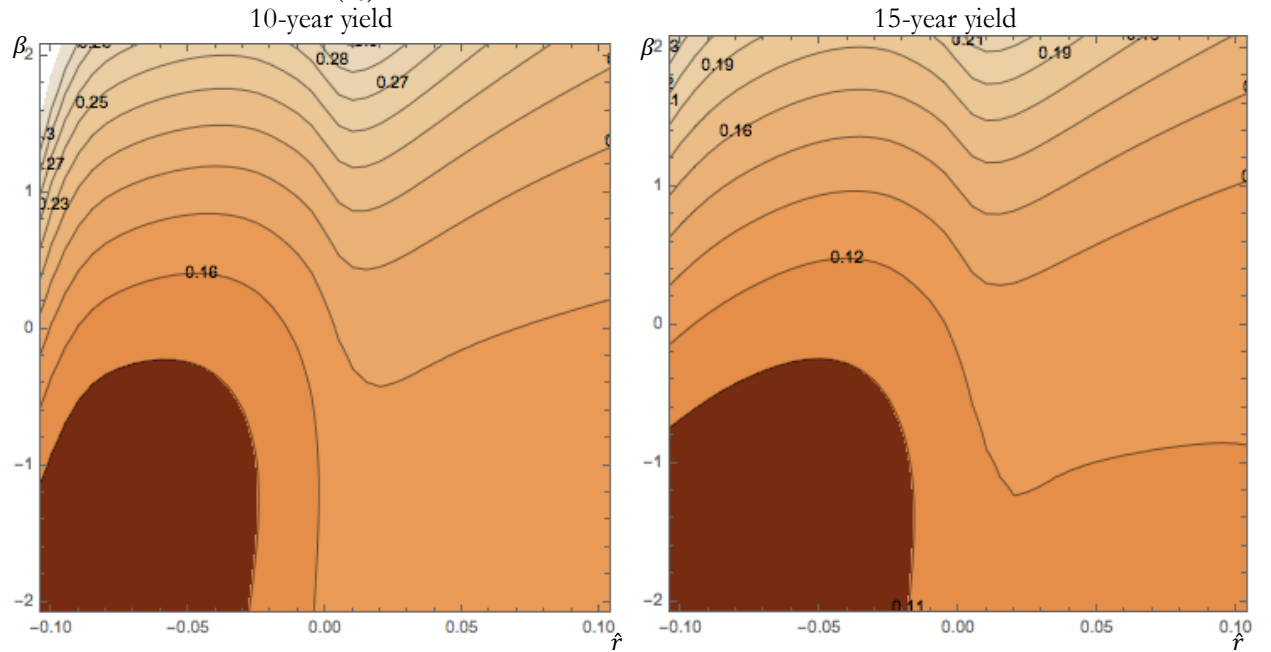


Figure 8. Relative efficacy of bond-supply shocks across state values. The figure shows contour maps of the effects of bond-supply shocks on 10- and 15-year yields, relative to the effects of shadow-rate shocks, in the baseline two-factor shadow-rate model. Relative efficacy is calculated, for each yield, as the size of the bond-supply shock that would be necessary to equal the effects of a -25-basis-point shock to the shadow rate. The values of this ratio are shown across different regions of the state space, with darker coloring indicating regions where the bond-supply shocks are relatively more powerful.