Expectation and Duration at the Effective Lower Bound

Thomas B. King[*]

March 15, 2017

Abstract

I study equilibrium bond pricing in a model with risk-averse arbitrageurs and an effective lower bound (ELB) on nominal rates. The model exposes nonlinear interactions among short-rate expectations, bond supply, and term premia that are absent from affine models, and these features help it replicate the observed behavior of the yield curve near the ELB, including event-study evidence on the effects of unconventional monetary policy. The pricing impact of both short-rate expectations and bond supply are attenuated at the ELB. However, under a calibration that matches the data, the model implies that shocks to the anticipated path of short-term interest rates were more important for long-term yields than shocks to investors' interest-rate risk exposures during the ELB period in the U.S..

1 Introduction

A growing body of research emphasizes that financial intermediaries are often the marginal investors in asset markets and that, consequently, asset prices may depend on how the portfolios of those intermediaries are structured (e.g., He and Krisnamurthy (2013), Adrian, Etula, and Muir (2014), He, Kelly, and Manela (2016)). For the default-free bond market, Vayanos and Vila (2009) and Greenwood and Vayanos (2014) present models in which the marginal investors are risk-averse arbitrageurs. When the arbitrageurs are given more long-term bonds to hold, the duration risk of their balance sheets rises and they demand higher risk premia. Hamilton and Wu (2012), Kaminska and Zinna (2014), King (2014), Altavilla, Carboni, and Motto (2015), Haddad and Sraer (2015), and Hayashi (2016), among others, have also studied variations of this framework. Although these models have been useful for formalizing and understanding the so-called "duration channel" of bond supply, one limitation is that they have almost exclusively been developed under an assumption of linearity.

This paper studies how the duration channel operates in a world with one prominent *non*linearity—an occasionally binding effective lower bound (ELB) on the short-term interest rate. From a practical perspective this is a particularly germane extension because central banks that have sought to exploit the duration channel through long-term bond purchases have typically done so only after they have cut traditional policy rates as far as possible. If the effects of the ELB were confined to relatively minor deviations from linearity or localized to very short-term interest rates, its presence might not be quantitatively important, and the models mentioned above might still do a good job of describing the relationships between monetary policy, Treasury supply, and bond yields. Alas, there are both theoretical and empirical reasons to believe that the ELB has first-order effects on those relationships.

As an illustration, the top panel of Table 1 reports regressions of long-term yields of different maturities on the ratio of maturity-weighted Treasury debt to GDP over the period 1971 - 2015. The regressions also include the one-year Treasury yield to control for the level and near-term expectation of short-term interest rates. The specification and data are the same as in Greenwood and Vayanos (2014), except that I allow the coefficients to differ after December 2008, when short rates first hit the ELB. (The Greenwood-Vayanos sample ended in 2007.) The remaining panels of the table check robustness by using the weighted-average maturity of outstanding Treasury debt in place of the debt-to-GDP ratio and the two-year Treasury rate in place of the

one-year.

Column 2 replicates the main Greenwood and Vayanos (2014) result: long-term yields were significantly positively related to Treasury supply in the pre-ELB period. A one-percentage-point increase in the maturity-weighted debt-to-GDP ratio increased the 10-year yield by 26.5 basis points in this sample, and, consistent with longer-term bonds having greater duration exposure, the coefficients were monotonically higher for longer maturities. But, as column 3 shows, those coefficients all fall to near zero and lose statistical significance during the ELB period. The t statistics reported in column 4 show that the declines in the coefficient values are significant at the 1% level.[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 short-term yield were in the range of 0.7 to 0.9 and were monotonically decreasing in the maturity of the dependent variable. At the ELB, the coefficient magnitudes increase by a factor of 2 to 6, and the pattern of coefficients across maturities reverses, with the 15-year yield now being more sensitive than the 5-year yield is. Again, the break statistics show that these differences are statistically significant.

These results strongly suggest both qualitative and quantitative changes in the behavior of the yield curve and its relationship to Treasury supply during the ELB period. Affine models—which, by definition, have constant factor loadings—are essentially incapable of generating such outcomes. This raises the possibility that any conclusions derived from such models about the effects of various types of shocks, including alternative monetary policies, may not be robust to the presence of the ELB.

The remainder of this paper studies effects of the ELB in a structural equilibrium model of bond pricing in the tradition of Vayanos and Vila (2009). I generalize the standard affine process for the short-term interest rate in such models to a "shadow rate" process, following successful use of that device in the recent empirical term-structure literature, such as Kim and Singleton (2012) and Krippner (2012).[2] Away from the ELB, a shadow-rate shock behaves just like a standard shock to the level of the short rate. At the ELB, it corresponds to changes in the length of time that the ELB is expected to bind. Anything that alters investors' beliefs about the persistence

---

[1] I calculate standard errors using the Newey-West (1994) procedure, with 36 lags, again following Greenwood and Vayanos.

[2] Bauer 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 atheoetical term structure models. This paper is among the first to incorporate a shadow-rate process into a structural model of the yield curve.

of the ELB can be modeled as a shadow-rate shock. This includes explicit forward guidance about the short rate, such as that articulated by the FOMC with various degrees of specificity throughout the 2008 – 2015 period. It may also include the "signaling channel" of asset purchases, through which expansions of the central bank's balance sheet might be viewed as committing it to keeping rates near zero for a long time.[3]

When parameterized to match the unconditional moments of Treasury yields since 1971, the model delivers quantitatively accurate results in three verifiable respects. First, it matches the *conditional* moments of the data well. Most importantly, it does a good job of replicating the empirical behavior of the yield curve when the short rate is near the ELB. (In contrast, a comparable affine model makes predictions about yields that differ wildly from what we have actually observed near the ELB.) Second, it reproduces the patterns reported in Table 1. The effects of shocks to Treasury supply become weaker at the ELB, and longer rates become "more sensitive" to shorter rates, with the coefficients switching from declining to increasing across maturities. Third, when subjected to shocks that approximate the Federal Reserve's unconventional monetary policy over the ELB period—a net decrease of about 35% in the duration-weighted quantity of government bonds outstanding and shadow-rate shocks that kept the short rate at the ELB for exactly seven years—the model reproduces the results of empirical event studies. In particular, it matches empirical estimates of (1) the cumulative yield-curve impact of unconventional-policy announcements (about -200 basis points on the ten-year yield), (2) the hump shape of the reaction in forward rates across maturities to these announcements, and (3) the approximate relative responses of the expectations and term-premium components of yields.

Key to achieving these successes are the nonlinearities that the ELB induces throughout the term structure. Those nonlinearities manifest themselves in a variety of subtle ways and can be relevant even when the short rate is well above zero, but they are, naturally, most pronounced in states of the world in which the ELB is binding. Three novel impications are particularly important. First, at the ELB, changes in future short-rate expectations have their largest effects on yields at intermediate maturities, in contrast to the monotonic effects of short-rate shocks in simple affine models. This

---

[3]Woodford (2012), Bauer and Rudebusch (2014), and Bhattarai, Eggertsson, and Gafarov (2015) argue for the importance of the signaling channel. As noted by Swanson (2016) and discussed further below, 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 the effects of the signalling channel and those of forward guidance.

phenomenon largely explains the empirical observation that policy announcements at the ELB move forward rates the most at the five- to seven-year horizon. Second, as noted above, the effects of bond-supply shocks on term premia are damped at the ELB. The reason is that the magnitude of the duration channel varies directly with interest-rate volatility, and volatility is lower when the short-rate distribution is truncated. Third, in the presence of the ELB, changes in the anticipated path of the short rate have effects on term premia. Again, the intuition is that interest-rate volatility moves together with short-rate expectations at the ELB, and term premia are proportional to volatility *ceteris paribus*. The impact of short-rate expectations on term premia has received relatively little attention in the literature (it does not exist in an affine model), but my results suggest that it was of comparable importance to the duration channel of asset purchases during the ELB period.

I use the model to address two quantitative questions. First, I ask through which channel unconventional monetary-policy shocks had their largest effects. The answer is that changes in the anticipated path of the short rate were responsible for the majority of the downward shift in the yield curve. Bond-supply shocks account for only about 1/4 of the total decline in the ten-year yield, and at shorter maturities their relative importance is even smaller. In other words, the model suggests that the duration channel of asset purchases was considerably less significant than implicit or explicit Fed communications about the course of the conventional policy instrument. About a third of the total effects of shadow-rate shocks operate through the term premium.

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. The efficacy of both types of shocks is attenuated in such a situation because of the damping effects associated with the ELB, but this attenuation is greater for the shadow-rate shocks. A negative shadow rate and a high supply 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 other in the recent literature. As noted above, a number previous studies have used variants of the Vayanos-Vila framework to analyze the effects of fluctuations in bond supply in a linear environment. Hamilton and Wu's

(2012) analysis included a version in which, once the short rate has reached the ELB, investors believe that it will 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, Hanson, and Vayanos (2015) noted the hump-shaped pattern in forward rates in response to asset-purchase announcements and, through the lens of a linear model, argued that duration effects must have been responsible. Although I also obtain a hump-shaped forward-curve response to bond-supply shocks, I find that, when the quantities of assets actually purchased by the Fed are fed into the model, its magnitude is relatively small. Instead, my model implies that empirical hump in the forward curve primarily results from the non-monotonic effects of changes in short-rate expectations due to the ELB, a mechanism that is absent from the affine model of Greenwood, Hanson, and Vayanos.

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 (2012) and Krishnamurthy and Vissing-Jorgensen (2013) argue, based on event studies, that the evidence for the duration channel is weak, consistent with the results presented here. Although they also argue that short-rate expectations changed by a smaller amount than my model suggests, that is likely in part because their sample excludes most of the important instances of explicit forward guidance by the Fed. Swanson (2016) 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, he notes that the decomposition relies on coefficients estimated during the pre-ELB period and may not be robust if factor loadings differ at the ELB. My model explores the potential for endogdous changes in loadings at the ELB and implies that the changes could be quite substantial.[4]

---

[4]D'Amico, English, Lopez-Salido, and Nelson (2012) and Cahill, D'Amico, Li, and Sears (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. Neither of those studies separately identifies the magnitude of the signaling channel.

2 Theoretical Framework

2.1 Investor behavior and equilibrium yields

I consider investors solving a standard Markowitz portfolio-choice problem. This setup is simple enough to capture the fundamental mechanisms at work in a transparent way, and the results in King (2014) suggest that modifying it with more-realistic features does not do much to alter its quantitative outcomes. Furthermore, this is the same optimization problem that forms the basis of the models in Vayanos and Vila (2009) and the several theoretical papers that have followed it, facilitating comparison to previous literature.

Investors have access to a continuum of zero-coupon bonds with maturities of 0 to T. At each point in time t, they choose to hold a market-value quantity $x_t(\tau)$ at each maturity τ. Let P_t^τ represent the time-t price of a bond with τ periods of remaining maturity. 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

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

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

$$\max_{x_t(\tau) \forall \tau} \ \mathrm{E}_t[dW_t] - \frac{a}{2} \mathrm{var}_t[dW_t] \tag{2}$$

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

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

$$E_t\left[dp_t^{(\tau)}\right] = r_t + a \int_0^T x_t(\tau') \mathrm{cov}_t\left[dp_t^{(\tau)}, dp_t^{(\tau')}\right] d\tau' \tag{3}$$

for all τ, where $p_t^{(\tau)}$ is the log price of the τ-maturity bond and cov_t denotes the covariance conditional on the time-t state. 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 wealth.

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

$$s_t(\tau) = x_t(\tau) \tag{4}$$

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. Solving the model is thus tantamount to solving for the conditional expectations and covariances in equation (3).

The exogenous state variables in the model are r_t and $s_t(\tau)$. I assume that the short rate and the par value of debt outstanding $\left(s_t(\tau)/P_t^{(\tau)}\right)$ are constant "within" periods. That is, they jump discretely at regular intervals, normalized to unit length. This discretization will be necessary for the numerical solution, but it can also be justified by the observation that monetary policy and debt issuance do not, in reality, adjust in infinitesimal increments in continuous time but rather move by sizeable amounts following periodic policy decisions. In any case, the discretization makes little quantitative difference, and it becomes irrelevant as the length of the time interval goes to zero. The discretization implies that $\frac{dP_h^{(\tau)}}{P_h^{(\tau)}} = p_{t+1}^{(\tau-1)} - p_t^{(\tau)}$ for all $h \in [t, t+1)$.

Since the τ-period bond yield has the usual relationship to prices, $y_t^{(\tau)} = -p_t^{(\tau)}/\tau$, it is straightforward to show that yields are given by:

$$y_t^{(\tau)} = \frac{1}{\tau}\sum_{h=0}^{\tau} \mathrm{E}_t\left[r_{t+h}\right] + a\frac{1}{\tau}\mathrm{E}_t\left\{\sum_{h=0}^{\tau-1}\int_0^T \tau' s_{t+h}(\tau')\,\mathrm{cov}_{t+h}\left[y_{t+h+1}^{(\tau-h-1)}, y_{t+h+1}^{(\tau')}\right]d\tau'\right\} \tag{5}$$

It will also be instructive to examine forward rates. The one-period forward rate τ periods ahead is given by

$$f_t^{(1,\tau)} \equiv \tau y_t^{(\tau)} - (\tau-1)y_t^{(\tau-1)} \tag{6}$$

The first term on the left-hand side of (5) is the expectations component of yields; the second term is the term premium. The basic intuition for how asset purchases (or other fluctuations in bond supply) affect yields in this model is that they change the

weights on the covariance terms in the term premium. A shock that shifts $s_t(\tau)$ toward lower-covariance assets—typically, those with shorter duration—will reduce yields through that term in period t. Note that today's term premium depends not just on today's bond supply s_t, but also on the expected future values of supply s_{t+h}. Thus, if bond-supply shocks are persistent, they will also affect the expected value of the integral term in subsequent periods, leading to a further reduction in time-t yields.[5]

2.2 The short rate

I assume that r_t follows the "shadow rate" process

$$r_t = \max[\widehat{r}_t, b] \qquad (7)$$

where b is the lower bound on the nominal short rate and

$$\widehat{r}_t = \mu_{\widehat{r}}(1 - \phi_{\widehat{r}}) + \phi_{\widehat{r}}\widehat{r}_{t-1} + e_t^{\widehat{r}} \qquad e_t^r \sim Niid(0, \sigma_{\widehat{r}}) \qquad (8)$$

for some parameters $\mu_{\widehat{r}}$, $\phi_{\widehat{r}}$, and $\sigma_{\widehat{r}}$. This is essentially the process used in the empirical shadow-rate literature mentioned in the introduction (although in those studies $\widehat{r}_t$ can sometimes depend on multiple factors). 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.[6] 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.

Since the shadow rate follows a Gaussian AR(1) process, the conditional distribution of future shadow rates h periods ahead is normal, with mean and variance given by the standard prediction equations

---

[5]Also note that the term premium at a given maturity increasing in risk aversion. Risk neutrality implies the expectations hypothesis.

[6]Of course, one can imagine alternative ways of modeling the behavior of short rates that also respect the ELB. For example, Monfort, Pegoraro, Renne, and Roussellet (2015) suggest a model in which, once the ELB is reached, the short rate stays there with some probability. At least qualitatively, the exact specification of the short rate at the ELB is of only of minor importance. 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 is reached. Any model that generates these properties (which are amply evident in the data) will produce results along the lines of those presented below, although of course exact magnitudes will differ with the specification.

$$\mathrm{E}_t\left[\widehat{r}_{t+h}\right] = \mu_{\widehat{r}}(1 - \phi_{\widehat{r}}^h) - \phi_{\widehat{r}}^h \widehat{r}_t \tag{9}$$

$$\mathrm{var}\left[\widehat{r}_{t+h}\right] = \sigma_{\widehat{r}}^2 \sum_{j=1}^{h} \phi_{\widehat{r}}^{2(j-1)} \tag{10}$$

In an affine model, since $\widehat{r}_t = r_t$ in all states of the world, these equations also describe the conditional distribution of future short rates. To begin to get a sense of why the shadow-rate model delivers qualitatively different results than the affine model, note that, once the ELB is imposed, the conditional distribution of the short rate at any period in the future becomes *truncated* normal. Consequently, the mean and variance of r_{t+h}, conditional on information at time t, are given by

$$\mathrm{E}_t\left[r_{t+h}\right] = \mathrm{E}_t\left[\widehat{r}_{t+h}\right] + L_t^{(h)} \sqrt{\mathrm{var}\left[\widehat{r}_{t+h}\right]} \tag{11}$$

$$\mathrm{var}_t\left[r_{t+h}\right] = \mathrm{var}_t\left[\widehat{r}_{t+h}\right] \left[1 - \left(\frac{\mathrm{E}_t\left[\widehat{r}_{t+h}\right] - b}{\sqrt{\mathrm{var}_t\left[\widehat{r}_{t+h}\right]}}\right) L_t^{(h)} - \left(L_t^{(h)}\right)^2\right] \tag{12}$$

where

$$L_t^{(h)} \equiv \frac{\varphi\left(\frac{b - \mathrm{E}_t[\widehat{r}_{t+h}]}{\sqrt{\mathrm{var}_t[\widehat{r}_{t+h}]}}\right)}{1 - \Phi\left(\frac{b - \mathrm{E}_t[\widehat{r}_{t+h}]}{\sqrt{\mathrm{var}_t[\widehat{r}_{t+h}]}}\right)} \tag{13}$$

with $\varphi\left(.\right)$ and $\Phi\left(.\right)$ denoting the standard normal PDF and CDF, respectively. The term $L_t^{(h)}$, which is strictly positive if b is finite, represents the effect of the ELB, relative to the affine case. In general, this term is larger for lower values of $\widehat{r}_t$, implying that short-rate expectations are pushed up by the proximity of the ELB, and short-rate volatilities are damped.

Figure 1 depicts these conditional moments of the forward short rate, across different values of $\widehat{r}_t$, comparing a model in which $b = 0.0017$ (solid lines) to one in which $b = -\infty$ (dashed lines). (The vertical lines indicate the location of the ELB in the former model.) All other parameters of the short rate processes are the same for both models.[7] The forward moments are shown for horizons of 2, 5, 10, and 20 years.

As shown in Panel A, in the model with $b = -\infty$, the forward expected short rate $\mathrm{E}_t[r_{t+h}]$ is an affine function of $\widehat{r}_t$, with the slope of that function decreasing in τ. The

---

[7]They are the calibrated values shown in the top row of Table 1 in Section 3; the details of the calibration are discussed there.

constrained model approaches these functions as $\widehat{r}_t$ moves far above b. This is intuitive, since, as the ELB gets farther away, it should have less influence on asset prices and the model should behave approximately linearly. However, near and below the ELB, the two models behave much differently. The affine model attaches no special importance to $\widehat{r}_t = b$, and it thus, for example, permits arbitrarily negative rates. Instead, in the constrained model, expected future short rates asymptote to the ELB as $\widehat{r}_t \to -\infty$. Because of this, their slopes with respect to $\widehat{r}_t$ decrease and eventually go to zero. In other words, when $\widehat{r}_t < b$, shocks to $\widehat{r}_t$ have smaller effects on expected future short rates, particularly at relatively short horizons.

Similarly, in the model without the ELB, the conditional variance of the short rate is a constant value for all values of $\widehat{r}_t$ at any given horizons. This is depicted by the horizontal dashed lines in Panel B. Again, in the constrained model, the conditional short-rate variances approach these values as $\widehat{r}_t$ gets much larger than b. However, as $\widehat{r}_t$ approaches b from above, and particularly as it falls below b, the conditional variances 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.

Returning to equation (5) with these observations in mind, we can see heuristically how the ELB will matter for the propagation of shocks. First, shocks to $\widehat{r}_t$ will have generally weaker effects on the expectations component of yields when $\widehat{r}_t < b$. 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 $\widehat{r}_t$ are always largest at the short end of the curve. Second, current and future short-rate volatilities are lower at the ELB, which will mean that the volatilities of *all* yields are lower, *ceteris paribus*. Thus the covariance terms that represent the multipliers on s_t in equation (5) will generally be smaller. This means that a given shift in the supply distribution will have a smaller effect on term premia at the ELB than it does away from the ELB (or in an affine model). Finally, for similar reasons, the covariance terms that determine term premia are increasing in the level of $\widehat{r}_t$. At and near the ELB there is a positive relationship between the expected future short rate and its variance, and, all else equal, this will translate into a positive relationship between short-rate expectations and term premia. Because of the constant conditional variance of r_t in the affine model, this channel does not exist there.

2.3 Bond Supply

The arguments just sketched for the qualitative effects of the ELB made no reference to the way in which asset supply $s_t(\tau)$ was determined. Indeed, they hold for a variety of possible processes for bond supply. Nonetheless, to obtain a quantitative assessment, we must specify a particular process.

Since bond supply is continuous across maturities, the object $s_t(\tau)$ is an infinite-dimensional vector. Clearly, it is desirable to reduce this dimension of to a manageable number of state variables. I follow previous literature (Vayanos and Vila, 2009, and subsequent papers that have adapted that model) and specify $s_t(\tau)$ as an affine function of a finite state vector β_t that follows a linear-Gaussian process:

$$s_t(\tau) = \zeta(\tau) + \theta(\tau)\beta_t \tag{14}$$

$$\beta_t = \phi_\beta \beta_{t-1} + e_t^\beta \qquad e_t^\beta \sim Niid(0, \sigma_\beta) \tag{15}$$

where $\zeta(\tau)$ and $\theta(\tau)$ are maturity-specific intercepts and coefficients. A major advantage of the affine specification for $s_t(\tau)$ in previous papers is that it allows for analytical tractability if the short-rate process is also affine. That property makes little difference here because the model will be solved numerically anyway, but it is still expositionally helpful to maintain the affine specification for supply. I have also compared specifications in which $s_t(\tau)$ follows a nonlinear process (see King, 2014, for example) and found results similar to those reported below.

I assume that the intercept is constant across maturities: $\zeta(\tau) = \zeta$. This involves only a small loss of generality, since $\zeta(\tau)$ is integrated out in equation (5) and is thus only a level shifter. Similarly, from equation (5), the individual factor loadings $\theta(\tau)$ do not matter for yields; only the weighted sum $\int_0^T \tau'\theta_t(\tau')\mathrm{cov}_t\left[y_{t+1}^{(\tau)}, y_{t+1}^{(\tau')}\right]d\tau'$ 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. For simplicity, I follow Greenwood, Hanson, and Vayanos (2015) by assuming that this function is linear across maturities:

$$\theta(\tau) = \left(1 - \frac{2\tau}{T}\right) \tag{16}$$

This specification is parsimonious in the sense that it does not involve any parameters beyond those already discussed. The supply process is depicted in Figure 2. Positive

supply shocks tilt the outstanding distribution of bonds toward shorter maturities.

A helpful summary measure of supply that is frequently used in the literature is the amount of "ten-year equivalent" bonds outstanding. This measure is equal to the dollar value of ten-year bonds that would produce the same duration-weighted value as the actual portfolio of outstanding 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 (10YE) is defined as

$$10\text{YE}_t = \frac{v}{10} \int_0^T \tau s_t(\tau)\, d\tau \tag{17}$$

where v is the length of one period, expressed in years.[8] Given (14) and (16), changes in β_t between two periods can be converted into percentage changes in ten-year equivalent bonds as follows:

$$
\begin{aligned}
\%\Delta 10\text{YE}_{t+1} &= \frac{\int_0^T \tau \left[\left(1 - \frac{2\tau}{T}\right)\left(\beta_{t+1} - \beta_t\right)\right] d\tau}{\int_0^T \tau \left[\zeta + \left(1 - \frac{2\tau}{T}\right)\beta_t\right] d\tau} \\
&= -\frac{1}{3\zeta - \beta_t}\Delta\beta_{t+1}
\end{aligned}
\tag{18}
$$

This equation provides a convenient way of translating real-world changes in the outstanding bond distribution into the asset-supply shocks of the model.

3 Calibration and Solution

There are eight parameters to be calibrated in this model. I consider two specifications: one in which b is calibrated to enforce a realistic lower bound on the short rate, and one in which b is set to $-\infty$, delivering an affine model. The two calibrations are summarized in Table 2.

I calibrate the autoregressive coefficient on the supply factor ϕ_β to match the persistence of the weighted-average maturity of outstanding Treasury debt. The Treasury duration series is depicted in Figure 3. (It is the same series used in the regressions in Table 1.) In the data, this variable is calculated as the value-weighted timing of all

---

[8]Note that, since all bonds are assumed to be zero-coupon, duration is simply equal to maturity.

nominal cash flows (principal and coupon payments) on all Treasury instruments held by the public, as reported in CRSP.[9] In the model, the weighted-average maturity of the debt held by investors at any point in time is linear in β_t,[10] so it has the same persistence. However, the results presented below are largely insensitive to variations in the size of the persistence parameter. The parameter σ is not identified separately from the risk-aversion parameter a. I follow Greenwood, Hanson, and Vayanos (2015) and assign it a value such that the unconditional standard deviation of β_t is 0.25, although this is arbitrary and without loss of generality given that a can be freely calibrated.

I calibrate the remaining parameters jointly to match the empirical features of the yield curve over the long run. Specifically, I use the Gurkaynak, Sack, and Wright (2007) zero-coupon yields available on the Federal Reserve Board's website and calculate the empirical moments from a sample beginning August 1971 because at that time 10-year yields become available. I also used a sample excluding the volatile period of the 1970s without much effect on the results. The sample ends in December 2015. Because, in the model, I will take a "period" to be one quarter, I average the daily data to obtain quarterly values.

The specific moments I match are the unconditional mean and standard deviation of the three-month Treasury yield (5.1% and 3.6% respectively), the unconditional mean and standard deviation of the ten-year yield (6.7% and 2.8%), and the correlation between the three-month and ten-year yield (0.94%).[11] For the model with the binding lower bound, I also match the average value of the three-month Treasury rate during the ELB period. Specifically, between December 2008 and December 2015, the three-month rate averaged 0.22%, with a maximum value of 0.68%. The calibration that achieves a mean short rate of 0.0022 conditional on $\widehat{r}_t < 0.0068$, given the other values of the short-rate parameters, is $b = 0.0017$.

In general, the model developed in Section 2 does not have an analytical solution. I solve globally by discretizing the state space and iteratively (a) calculating state-contingent yields in equation (5) given conditional expectations, and (b) calculating conditional expectations given state-contingent prices using the transition densities

---

[9]Center for Research in Security Prices, Booth School of Business, The University of Chicago. Used with permission. All rights reserved. crsp.uchicago.edu. I exclude callable bonds from this calculation.

[10]Specifically, it is equal to $T(\frac{1}{2} - \frac{1}{6\zeta}\beta_t)$.

[11]Note 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.

implied by equations (8) and (15). The details of the solution method are described in the appendix.

4   Results

4.1   Model fit

Table 3 summarizes the properties of bond yields produced by the calibrations with and without the ELB imposed, for a variety of maturities, 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 $\widehat{r}_t$ and β_t. To illustrate the importance of the ELB, I report the results in three different regions of the state space: (1) short rate below 0.68%; (2) short rate between 0.68% and its unconditional mean of 5.1%; and (3) short rate above 5.1%. The reason for the choice of 0.68% for the first threshold is that, as noted above, this 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 bp range (December 2008 through December 2015).

Not surprisingly, the constrained model performs much better than the unconstrained model when the short rate is at its lower bound. The affine model predicts that, conditional on being less than 0.68%, the short rate will average -1.0%. The shadow-rate model exactly matches the observed conditional average of 0.2%, since the value of b is calibrated to do so. However, adding this one parameter also allows the shadow-rate model to perform better in nearly every other respect near the ELB. In particular, it does a better job of capturing the frequency with which this region of the space is visited, the relative flatness of the slope in this region, and, apart from the very long end of the curve, the volatility of the slope.

The shadow-rate model achieves these successes near the ELB without sacrificing performance relative to the affine model in other regions of the state space. Indeed, again with the exception of the 20-year maturity, where neither model does particularly well, the shadow-rate model generally performs at least as well as—and in some dimensions significantly better than—the affine model. There are two reasons why this is so. First, even when the short rate is not immediately constrained, the existence of the lower bound introduces a positive probability that it will be constrained in the future. This reduces the variance of future short rates, roughly in proportion to the

lower bound's proximity, with consequences for term premia that do not appear in the affine model. Second, since the assumption of $b = 0.0017$ allows the model to fit the data well in the constrained region of the space, the other parameters are free to adjust to fit the unconstrained regions. In the affine model, those parameters must struggle to trade off the different behavior of yields across the possible values of $\widehat{r}_t$.

4.2   Yield determination in the shadow-rate model

I now consider the effects of the ELB on the yield curve and on the propagation of the two shocks in different regions of the state space. To facilitate this discussion, consider the first-order Taylor series expansion of the τ-maturity yield. For arbitrary state values, $\widehat{r}$ and β, we have

$$y_t^{(\tau)} \approx C_t^{(\tau)} + A_{\widehat{r},t}^{(\tau)}\widehat{r} + A_{\beta,t}^{(\tau)}\beta \tag{19}$$

where $A_{\widehat{r},t}^{(\tau)} \equiv \partial y_t^{(\tau)}/\partial \widehat{r}_t$ and $A_{\beta,t}^{(\tau)} \equiv \partial y_t^{(\tau)}/\partial \beta_t$, with both derivatives evaluated at the time-t values of the states, and $C_t^{(\tau)}$ is the corresponding intercept term. In general, these derivatives vary over values of the state variables; hence, their time subscripts. However, in the case of the affine process for the short rate ($b = -\infty$), yields themselves are affine in the states (as shown, for example, in Greenwood and Vayanos, 2014). Thus, equation (19) holds exactly, and the derivatives are constant. For example, consider the calibration of the shadow-rate model (line 1 of Table 2), but let $b = -\infty$. Then $A_{\widehat{r}}^{(10y)} = 0.72$ and $A_{\beta}^{(10y)} = -0.033$, meaning that a one-percentage-point increase in the short rate in that model translates into a 72-basis-point increase in the ten-year yield and a one-unit change in the bond-supply factor (four unconditional standard deviations) translates into a 3.3-percentage-point decline in the ten-year yield. Again, because of the linearity, these magnitudes do not depend on the initial values of the state variables.

Once we allow for a finite ELB, the coefficients in equation (19) are no longer constant. This is illustrated by the solid lines in Figures 4 and 5, which show the factor loadings across different values of the state variables when we set $b = 0.0017$. The dashed lines in the figures show the factor loadings in the same model, with the ELB removed; for any given maturity, as just noted, these are constant across all values of $\widehat{r}_t$ and β_t.[12]   In particular, Figure 4 shows the loadings at various maturities

---

[12]Note that this is not the same model that is described in the second line of Table 1, which was

for the shadow rate $A_{\widehat{r},t}^{(\tau)}$ (on the left) and the bond-supply factor $A_{\beta,t}^{(\tau)}$ (on the right) across a range of values for $\widehat{r}_t$, holding β_t fixed at its mean value of zero. Figure 5 shows these two factor loadings across a range of values for β_t, holding $\widehat{r}_t$ fixed at either the unconditional mean of the short rate (Panel A) or at a value of -2.7% (Panel B). I choose the latter value for illustration of the ELB environment because it is the average of the shadow rate estimated by Krippner (2012). To interpret its magnitude, it implies a modal time to "liftoff" of approximately 7 quarters, which is roughly consistent with the outcomes of surveys of market participants' expectations and other evidence collected during much of the ELB period (see Femia, Friedman, and Sack (2013)).

The shapes of the factor loadings across state values in the shadow-rate model follow from the nonlinear conditional moments of the short rate that were depicted in Figure 1. As in that figure, the behavior of the shadow-rate model approaches that of the affine model as $\widehat{r}_t$ moves arbitrarily far above b. But for values of the shadow rate near and below b, nonlinearities become material, and the behavior of yields differs both quantiatively and qualitatively from that implied by the affine model.

As shown in Figure 1, for lower values of the shadow rate, both the conditional mean and the conditional variance of r_{t+h} in future periods falls. The lower mean leads to lower yields through the expectations component, and the lower variance leads to lower yields through the term premium. Thus, as shown in Figure 4, $A_{\widehat{r},t}^{(\tau)}$ is monotonically increasing in $\widehat{r}_t$. Furthermore, for $\widehat{r}_t$ low enough, $A_{\widehat{r},t}^{(\tau)} < A_{\widehat{r},t}^{(\tau')}$ when $\tau < \tau'$. 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 $\widehat{r}_t > b$, and it is the opposite of what an affine model would predict.

Similarly, in the shadow-rate model, $A_{\beta,t}^{(\tau)}$ is monotonically decreasing (i.e., becoming more negative) in $\widehat{r}_t$. Again, this is because $\mathrm{var}_t[r_t]$ falls as $\widehat{r}_t$ moves below the ELB, causing the covariance terms in (5) to become smaller. Consequently, at the ELB— and particularly when $\widehat{r}_t$ is deeply negative—bond-supply shocks have smaller effects on yields than they do in an affine model.

These 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

---

recalibrated to match the moments of the empirical yield curve. Rather, in order to isolate the effect of the ELB, all parameters other than b are the same as in the *top* line of Table 1. The derivatives are all calculated numerically, based on the model solution discussed in the appendix.

prior sample. This is exactly the result predicted by model in the right-hand panel 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 left-hand panel, and in particular how they cross in the sub-ELB region. For example, the two-year yield is more responsive to shadow-rate shocks than the 10-year yield is during normal times, but it less responsive than the 10-year yield when the shadow rate is negatve. Indeed, the model predicts that, at the ELB, the sensitivity of the 10-year yield to the 2-year yield should be higher than the sensitivity of the 5-year yield to the 2-year yield. This increasing sensitivity across maturities was also evident in the empirical results in Table 1.

Turning to Figure 5, the existence of the ELB also creates nonlinearities across different values of the supply factor. The differences from the affine model are non-negligible (though relatively modest) even when the short rate is at its unconditional mean. But they are of first-order importance when the shadow rate is negative. Two particularly notable results stand out in that region. First, in addition to being generally lower, $A_{\hat{r},t}^{(\tau)}$ is 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 rate volatility.

The second notable result is that, over most regions of the state space, $A_{\beta,t}^{(\tau)}$ is increasing (becoming less negative) in β_t. This implies, for example, that the marginal effects of asset purchases decline as the central bank does more of them. Intuitively, at the ELB, higher levels of β_t (less duration in the market) are associated with smaller term premium effects of shocks to the shadow rate. Consequently, positive shocks to β_t reduce the volatility of yields, making *further* shocks to β_t less potent. This result will also be important for analyzing the relative effectiveness of alternative policies in different environments in Section 6.

4.3   Dynamic effects of shocks

To see the results in another way, Figure 6 shows the dynamic effects of shocks in the shadow-rate model, both with and without the ELB in place. I consider independent

one-standard-deviation shocks to the shadow rate and bond supply in directions that lower yields ($e_t^r = -.0079$ or $e_t^\beta = 0.086$). Since the nonlinearities can be important, starting values matter. I consider two cases, one in which the short-rate begins at its unconditional mean of 5.1%, and one in which it is constrained by the ELB. (In both cases, I let β_t start at is mean value of zero, since the nonlinearities are relatively minor for bond-supply shocks of average size.) More specifically, for the ELB-constrained case, I start the shadow rate at -2.7%, its ELB-period average according to the Krippner estimates.

For each set of starting values, I simluate the model forward ten years, both with the one-standard-deviation shocks in the first period and without them. I compute impulse-response functions as the difference between those two simulations. In the figure, each IRF depicts the response of the entire yield curve over the ten-year period, with maturities on the lower-left axis and calendar time on the upper-left axis. Panel A shows the response of spot yields, while Panel B presents the same information in terms of forward rates, where the patterns are often easier to see.

When the short-rate is unconstrained, the shock to $\widehat{r}_t$ has a monotonic effect on the yield curve that reflects only the expectations of its first-order autoregressive dynamics. It lowers the ten-year yield by 50 basis points. On the other hand, the response to the bond-supply shock displays nearly the opposite pattern, with no effect on short-term yields and fairly large effects at the long end. The supply shock initially lowers the ten-year yield by 28 basis points.[13] The effects of both shocks decay monotonically over time, with the yield curve coming most of the way back to its starting position by the end of the ten years.

When the shadow rate starts from a value notably below b, the situation is much different. For one thing, the effects of both shocks are smaller; the shadow-rate shock now reduces the ten-year yield by just 38 basis points, while the bond-supply shock reduces it by 19 basis points. For the shadow-rate shock, the damped response occurs because the short rate was already expected to be near zero with near certainty for several quarters; the shock cannot move the expectations component at these horizons any lower. For the bond-supply shock, as discussed above, the attenuation is due to interest-rate volatility being lower at the ELB. This causes the covariance terms that

---

[13]The bond-supply shock has a hump-shaped effect on forward rates, similar to that discussed in Greenwood et al. (2015b), with the peak occuring around the 12-year maturity. As discussed in King (2014), this is because the change in term premia depends on the expected change in duration risk over the life of each bond. Since shocks to β_t have a half-life of about four years, very long-term bonds see less of a change in this lifetime average than medium-term bonds do.

multiply β_t in (5) to be smaller, particularly for bonds with relatively short maturities, which can expect to be near the ELB for their entire remaining lives.

Another important feature of the results at the ELB is that the shock to $\widehat{r}_t$ now produces a hump-shaped reaction across maturities in forward rates. This pattern will be critical for matching the observed forward-curve response in event studies, considered below. Again, it occurs because shorter-term forwards cannot move much lower to begin with; the deeply negative shadow rate has already depressed them to near b. Meanwhile, as in the unconstrained case, long-term forwards are not much affected by shadow-rate shocks, because they depend mostly on expectations of the short rate in the far future. Consequently, both the expectations and term-premium components of shadow-rate shocks have their largest effects on medium-term yields. This prediction of the model is consistent with the evidence presented by Swanson and Williams (2014), who show that responses of shorter-term yields to macroeconomic shocks were muted during the ELB period, and Carvalho, Hsu, and Nechio (2016), who show that Federal Reseve communications had their largest effects in the 2- to 10-year maturity range during the ELB period. The exact maturity of the peak of the hump (about 7 years, in the case shown in the figure) depends on the size of the shock and how far below b the shadow rate initially is.

5   Assessing Unconventional Policy

5.1   Event-study evidence

To assess whether the model can generate realistic effects of unconventional policy, I begin with a simple event study. Table 4 lists 29 dates during the period November 2008 to September 2015 on which the FOMC or its chair delivered important communications related to unconventional monetary policy. This list was compiled by taking the union of dates used in event studies by Fawley and Neeley (2013), Greenwood, Hanson, and Vayanos (2015), and Swanson (2016). Those studies exclude a few important forward-guidance-related events late in the ELB period, and I therefore also incorporate the list of important unconventional-policy dates maintained by the Reserve Bank of New Zealand.[14] The table includes announcements pertaining to asset purchases ("QE"), forward guidance ("FG"), or both, with the type of each

---

[14]http://www.rbnz.govt.nz/-/media/ReserveBank/Files/Publications/Research
/Additional%20research/Leo%20Krippner/US-monthly-update.xls?la=en

announcement indicated in the last column.

One might like to be able to perform separate event studies for asset purchases and forward guidance in order to estimate the effects of each policy separately. Unfortunately, this is not possible because, as noted by Swanson (2016), many of the most important communications during this time contained news about both types of policy. For example, the largest change in yields in the sample occurred following the FOMC statement of March 18, 2009. This statement announced both a large expansion of the Fed's first asset-purchase program ("QE1") and introduced language indicating that the Committee expected the short rate to remain near zero for "an extended period," the first instance of explicit forward guidance during this time, which was widely discussed in the press. This event and others like it create an identification problem that cannot be solved in a model-free way. For this reason, I focus on fitting and explaining the joint effects of both types of unconventional policy.

I compute two-day changes in the Gurkaynak et al. (2008) zero-coupon yields surrounding each of these events. Assuming that the communications listed in the table contained the most important innovations to market expectations about unconventional policies, the sum of the yield changes provides a rough estimate of the total effects of these policies on the term structure.

In addition to changes in yields, I compute two-day changes in the expectations and term-premium components of yields, based on the Kim and Wright (2005) model. The Kim-Wright model is a standard, three-factor term-structure model that is widely used and cited in the literature. (See, for example, Gagnon et al., 2011, and D'Amico et al., 2012.) It has the advantage that its components are identified, in part, by survey data on short-rate expectations. Thus, although it does have an affine structure that introduces measurement error near the ELB, it is still informed by market participants' actual short-rate expectations, providing discipline that other affine models do not have. Still, the model tends to associate changes in the expectations component of interest rates with changes in shorter-term yields. Since shorter-term yields move little when they are constrained by the ELB, the model likely understates the amount by which expectations change.[15]

---

[15]The Kim-Wright data are available on the Federal Reserve Board's website for maturities of up to ten years. Although the Kim-Wright model is estimated on the same yield data that I use, it has substantial fitting errors on some of the event dates due to market turbulence during the crisis. For the purposes of presentation, I normalize the changes in the expectations and term-premium components on each day such that they sum to that day's total change in yields at each maturity. This normalization makes no material difference for the results.

Figure 7 reports the sum of the changes in the 29 event windows, both as spot yields and forward rates. The thick black lines show the net change in yields (or forward rates), while the red and blue regions show the Kim-Wright decomposition into changes in short-rate expectations and changes in term premia. Three stylized facts emerge from this event study:

1. Unconventional policy shocks collectively resulted in decreases in yields across the curve, with the ten-year yield falling by about 200 basis points.

2. The largest decreases in forward rates occurred at maturities in the five- to ten-year range.

3. For maturities between 2 and 10 years, approximately a third of the decline in yields reflected a decrease in the Kim-Wright expected short-rate path, with the rest reflecting term premia. Again, this might be considered a lower bound on the size of the expectations component, given the linearity of the model.

The next section tests whether the model presented above can replicate these findings.

5.2 Simulating unconventional policy

To replicate the effects of unconventional policy, we first must translate the actions taken by the Federal Reserve during the ELB period into shocks that can be fed into the model. For shocks to the shadow rate, this is straightforward. We know exactly how long r_t remained at the ELB in practice (seven years, to the day), and the cumulative effect of the shadow-rate shocks must keep r_t at b for the same amount of time when starting from a value of $\widehat{r}_t = b$.[16]

To calculate the appropriate size of the bond-supply shocks, I consider the duration-weighted quantity of debt purchased on net by the Federal Reserve during the time that the ELB was binding. Greenwood, Hanson, Rudolph, and Summers (2015) calculate that the Fed removed approximately \$2.7 trillion of ten-year-equivalent bonds from the market during this time, including Treasuries, agency debt, and MBS. The Treasury data show that the amount of ten-year-equivalent Treasury bonds outstanding was \$2.3 trillion at the beginning this period, rising to \$4.9 trillion at the time of liftoff.

---

[16]The Federal Reserved moved its federal funds target to the 0-25 basis point range on December 16, 2008, where it remained until December 16, 2015, when it was raised to the 25-50 basis point range.

Meanwhile, SIFMA data show $6.3 trillion of agency-backed MBS and CMOs and $2.1 trillion of long-term agency debt outstanding in Q4 2008. In Q4 2015, these values were $7.2 trillion and $1.3 trillion, respectively.[17] 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, the $2.7 trillion purchased by the Fed represented approximately 50% of the $5.4 trillion of ten-year-equivalent government-backed debt outstanding at the beginning of the period. Looked at another way, ten-year equivalents in investors' hands would have been $7.7 trillion, instead of $5.0 trillion, at the end of 2015 if the Fed had not engaged in QE—-a difference of about 35%. Equation (18) translates these percentage changes into changes in β_t.

If the model were linear, the impact of a 100 bp shock would be the same as the sum of the impact of ten 10-basis point shocks. But, because of the nonlinearities, how the cumulative changes in shadow rates and bond supply are divided over time matters. I consider two polar scenarios, one in which a large shadow-rate and a large bond-supply shock occur all at once, and one in which these shocks arrive in equal quarterly installments over a seven-year period. The reality was between these two cases—the Fed spread its unconventional policy over seven years, but it did so in irregular, relatively large moves.

In both the "all at once" and "quarterly installments" scenarios, I initialize $\widehat{r}_t$ at a value of 0.17%, exactly at the ELB, and I initialize β_t at a value of 0.009, which produces a 10-year slope of the yield curve of 340 basis points, the observed value in November 2008, just before unconventional policy began. In the "all at once" scenario, I set the supply shock to $e_t^\beta = 0.15$ in the first period of the simulation, corresponding to a decline in 10-year equivalents of 50%, and $e_t^\beta = 0$ for all periods thereafter. I set the shadow-rate shock to $e_t^r = -0.035$ in the first period and zero for all periods thereafter, which is sufficient to keep the short rate at the ELB for exactly 28 quarters. In the "quarterly installments" case, I parcel out the total change in ten-year equivalents into 28 equal supply shocks that result in a cumulative decline in ten-year equivalents of 35%, relative to a baseline simulation in which there are no shocks. This gives $e_t^\beta = 0.008$ for each of the 28 periods. Similarly, I introduce 28 sequential and equal shadow-rate shocks, each of which is sufficient to lower $\widehat{r}_t$ by just enough that it returns exactly to the ELB in the next quarter. This gives each of the e_t^r equal to -0.0009. To report the results, I sum the effects of all the shocks across the

28 quarters, analogously to the way that the cumulative responses were constructed in the event study of the previous section.

Figure 8 shows the total simulated effects of unconventional policy on the spot and forward curves under each scenario and breaks these responses into the expected path of short rates and the term premium, as was done in the event study.[18] The model reproduces the three stylized facts noted in the previous section. The 10-year yield falls by 187 basis points in the "all at once" scenario and 206 basis points in the "quarterly installments" scenario. The decrease in forward rates is largest at the 7-year maturity in the former case and the 5-year maturity in the latter case. The expectations component of yields is a bit larger than in the event study—about half of the overall response, rather than one-third—but, as noted earlier, the Kim-Wright decomposition likely understates the contribution of expectations in this case. The close match between the event study and the model occurs despite the fact that none of the model's parameters (other than the level of the ELB itself) were specifically calibrated to match the effects of unconventional policy or the environment under which it was conducted. Together with the model's ability to match the patterns in Table 1 and the evidence of its overall fit provided in Table 3, they suggest that it does a good job of describing the behavior of the yield curve and its response to policy shocks of different kinds at the ELB.

Table 5 shows how the total change in yields breaks down between various channels in each scenario. This breakdown is calculated by computing what the change in yields would be if only the shadow rate or the bond-supply shocks occurred. 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. The effects of these nonlinearities, which are reported in the "interaction" column of the table, loom a bit larger for the "all at once" case than for the series of small shocks, since the latter maintain the state variables in a localized region.

Overall, the shadow-rate shocks have larger effects than the bond-supply shocks. In the "quarterly installments" scenario, they account for 141 of the total 206-basis-point change in the ten-year yield; in the "all at once" scenario, they account for 160 of the total 187 basis point change. At shorter maturities, they explain nearly the

---

[18]The change in short-rate expectations is given by equation (11), and the change in the forward term premium is simply the change in the forward rate itself minus that value.

entirety of the yield decline. The majority of the response to the shadow-rate shocks comes through the expectations component of yields. However, the term-premium response is also significant—at the ten-year maturity, it is responsible for about half a percentage point of the decline in both scenarios. On average across both scenarios, this channel of unconventional monetary policy is approximately as important for term premia as is the "duration channel" through which the bond-supply shocks operate.

6 Evaluating Policy Options

I now ask whether the *relative* efficacy of bond-supply shocks and shadow-rate shocks depends on whether the ELB is binding. One reason that this question might be of interest is that the answer may be useful for guiding policy during future ELB episodes. 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 short rate has. In an affine model, this quantity is 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 the two shocks differ in different areas of the state space, and therefore their relative potency also differs.

Figure 9 presents contour maps of this relative efficacy for 5- 10- and 20-year yields, with lighter colors indicating less-negative values—i.e., areas of the space in which bond-supply shocks have relatively large effects compared to those of shadow-rate shocks.[19] Bond-supply shocks achieve their greatest relative efficacy in the southwest quadrant of the maps, where both $\widehat{r}_t$ and β_t are deeply negative. As noted earlier, both types of 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. (This result can be seen to some extent in Figure 5, panel B.) Thus, for example, a change in β_t of about -0.04 in this region is sufficient to lower the ten-year yield by as much as a 25-basis-point shock to $\widehat{r}_t$. In contrast, at the unconditional means of the states, the size of the shock required is closer to -0.06.

Interestingly, this high-relative-potency region for the bond-supply shocks is ap-

---

[19]I do not report the map for the two-year yield because, at that maturity, the responses to both types of shocks is effectively zero for negative values of $\widehat{r}_t$, and small numerical errors consequently generate large swings in the ratios of those responses. Nonetheless, the broad patterns are similar to those for longer maturities.

proximately 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, shadow-rate models show $\widehat{r}_t$ near its nadir, with the Krippner (2012) estimate, for example, averaging -4.5% over those three years. Meanwhile, as suggested by Figure 2, the Treasury was lengthening the maturities of its issuance, so that the average duration outstanding stood near the upper end of its historical range. Furthermore, fiscal expansion increased the total quantity of Treasury debt outstanding, further expanding the amount of interest-rate risk held by investors. 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 Committee perceived that the marginal benefits of forward guidance declined sufficiently relative to those of asset purchases.

7   Conclusion

This paper has augmented the standard model of risk-averse arbitrage in the bond market to account for the effective lower bound on nominal interest rates. The model successfully reproduces the conditional moments of the yield curve, particularly near the ELB. To the extent that such models are used to study the effects of unconventional monetary policy, incorporating the ELB is of first-order importance, since it is in environments in which the short rate is constrained that such policies have been implemented in the past and are most likely to be implemented in the future. Affine models do a poor job of matching the behavior of yields in this region of the space and produce misleading predictions about the effects of policy.

The model was used to explore the channels through which unconventional monetary policy operated in the U.S.. The main finding was that the majority of the effects of such policies came through the expectations component of yields. The term premium effects of changes in policy expectations—a channel that does not exist in affine models and has been largely ignored by previous literature—also played a significant role. The duration effects of bond-supply shocks were relatively weak, likely accounting for only about a quarter of the overall change in the ten-year yield and even less at shorter maturities. The Fed bought approximately $3 trillion of longer-term bonds on net and

reduced the ten-year yield through this channel by at most 60 basis points, about 2 basis points per $100 billion. The finding that the duration channel of asset purchases is likely small suggests that policymakers may be hesitant to rely on that channel if they have other tools—such as forward guidance—at their disposal. However, the results of Section 6 suggest that, in states of the world in which the ELB is a serious constraint, asset purchases may become more attractive, particularly if the quantity of duration risk in the market is high.

Most empirical papers estimate an effect of asset purchases that is much larger than the duration effect implied by the model. For example, surveying the literature, Williams (2014) suggests that the Fed's programs taken together might have resulted in a decrease of 150 basis points in the ten-year yield. This is about three times larger than the duration effects suggested by the model. Since changes in short-rate expectations were the most important driver of yields in the model, the results here could be reconciled with the empirical evidence through a signaling channel of asset purchases, although neither event studies nor models can distinguish such a channel from explicit forward guidance. It may also be that other channels that are not captured by the model are at work. For example, D'Amico and King (2013) and Cahill et al. (2015) find a large "local supply" effect of Treasury purchases that is orthogonal to the duration channel. Incorporating such channels into strucutral models is an important direction for future research.

References

ADRIAN, T., E. ETULA, AND T. MUIR (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance*, 64(6), 2557–2596.

ALTAVILLA, C., G. CARBONI, AND R. MOTTO (2015): "Asset purchase programmes and financial markets: lessons from the euro area," European Central Bank Working Papers, no 1864.

BAUER, M. D., AND G. D. RUDEBUSCH (2014): "The Signaling Channel for Federal Reserve Bond Purchases," *International Journal of Central Banking*, 10(3), 233–289.

BHATTARAI, S., G. B. EGGERTSSON, AND B. GAFAROV (2015): "Time Consistency

and the Duration of Government Debt: A Signalling Theory of Quantitative Easing," NBER Working Paper 21336.

CAHILL, M. E., S. D'AMICO, C. LI, AND J. S. SEARS (2013): "Duration Risk versus Local Supply Channel in Treasury Yields: Evidence from the Federal ReserveâĂŹs Asset Purchase Announcements," Federal Reserve Board Finance and Economics Discussion Series.

CARVALHO, C., E. HSU, AND F. NECHIO (2016): "Measuring the Effect of the Zero Lower Bound on Monetary Policy," Federal Reserve Bank of San Francisco Working Paper Series.

D'AMICO, S., W. ENGLISH, D. LOPEZ-SALIDO, AND E. NELSON (2012): "The Federal Reserve's Large-scale Asset Purchase Programmes: Rationale and Effects," Federal Reserve Board Finance and Economics Discussion Series.

D'AMICO, S., AND T. B. KING (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., S. FRIEDMAN, AND B. SACK (2013): "The Effects of Policy Guidance on Perceptions of the FedâĂŹs Reaction Function," Federal Reserve Bank of New York, Staff Report no.652.

GREENWOOD, R., S. HANSON, J. S. RUDOLPH, AND L. H. SUMMERS (2015): "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, 43âĂŞ89. Brookings Institution Press.

GREENWOOD, R., S. HANSON, AND D. VAYANOS (2015): "Forward Guidance in the Yield Curve: Short Rates versus Bond Supply," NBER Working Paper No. 21750.

GREENWOOD, R., AND D. VAYANOS (2014): "Bond Supply and Excess Bond Returns," *Review of Financial Studies*, 27(3), 663–713.

GURKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2007): "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 54(8), 2291âĂŞ2304.

HADDAD, V., AND D. A. SRAER (2015): "The Banking View of Bond Risk Premia,"
.

HAMILTON, J. D., AND J. C. WU (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.

HAYASHI, F. (2016): "Affine Term Structure Pricing with Bond Supply as Factors,"
Working Paper.

HE, Z., B. KELLY, AND A. MANELA (2016): "Intermediary Asset Pricing: New
Evidence from Many Asset Classes," Working Paper.

HE, Z., AND A. KRISNAMURTHY (2013): "Intermediary Asset Pricing," *American
Economic Review*, 103(3), 732–70.

KAMINSKA, I., AND G. ZINNA (2014): "Official Demand for U.S. Debt: Implications
for U.S. Real Rates," IMF Working Paper.

KIM, D. H., AND K. J. SINGLETON (2012): "Term Structure Models and the Zero
Bound: An Empirical Investigation of Japanese Yields," *Journal of Econometrics*,
170(1), 32–49.

KIM, D. H., AND J. H. WRIGHT (2005): "An Arbitrage-Free Three-Factor Term
Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon
Forward Rates," Finance and Economics Discussion Series (The Federal Reserve
Board).

KING, T. B. (2014): "A Portfolio-Balance Approach to the Nominal Term Structure,"
Working Paper.

KRIPPNER, L. (2012): "Modifying Gaussian Term Structure Models When Interest
Rates Are Near the Zero Lower Bound," Reserve Bank of New Zealand Working
Paper.

KRISHNAMURTHY, A., AND A. VISSING-JORGENSEN (2012): "The Aggregate Demand
for Treasury Debt," *Journal of Political Economy*, 120(2), 233–267.

——— (2013): "The Ins and Outs of Large Scale Asset Purchases," Kansas City Federal Reserve Symposium on Global Dimensions of Unconventional Monetary Policy.

MONFORT, A., F. PEGORARO, J.-P. RENNE, AND G. ROUSSELLET (2015): "Staying at Zero with Affine Processes: An Application to Term Structure Modeling," Banque de France Working Paper.

SWANSON, E. (2016): "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets," Working Paper.

SWANSON, E. T., AND J. C. WILLIAMS (2014): "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review*, 104(10), 3154–3185.

VAYANOS, D., AND J.-L. VILA (2009): "A Preferred-Habitat Model of the Term Structure of Interest Rates," NBER Working Paper 15487.

WILLIAMS, J. C. (2014): "Monetary Policy at the Zero Lower Bound: Putting Theory into Practice," Hutchins Center on Fiscal and Monetary Policy at Brookings.

WOODFORD, M. (2012): "Methods of Policy Accommodation at the Interest-Rate Lower Bound," Presented at the Jackson Hole symposium.

Appendix: Solution Algorithm

Transform log prices from the time/maturity domain to the state-space domain by writing them as a function $p(.)$:

$$p_t^{(\tau)} = p\left(\tau, \widehat{r}_t, \beta_t\right) \tag{20}$$

We want to solve for this function.

Discretize the maturity space into T bonds and the state space into N nodes. Let $\widehat{r}^n$ and β^n denote the values of the state variables at node n. Then, from equations (3), (14), and (16), we have

$$
\begin{aligned}
p\left(\tau, \widehat{r}^n, \beta^n\right) &= E\left[p\left(\tau - 1, \widehat{r}_{t+1}, \beta_{t+1}\right) | \widehat{r}^n, \beta^n\right] + \widehat{r}^n \\
&+ a\sum_{\tau'=1}^{T} \beta^n \left(\zeta + 1 - \frac{2\tau'}{T}\right) \mathrm{cov}\left[p\left(\tau', \widehat{r}_{t+1}, \beta_{t+1}\right), p\left(\tau - 1, \widehat{r}_{t+1}, \beta_{t+1}\right) | \widehat{r}^n, \beta^n\right]
\end{aligned} \tag{21}
$$

Rewrite the law of motion for the state vector as the conditional probability function

$$\pi\left(\widehat{r}_{t+1}, \beta_{t+1} | \widehat{r}_t, \beta_t\right) = \varphi\left(\widehat{r}_{t+1} - \mu^{\widehat{r}}(1 - \phi^{\widehat{r}}) - \phi^{\widehat{r}}\widehat{r}_t\right) \varphi\left(\beta_{t+1} - \phi^\beta \beta_t\right) \tag{22}$$

where $\varphi(.)$ is the standard-normal PDF (and the separability makes use of the independence of $\widehat{r}_t$ and β_t).

The algorithm proceeds as follows:

Step 0. Set $i = 0$. Begin with an initial guess of the pricing function $p^0(.)$. For example, choose $p^0(\tau, \widehat{r}_t, \beta_t) = 0$ for all $\tau, \widehat{r}_t, \beta_t$.

Step 1. At each node n, evaluate the functions

$$
\begin{aligned}
\mu^i\left(\tau, \widehat{r}^n, \beta^n\right) &\equiv \int\int \pi\left(\widehat{r}, \beta | \widehat{r}^n, \beta^n\right) p^i\left(\tau - 1, \widehat{r}, \beta\right) d\widehat{r} d\beta \\
&= E\left[p^i\left(\tau - 1, \widehat{r}_{t+1}, \beta_{t+1}\right) | \widehat{r}^n, \beta^n\right]
\end{aligned} \tag{23}
$$

and

$$
\begin{aligned}
\omega^i\left(\tau, \tau', \widehat{r}^n, \beta^n\right) &\equiv \int\int \{\pi\left(\widehat{r}, \beta | \widehat{r}^n, \beta^n\right) \left[p^i\left(\tau - 1, \widehat{r}, \beta\right) - \mu^i\left(\tau, \widehat{r}^n, \beta^n\right)\right] && (24) \\
&\quad \times \left[p^i\left(\tau', \widehat{r}, \beta\right) - \mu^i\left(\tau', \widehat{r}^n, \beta^n\right)\right]\} d\widehat{r} d\beta && (25) \\
&= \mathrm{cov}\left[p^i\left(\tau', \widehat{r}_{t+1}, \beta_{t+1}\right), p^i\left(\tau - 1, \widehat{r}_{t+1}, \beta_{t+1}\right) | \widehat{r}^n, \beta^n\right]
\end{aligned}
$$

31

Step 2. Update the pricing function by calculating, at each node,

$$p^{i+1}\left(\tau,\widehat{r}^{n},\beta^{n}\right) = \mu^{i}\left(\tau,\widehat{r}^{n},\beta^{n}\right) + \widehat{r}^{n} + a\sum_{\tau'=1}^{T}\beta^{n}\left(\zeta + 1 - \frac{2\tau'}{T}\right)\omega^{i}\left(\tau,\tau',\widehat{r}^{n},\beta^{n}\right) \quad (26)$$

Set $i = i + 1$.

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

The expectations in Step 1 are computed numerically using the probability function $\pi(.)$ and the pricing function $p^{i}(.)$. The integration is performed by quadrature and, to ensure accuracy, relies on a much finer grid than the price computation does. To obtain bond prices over this refinement of the space, the values of $p^{i}(.)$ 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. (Note that, 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 $T = 80$ (bonds maturing every quarter up to twenty years) and calculate bond prices at $N = 625$ nodes distributed uniformly over $\widehat{r} = (-0.20, 0.30)$ and $\beta = (-1.5, 1.5)$. 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, taking about 6 hours on a quad-core Intel i7 4.0-GHz processor with 1867 MHz DDR3 RAM running Mathematica version 11.0.

Table 1. Relationship of long-term yields to Treasury supply and short-term yields

| [1] Dep. Var. | Independent variables | | | | | | [8] Adj. R² |
| | Maturity-weighted debt/GDP | | | 1y yield | | | |
	[2] Pre-ELB	[3] ELB	[4] Break t-stat	[5] Pre-ELB	[6] ELB	[7] Break t-stat	
5y yield	0.184* (0.101)	-0.058 (0.085)	-2.72	0.852*** (0.050)	2.315*** (0.607)	2.50	0.951
10y yield	0.265** (0.130)	-0.051 (0.096)	-2.93	0.749*** (0.063)	3.526*** (0.713)	4.07	0.901
15y yield	0.301** (0.140)	-0.019 (0.106)	-2.85	0.703*** (0.070)	3.698*** (0.715)	4.44	0.871

| Dep. Var. | Independent variables | | | | | | Adj. R² |
| | WAM of Treas. debt | | | 1y yield | | | |
	Pre-ELB	ELB	Break t-stat	Pre-ELB	ELB	Break t-stat	
5y yield	0.145 (0.097)	-0.002 (0.103)	-2.33	0.844*** (0.050)	2.623*** (0.644)	2.78	0.950
10y yield	0.234* (0.121)	0.053 (0.117)	-2.54	0.739*** (0.060)	3.875*** (0.700)	4.55	0.900
15y yield	0.278** (0.133)	0.106 (0.128)	-2.31	0.692*** (0.065)	3.998*** (0.643)	5.27	0.870

| Dep. Var. | Independent variables | | | | | | Adj. R² |
| | Maturity-weighted debt/GDP | | | 2y yield | | | |
	Pre-ELB	ELB	Break t-stat	Pre-ELB	ELB	Break t-stat	
5y yield	0.129* (0.067)	-0.039 (0.053)	-2.82	0.908*** (0.033)	1.886*** (0.199)	5.32	0.981
10y yield	0.217** (0.103)	-0.034 (0.077)	-2.79	0.804*** (0.050)	2.472*** (0.265)	7.15	0.943
15y yield	0.252** (0.117)	0.006 (0.095)	-2.43	0.756*** (0.060)	2.433*** (0.309)	6.24	0.915

| Dep. Var. | Independent variables | | | | | | Adj. R² |
| | WAM of Treas. debt | | | 2y yield | | | |
	Pre-ELB	ELB	Break t-stat	Pre-ELB	ELB	Break t-stat	
5y yield	0.105* (0.062)	-0.001 (0.061)	-2.60	0.902*** (0.032)	2.014*** (0.186)	6.13	0.981
10y yield	0.196** (0.093)	0.058 (0.087)	-2.35	0.796*** (0.048)	2.578*** (0.224)	8.31	0.942
15y yield	0.239** (0.108)	0.121 (0.105)	-1.72	0.748*** (0.055)	2.473*** (0.255)	7.11	0.915

Notes: Each row 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 maturity-weighted Treasury-debt-to-GDP ratio or the weighted-average maturity of Treasury debt in public hands 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. 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. Calibration

	b	μ^r	ϕ^r	σ^r	ϕ^β	σ^β	a	ζ
Shadow-rate model	0.17%	4.2%	0.98	0.79%	0.94	0.086	0.38	0.10
Affine model	-∞	4.4%	0.98	0.73%	0.94	0.086	0.37	0.12

Notes: The table shows the calibrated values of the parameters in the baseline shadow-rate model, as well as in an affine model. In both cases values are chosen to match the same moments of the data, as described in the text.

Table 3. Conditional Moments of Yield Curve

Short rate below 0.68%

	Fraction of obs.	Short rate	Slopes			
			2Y	5Y	10Y	20Y
Mean						
Data	*20%*	*0.2%*	*0.3%*	*1.3%*	*2.6%*	*3.4%*
Shadow-Rate Model	*18%*	*0.2%*	*0.5%*	*1.2%*	*2.5%*	*4.8%*
Affine Model	*13%*	*-1.0%*	*0.8%*	*2.1%*	*3.8%*	*6.0%*
Standard Deviation						
Data		*0.1%*	*0.3%*	*0.6%*	*0.8%*	*0.7%*
Shadow-Rate Model		0.1%	0.2%	0.5%	0.9%	1.1%
Affine Model		1.4%	0.1%	0.3%	0.5%	0.8%

Short rate between 0.68% and 5.1%

	Fraction of obs.	Short rate	Slopes			
			2Y	5Y	10Y	20Y
Mean						
Data	*25%*	*2.5%*	*0.7%*	*1.7%*	*2.7%*	*3.3%*
Shadow-Rate Model	35%	2.7%	0.6%	1.4%	2.5%	4.4%
Affine Model	37%	2.7%	0.6%	1.4%	2.5%	3.8%
Standard Deviation						
Data		*1.1%*	*0.5%*	*0.9%*	*1.1%*	*1.2%*
Shadow-Rate Model		1.1%	0.1%	0.5%	0.8%	1.1%
Affine Model		1.1%	0.1%	0.2%	0.4%	0.6%

Short rate above 5.1%

	Fraction of obs.	Short rate	Slopes			
			2Y	5Y	10Y	20Y
Mean						
Data	*54%*	*6.9%*	*0.6%*	*0.9%*	*1.2%*	*1.5%*
Shadow-Rate Model	47%	7.3%	0.4%	0.7%	1.2%	2.0%
Affine Model	50%	7.1%	0.4%	0.6%	1.0%	1.4%
Standard Deviation						
Data		*2.3%*	*0.7%*	*1.0%*	*1.2%*	*1.3%*
Shadow-Rate Model		2.2%	0.2%	0.6%	1.0%	1.5%
Affine Model		2.0%	0.1%	0.4%	0.7%	1.1%

Notes: The table shows conditional moments of zero-coupon yields simulated from the shadow-rate and affine models, based on the calibrations shown in Table 1, together with the corresponding moments from the data. Model results are based on 2 million simulations. Yield data are quarterly, 1971:3 – 2015:4.

Table 4. Event days

Date	Event	Type of policy action
11/25/08	First MBS and agency-debt purchases announced	QE
12/1/08	Chair mentions possibility of Treasury purchases	QE
12/16/08	FOMC cuts short rate to ELB, mentions possibility of Treasury purchases	Both
3/18/09	FOMC says it expects low rates for "extended period," expands existing QE and announces it will buy Treasuries	Both
11/4/09	Tapering of QE1 announced	QE
8/10/10	Reinvestments of principal and coupon payments announced	QE
8/27/10	Chair says more QE if it "proves necessary"	QE
9/21/10	Reinvestments of principal and coupon payments will continue	QE
11/3/10	QE2 announced	QE
8/9/11	Calendar-based forward guidance adopted (short rate will likely remain low through mid-2013)	FG
8/26/11	Bernanke Jackson Hole speech interpreted as signaling additional accommodation	Both
9/21/11	Maturity Extension Program announced	QE
1/25/12	Guidance extended to late-2014	FG
6/20/12	MEP will continue through 2012	QE
8/22/12	FOMC statement says additional accommodation "would likely be warranted fairly soon"	Both
9/13/12	FOMC introduces QE3, extends short-rate guidance to mid-2015	Both
12/12/12	FOMC extends QE3 and introduces unemployment-threshold-based FG	Both
5/22/13	Bernanke testimony hints at end to QE3	QE
6/19/13	Bernanke press conference suggests liftoff in 2015	FG
12/18/13	Tapering of QE3 announced	QE
3/19/14	Switch to qualitative guidance	FG
12/17/14	FOMC announces it can be "patient" in removing accommodation	FG
3/18/15	Removes "patient" language	FG
9/17/15	Liftoff expected but does not occur	FG

Notes: The table reports the 29 days used to construct the empirical event study on the effect of unconventional-policy announcements. The final column reports whether each event was associated with a change in asset-purchase policy (QE), forward guidance (FG), or both.

Table 5. Decomposition of yield responses to unconventional policy shocks

Quarterly installments

Maturity	Shadow-rate shocks		Bond-supply shocks (term prem.)	Interaction	Total
	Expectations component	Term premium			
2 years	-0.94	-0.21	-0.16	+0.06	-1.25
5 years	-1.20	-0.37	-0.38	+0.08	-1.87
10 years	-1.11	-0.40	-0.60	+0.06	-2.06
20 years	-0.84	-0.33	-0.69	+0.02	-1.84

All at once

Maturity	Shadow-rate shocks		Bond-supply shocks (term prem.)	Interaction	Total
	Expectations component	Term premium			
2 years	-0.52	-0.18	-0.08	+0.07	-0.71
5 years	-0.92	-0.44	-0.23	+0.13	-1.45
10 years	-1.04	-0.56	-0.39	+0.12	-1.87
20 years	-0.89	-0.49	-0.47	+0.07	-1.78

Notes: The table reports the outcome of model simulations that approximate the size of the unconventional-policy shocks introduced by the Federal Reserve during the period at which the short rate was constrained by the ELB (2008:4 – 2015:4). In the top panel, the policy interventions are modeled as a series of shocks to the shadow rate and bond supply that occur in equal amounts for 28 consecutive quarters, and the response reported in the table is the sum of the quarterly responses (analogous to the way the responses are constructed in the empirical event study). In the second panel, the shocks are modeled single, large shocks, and the table reports the initial response. The sizes of the shocks are calculated as described in the text, and their total effect on yields of each maturity are shown in the last column. In each case, the bond supply and shadow-rate shocks are simulated both separately and together to obtain the decompositions 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 the joint simulation. For the shadow-rate shocks, the expectations component is calculated from equation (11), while the term-premium component is calculated as the difference between the total change in yields and the expectations component.

Figure 1. Conditional moments of short rate at various horizons
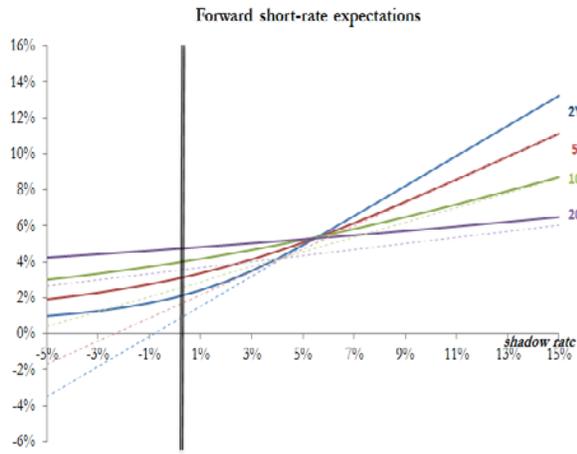
A.   B.
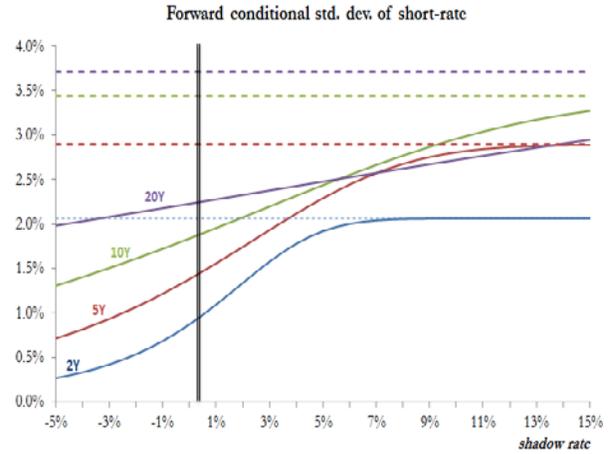
Notes: The figure shows the conditional mean (panel A) and standard deviation (panel B) of the time $t+h$ short rate, conditional on the value of the shadow rate in time t, where $h = 2, 5, 10,$ and 20 years. The solid lines show these moments in the shadow-rate model, under the calibration shown in the top line of Table 1. The dashed lines show the moments in an affine model with the same calibration but with the ELB removed.

Figure 2. Bond supply process
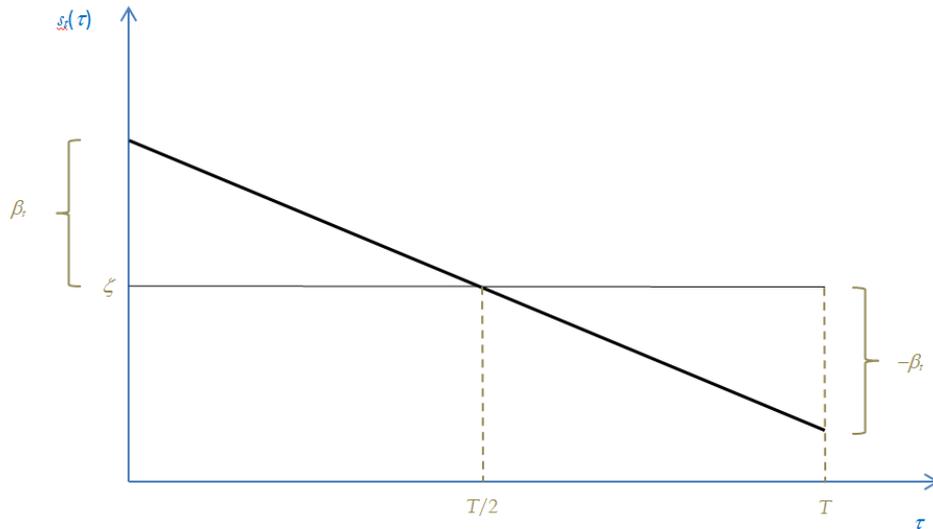
Notes: The figure depicts how the maturity distribution of bond supply in the model relates to the supply factor β_t. The horizontal axis corresponds to maturity, while the vertical axis is the amount of supply outstanding at each maturity. If the supply factor is at its average value, bond supply is equal to ζ at all maturities. Positive shocks to β_t tilt the supply distribution toward shorter maturities. This process is adapted from Greenwood, Hanson, and Vayanos (2015).

Figure 3. Average duration of Treasury securities held by the public
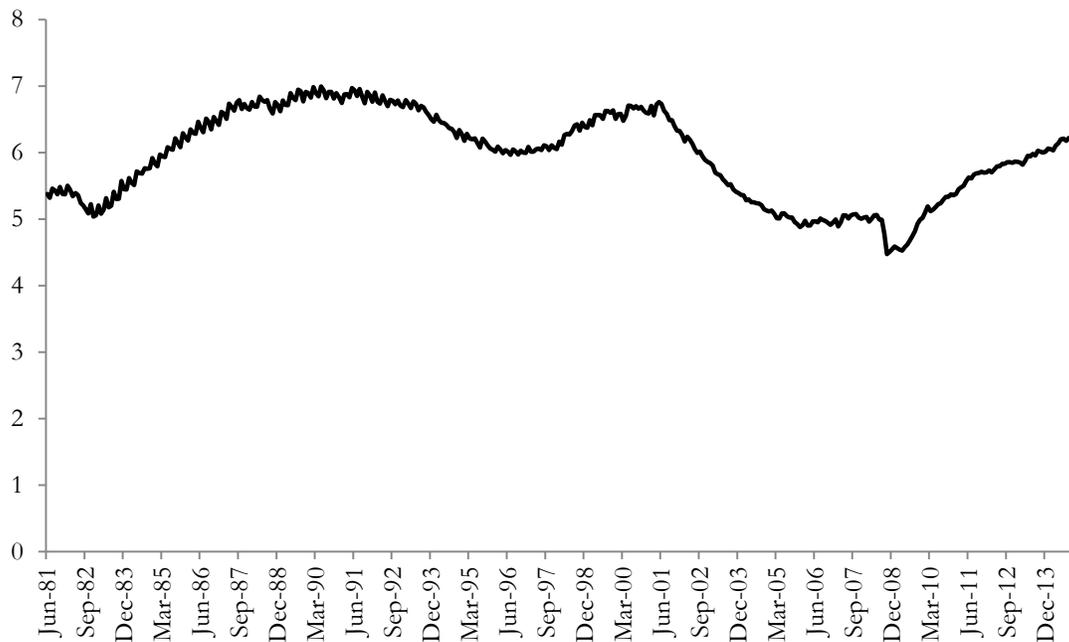
Notes: The figure shows the average duration of outstanding Treasury debt, in years, accounting for both principal and coupon payments. Callable bonds are excluded. Source: CRSP®, Center for Research in Security Prices, Booth School of Business, The University of Chicago. Used with permission. All rights reserved. crsp.uchicago.edu.
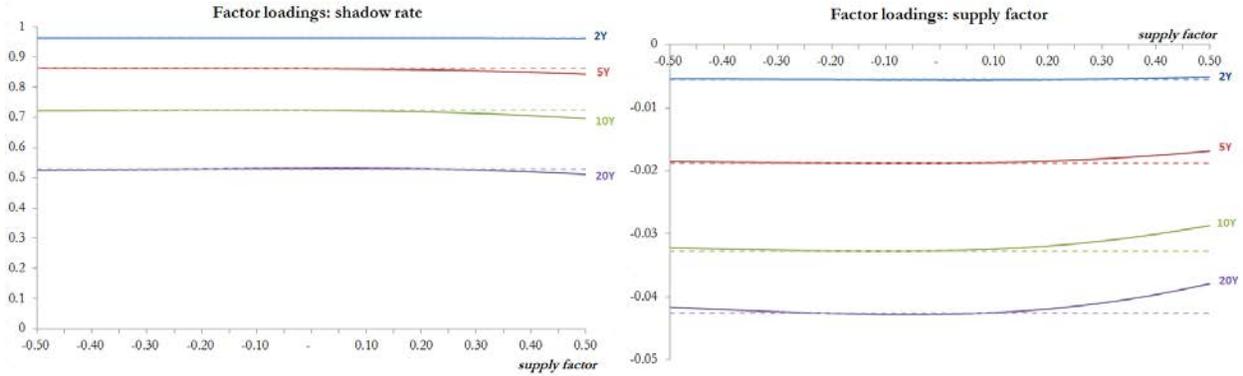
Figure 4. Factor loadings across values of the shadow rate at $\beta_t = 0$
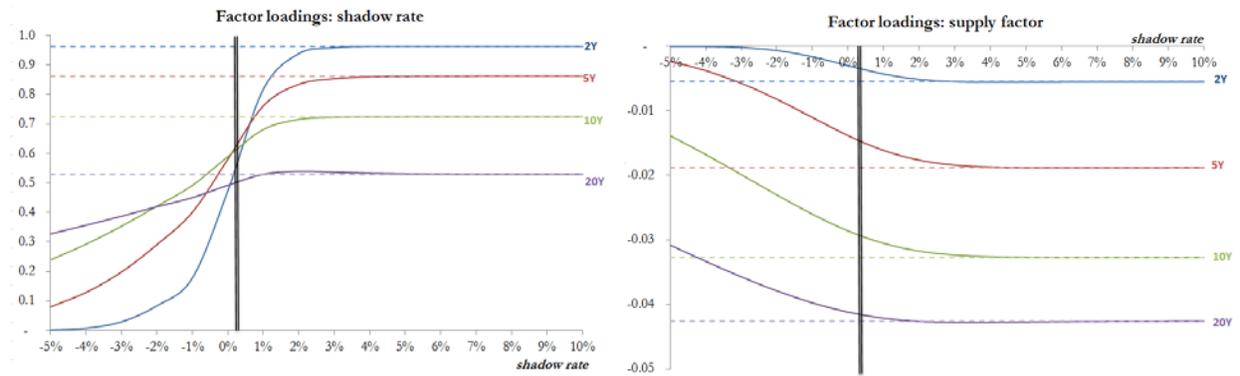
Notes: The figure shows the model-implied factor loadings for the h-period yield, conditional on the time-t value of the shadow rate, where h = 2, 5, 10, and 20 years. The solid lines show loadings in the shadow-rate model, under the calibration shown in the top line of Table 1. The dashed lines show the loadings in an affine model with the same calibration but with the ELB removed. The bond-supply factor β_t is held fixed at its mean of zero.

Figure 5. Factor loadings across values of the supply factor

A. $\hat{r}_t = 5.1\%$


B. $\hat{r}_t = -2.7\%$
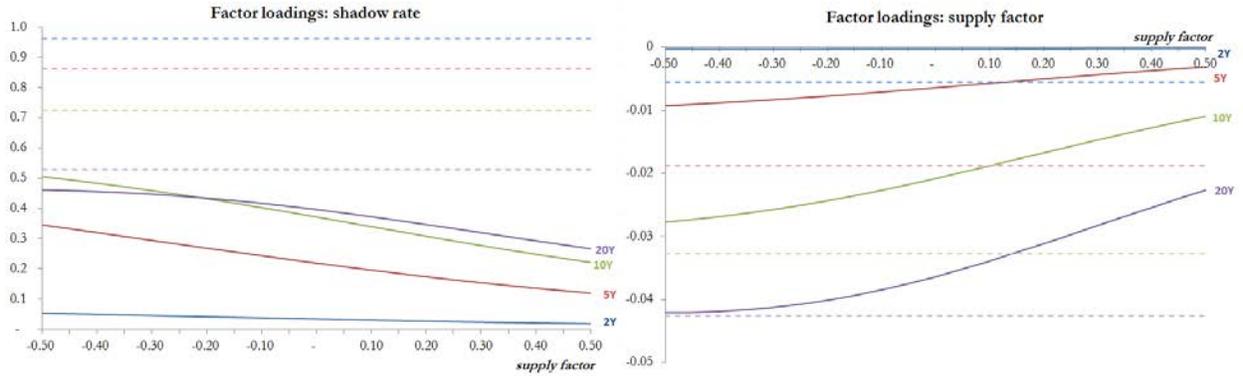
Notes: The figure shows the model-implied factor loadings for the h-period yield, across time-t values of the bond-supply factor, where $h = 2, 5, 10,$ and 20 years. The solid lines show the loadings in the shadow-rate model, under the calibration shown in the top line of Table 1. The dashed lines show the loadings in an affine model with the same calibration but with the ELB removed. In panel A, the shadow rate is held fixed at the mean value of the short rate (5.1%), while in panel B it is held fixed at a value of -2.7%, its average during the ELB period according to the Krippner (2012) estimates.

Figure 6. Impulse-Response Functions

A. Spot Yield Curve

B. Forward Rate Curve
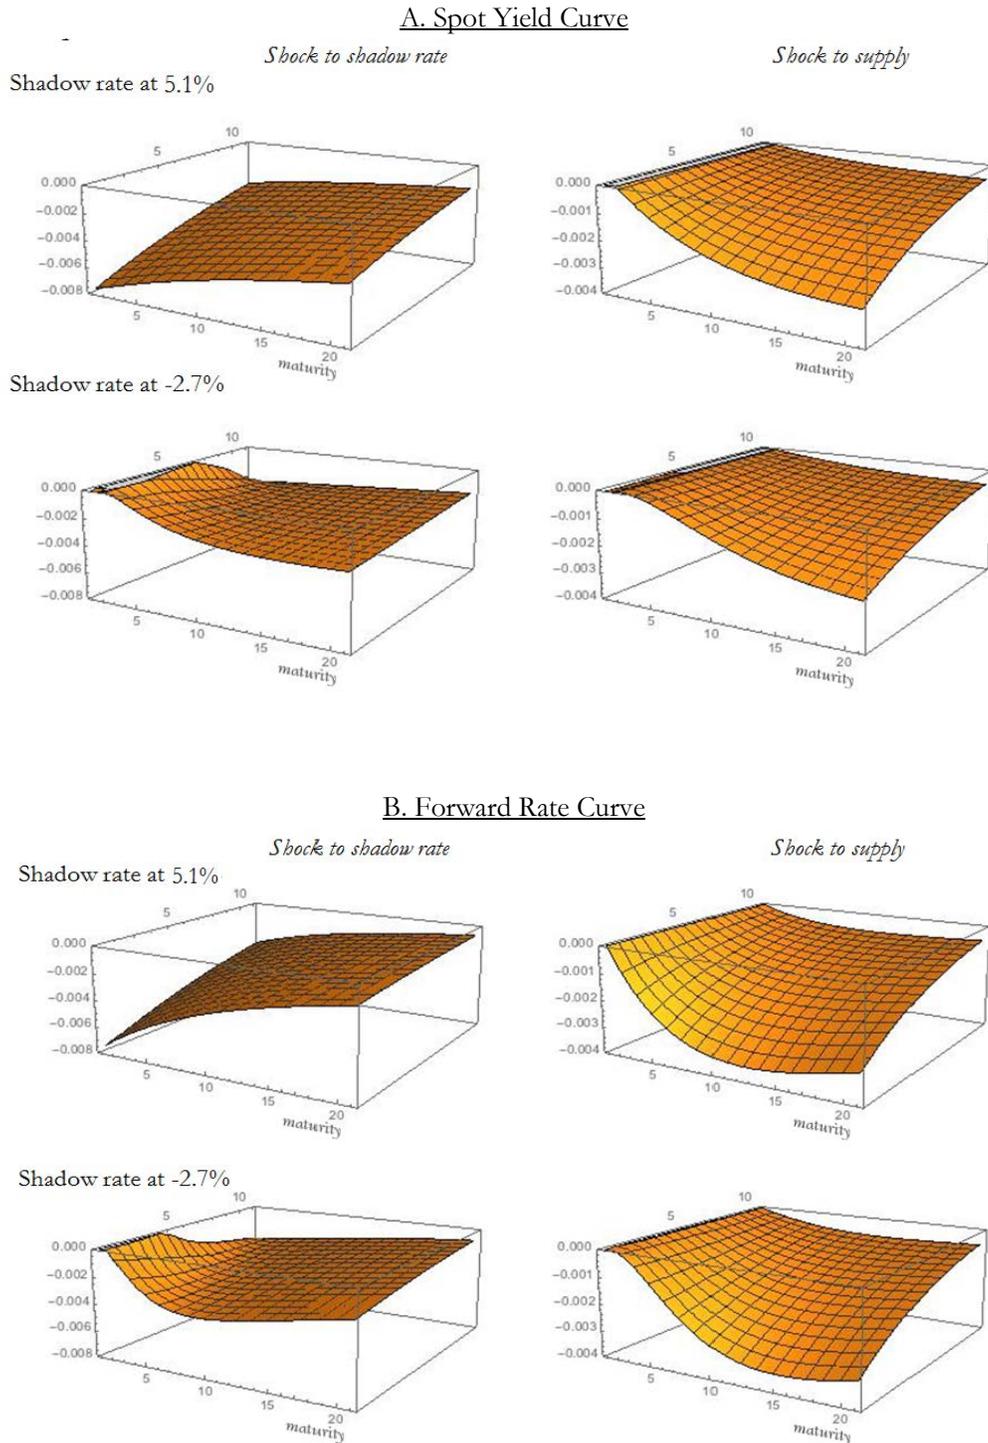
Notes: The figure shows model-implied responses of yields (panel A) and forward rates (panel B) to one-standard-deviation shocks to each of the factors over the subsequent 10 years. Maturity in years is shown on the lower-left axis in each graph, while calendar time is on the upper-left axis. Responses are evaluated starting both from a shadow rate at the mean value of the short rate (5.1%) and a value of -2.7%, its average during the ELB period according to the Krippner (2012) estimates. In both cases, the starting value of the bond-supply factor is zero.
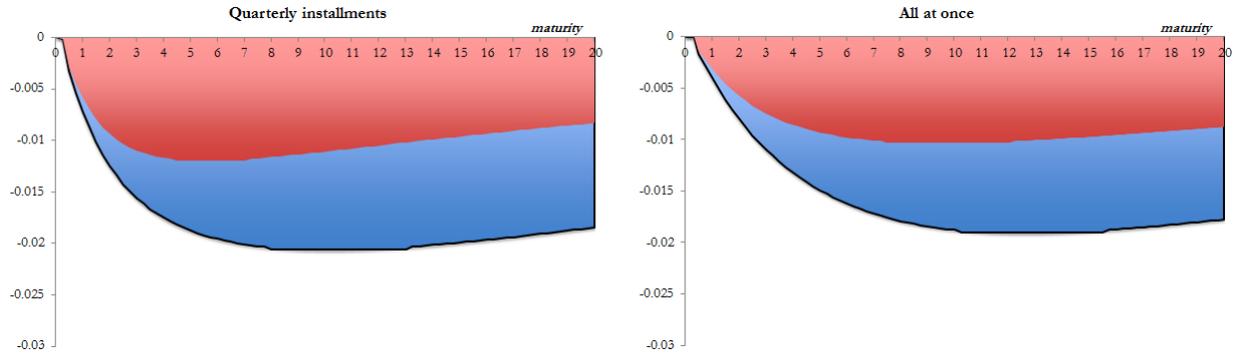
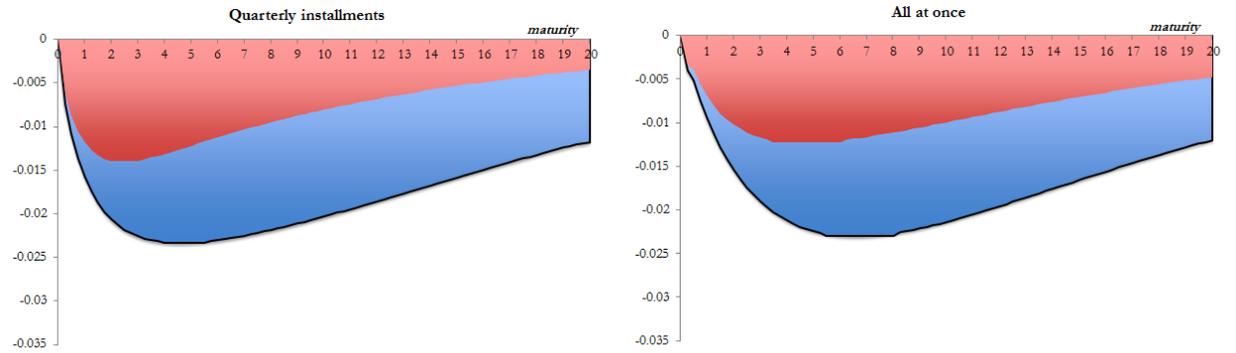Figure 7. Yield curve response in empirical event study
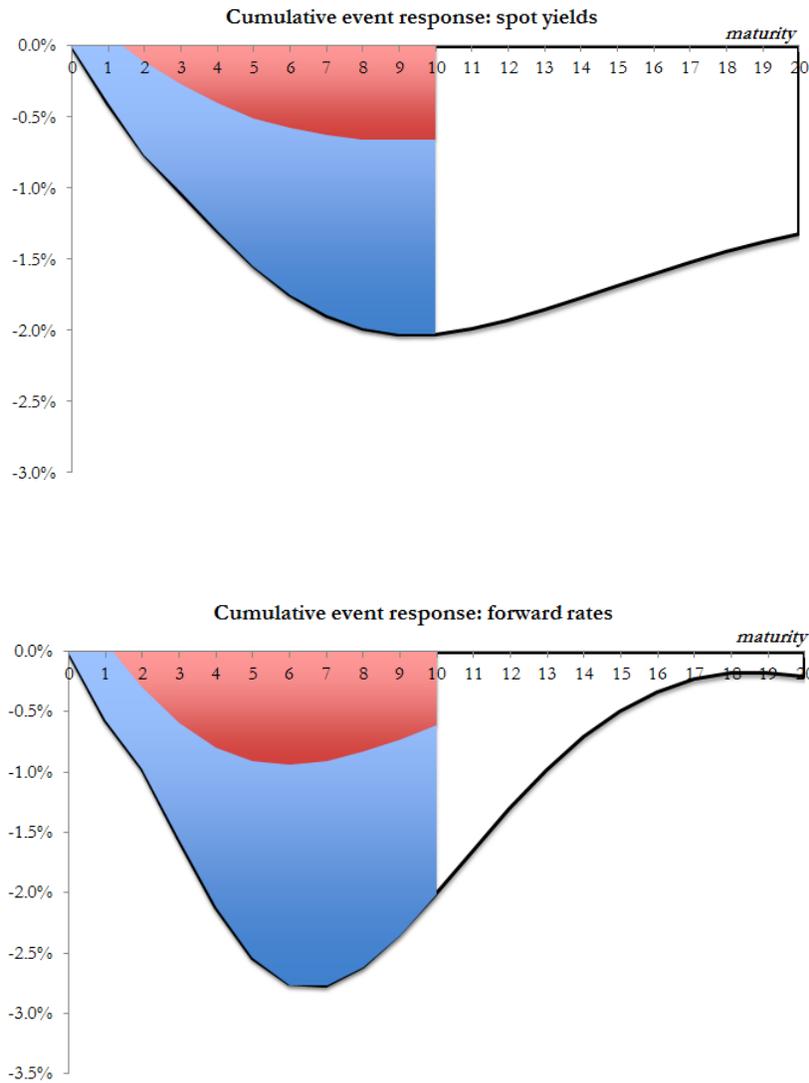
Notes: The figure shows the cumulative empirical response of zero-coupon yields and forward rates to various events related to the Federal Reserve's unconventional monetary policy. Responses are calculated based on two-day windows, using the Gurkaynak, Sack, and Wright (2007) data maintained on the Federal Reserve Board's website. The shaded regions decompose the total response into the term premium (blue) and expectations component (red), based on the Kim-Wright decomposition through the ten-year maturity, also available from the Federal Reserve Board. See text for additional details.

## Figure 8. Yield curve response in model-simulated event studies

Spot rates

Forward rates

Notes: The figure shows the cumulative response of zero-coupon yields and forward rates in shadow-rate model simulations that approximate the size of the unconventional-policy shocks introduced by the Federal Reserve during the period at which the short rate was constrained by the ELB (2008:4 – 2015:4). In the graphs labeled "quarterly installments," the policy interventions are modeled as a series of shocks to the shadow rate and bond supply that occur in equal amounts for 28 consecutive quarters, and the response shown in the figure is the sum of the quarterly responses (analogous to the way the responses are constructed in the empirical event study). In the graphs labeled "all at once," the shocks are modeled single, large shocks, and the figure shows the initial response. The sizes of the shocks are calculated as described in the text. The shaded regions decompose the total response into the term premium (blue) and expectations component (red). The expectations component is calculated from equation (11), while the term-premium component is calculated as the difference between the total change in yields and the expectations component.

Figure 9. Relative Efficacy of Bond-Supply Shocks across State Values
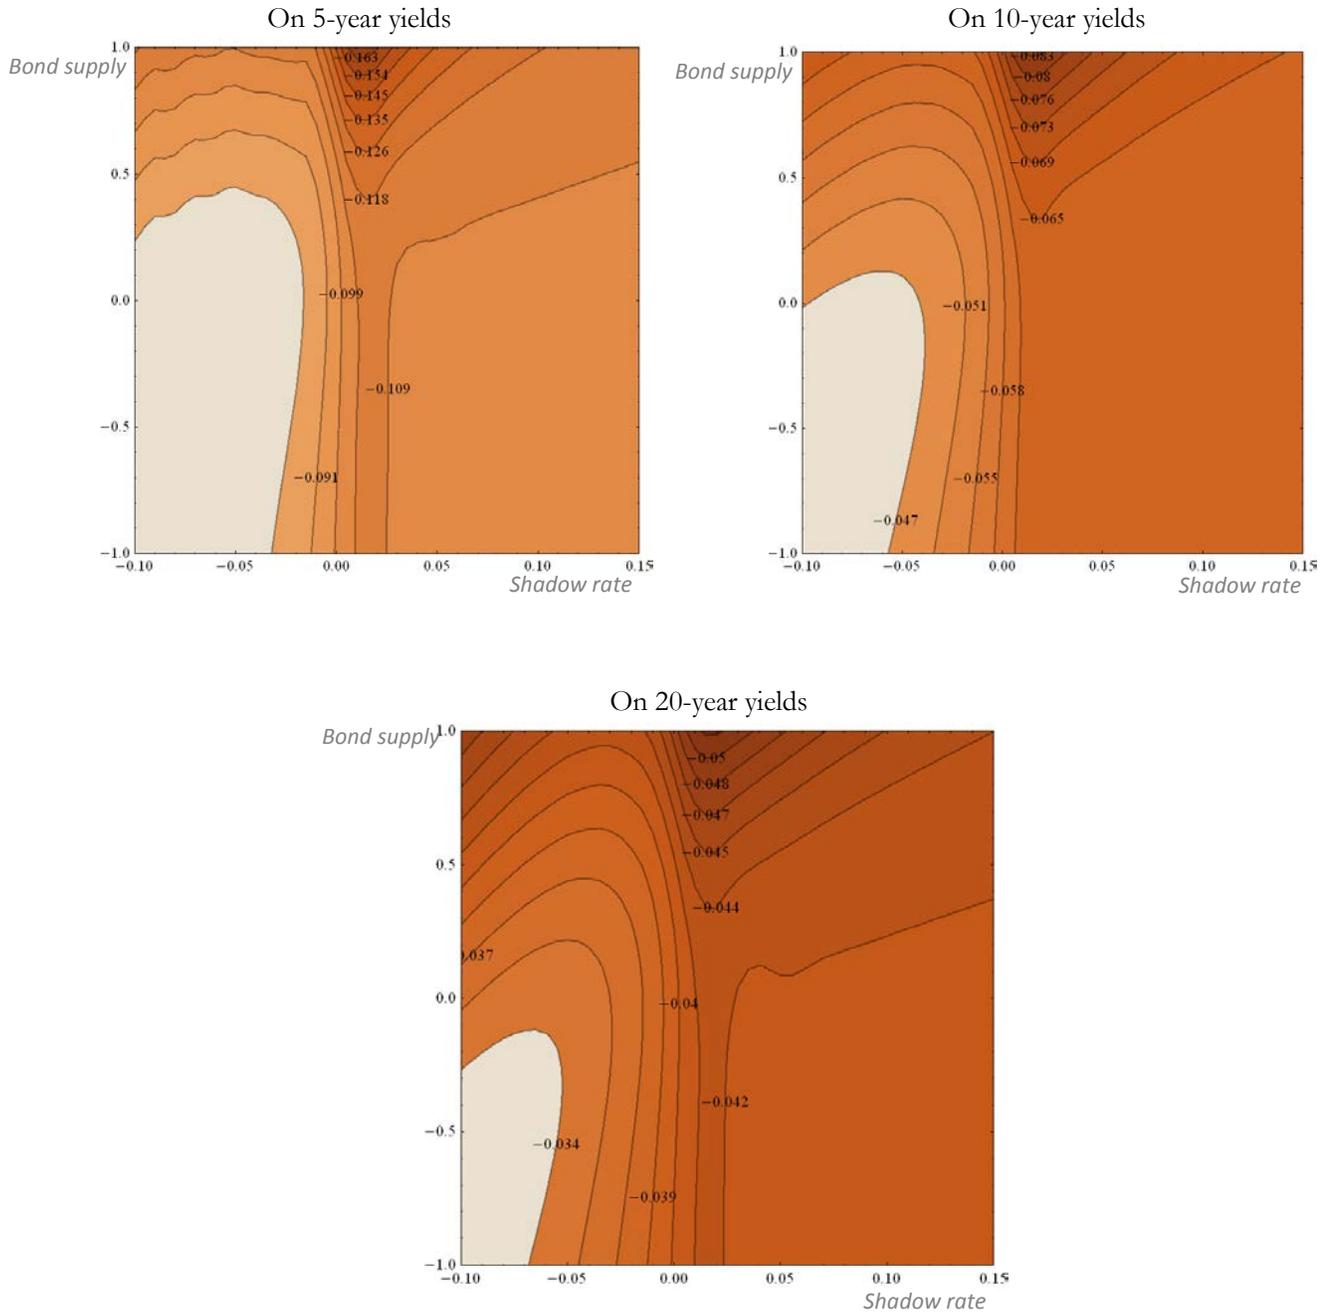
Notes: The figure shows contour maps of the effects of bond-supply shocks on various yields, relative to the effects of shadow-rate shocks. Relative efficacy is calculated, for each yield, as the size of the bond-supply shock that would be necessary to equal the effects of a -25-basis-point shock to the shadow rate. The values of this ratio are shown across different regions of the state space, with lighter-colored values indicating regions where the bond-supply shocks are relatively more powerful.