

# How Likely Is the Zero Lower Bound?

---

Thomas A. Lubik and Christian Matthes

**D**uring the course of the Great Recession and for long after, the Federal Reserve kept the main monetary policy rate at the zero lower bound (ZLB).<sup>1,2</sup> This policy was pursued in order to fight the deepest recession since the Great Depression and to support the budding recovery. The Federal Reserve finally abandoned its low interest rate policy and exited from the ZLB in December 2015 as the expansion gathered pace. It is now one of the longest on record in US economic history. Yet, given the length of the expansion and its recent strength, the level of the policy rate is arguably still low when compared with the historical experience.<sup>3</sup>

---

■ We are grateful to Caroline Davis for outstanding research assistance and to Felix Ackon, James Geary, John Jones, and John Weinberg, whose comments are greatly appreciated. The views expressed in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System. An updated version of this paper, with a correction to footnote 1, was posted on October 24, 2019. DOI: <https://doi.org/10.21144/eq1050102>

<sup>1</sup> Strictly speaking, the Federal Reserve's key policy rate, the federal funds rate, was maintained in a range between zero and twenty-five basis points and never was actually at zero. Other major central banks, such as the European Central Bank, the Bank of Japan, and the Swedish Riksbank set policy rates to zero or even negative values. The label ZLB is thus shorthand for rates that are effectively zero.

<sup>2</sup> The ZLB is often taken to coincide with the ELB, or effective lower bound, but experience has shown that nominal policy rates can be negative for extended periods. For instance, the European Central Bank and the Swedish Riksbank have maintained negative rates on excess deposits held willingly by banks in their reserve accounts. In that sense, the ELB is below the ZLB, but there is no consensus in the economics profession on how low nominal policy rates could go. For the purpose of this article, we assume, however, that the ELB and the ZLB coincide, since the Federal Reserve is unlikely to consider negative policy rates.

<sup>3</sup> One explanation of why the nominal policy rate is low is that with inflation expectations anchored at the Federal Reserve's target of 2 percent, the natural real rate of interest is lower than in prior expansions. There is substantial evidence (e.g., Lubik and Matthes 2015b; Laubach and Williams 2016) that the natural rate exhibits secular decline over the last thirty years, which limits how high the equilibrium policy rate can go.

Naturally, this raises the question of how likely it is that monetary policy will again be subject to the ZLB in the coming years. More specifically, policymakers may wonder when the current expansion might end and whether they may have to pursue accommodative policies, possibly in a preemptive manner. During contractions, the Fed has traditionally lowered policy rates and kept them low to support the recovery. However, when policy rates are already low during the expansionary phase, as is currently the case, this potentially limits the ability of the Fed to provide accommodation because of the presence of the ZLB. This might arguably put the Fed in a dilemma regarding the pace of interest rate increases.

On the one hand, the Fed could raise rates faster than is usually warranted in order to create more distance from the ZLB as insurance against the possibility that it might have to lower rates significantly to stem against a contraction. On the other hand, the Fed could go slower so as not to endanger a budding recovery and not face a contraction at all. At the same time, the likelihood of ZLB episodes is also central to debates about whether the Fed should replace its current 2 percent inflation target with either a higher target or a different framework. But all of these discussions center around the idea of insurance against being too close to the ZLB.

In this article, we therefore investigate the likelihood that the economy may be subject to the ZLB again. In this sense, we provide a quantification of the insurance aspect against the ZLB in terms of a forecast of that uncertainty. Specifically, we focus on the forecasting framework surveyed by Lubik and Matthes (2015a), which was previously applied to estimation of the natural rate of interest (Lubik and Matthes 2015b). We specify a time-varying parameter vector autoregressive model (TVP-VAR) for a set of key macroeconomic variables and estimate the model on the available data. Interpreting the model as an acceptable representation of the underlying structure of the economy and the effects of monetary policy, we then simulate the estimated TVP-VAR forward based on the estimated posterior distribution of its parameters. This generates a distribution of trajectories for macroeconomic outcomes, including the path of the federal funds rate (FFR), which we use as a monetary policy variable. From this distribution, we can then compute the probability that the interest rate will be at the ZLB in the future.

Our main result is that the probability of the ZLB is negligible over the next two years. It is only during 2022 that the probability rises above 5 percent. Depending on the specific interpretation of the ZLB probability, whether it is date-specific or horizon-specific, the probability rises at most to 15 percent by the end of the 2020s. While these

numbers are not negligible, they do not appear large enough to cause undue alarm. We also find that the ZLB probabilities have declined over the last three quarters. Our main findings are based on data up to and including the third quarter of 2013, which saw robust growth. The same exercise with a sample ending in 2018Q1 yields longer-term ZLB estimates of around 25 percent while still being negligible at a short horizon. The strength of the recent data flow thus makes the ZLB less likely since the model incorporates the possibility of an ever-so-slight trend GDP growth shift.

As a robustness check and an assessment of the overall validity of our forecasting model, we also investigate how well the TVP-VAR has performed in the past, specifically during the Great Recession. To do so, we conduct a pseudo out-of-sample exercise where we carry out our forecasting exercise while conditioning on the parameter estimates at that point in time. The ZLB probabilities are computed in a like manner, that is, as if the subsequent data were unknown to the researcher. We find that right at the onset of the downturn, the model predicts the ZLB with a very high probability of 80 percent on account of the dramatic decline in real GDP growth. Going further out, however, and as the subsample expands, this probability drops to below 40 percent, even as the policy rate remains at the ZLB and has been there for a while.

This observation reveals a feature of the data that even a flexible nonlinear model such as the TVP-VAR has difficulty dealing with, namely reversion to the mean. In other words, the Great Recession and the subsequent ZLB period are such unusual events that the model has a tendency to discount their impact going forward. At best, this feature of the data is reflected in a wider distribution of the forward simulation that underlies the computation of the ZLB probabilities. In that sense, a shift in the ZLB probabilities observed in 2018 can be seen as evidence of an underlying trend shift.

Perhaps surprisingly, the question of how likely the ZLB is has not attracted much attention in the empirical macroeconomics literature, specifically the forecasting literature. While there is much research on the effects of the ZLB in traditional New Keynesian models, these studies are not forecast-based but instead study the probability of being at the ZLB generically. Perhaps closest to our exercise is Chung et al. (2012), who use several forecasting models used in the policy process, such as the Federal Reserve's own large-scale macroeconometric model FRB/US, two canonical New Keynesian dynamic stochastic general equilibrium (DSGE) models, and also a TVP-VAR closely related to ours. In a similar stochastic simulation exercise, they construct forecast densities based on data up to and including 2007Q4. None of the

models, perhaps surprisingly least of all the TVP-VAR, include the ZLB in their 95 percent coverage region, which echoes some of our findings. However, they focus on this one base year only, whereas we compute densities for forty quarters out and also conduct a model validation exercise.

In a more recent study, Kiley and Roberts (2017) simulate both the FRB/US policy model and a standard DSGE model often used in the policy process with shocks drawn from estimated distributions over the 2000–15 period. They find that the ZLB probabilities are small, reaching at best 20 percent for levels of the natural rate at 3 percent, which is consistent with the natural rate forecast embedded in our model.<sup>4</sup> In a pseudo out-of-sample exercise for 2013 that is similar to ours, they also have the feature of mean reversion unless they strongly force the policy rule to follow the ZLB.<sup>5</sup>

We proceed as follows. In the next section, we discuss our empirical approach in more detail. We introduce our statistical forecasting model, a TVP-VAR, and then describe two alternative measures of the ZLB probabilities. Section 3 contains the results of the paper, including a pseudo out-of-sample exercise to assess the quality of the forecasting model. The final section concludes.

## 1. METHODOLOGY

We estimate the probability that the federal funds rate will be at or below the ZLB from a statistical model of the US economy. The first step of our analysis is therefore to develop a model that captures the behavior of key macroeconomic variables well, especially during previous ZLB episodes. Since an assessment of such probability involves a forecast, a desirable property of the statistical model is a good forecasting performance. For this purpose we use a TVP-VAR, which has become widely used for economic and policy analysis and is a flexible framework to address the kinds of issues discussed in this article.<sup>6</sup>

The advantage of a TVP-VAR is that it is a largely atheoretic time-series model, which absolves the researcher from taking a stand on the

---

<sup>4</sup> See Lubik and Matthes (2015b) and updates thereof at: [https://www.richmondfed.org/research/data\\_analysis](https://www.richmondfed.org/research/data_analysis).

<sup>5</sup> Jones (2017) finds similar results in a fully estimated DSGE model that accounts for possible trend breaks associated with the secular decline in the natural real rate of interest. He is able to match the data with a forward guidance policy rule.

<sup>6</sup> Doh and Connolly (2012) and Lubik and Matthes (2015a) provide an overview of the methodology and a step-by-step guide to its implementation. Examples of its use in the monetary policy process are discussed in Clark and Ravazzolo (2015) and Lubik and Matthes (2015b), while Canova and Gambetti (2009) and Lubik et al. (2016) detail some of its limitations.

deep, underlying relationships that govern the joint behavior of aggregate variables. Perhaps more importantly, a TVP-VAR can in principle capture nonlinear behavior in the underlying time series, such as the ZLB, where movements in the interest rate are capped by a lower bound of 0 percent, without specifying the precise source of the nonlinearity. TVP-VARs have also proved useful in forecasting because they allow researchers to distinguish between structural or long-lasting changes in the economy and shorter-term fluctuations in a consistent and transparent manner. The former affect trends and forecasts thereof, while the latter are often driven by changes in the volatility of shocks hitting the economy. Allowing for time variation in both elements of the model helps researchers differentiate these sources of aggregate fluctuations.

### A TVP-VAR for the US Economy

We specify a TVP-VAR in quarterly data on real GDP growth, PCE inflation, and the federal funds rate, which are collected in a column vector  $y_t$ . We assume that the joint evolution of these variables is governed by the law of motion:

$$y_t = \mu_t + \sum_{j=1}^2 A_{j,t} y_{t-j} + e_t. \quad (1)$$

$\mu_t$  is a drift term that can contain deterministic and stochastic components. It is of particular importance for capturing the changing trends in the variables, such as the decline of GDP growth over the last fifty years or the ZLB, which can be regarded as trend break in this context. The  $A_{j,t}$  are conformable coefficient matrices that contain time-varying parameters, the evolution of which we detail below.  $e_t$  is a vector of residuals. We set the lag length equal to two, which is standard for quarterly data in the TVP-VAR literature (see Primiceri 2005). We can define  $X_t' \equiv I \otimes (1, y_{t-1}', \dots, y_{t-2}')$  to provide a compact representation of the dynamics of  $y_t$ . We then rewrite equation (1) as:

$$y_t = X_t' \theta_t + e_t. \quad (2)$$

We assume that the law of motion for the time-varying parameters in the coefficient matrices  $A_{j,t}$  is given by a random walk process:

$$\theta_t = \theta_{t-1} + u_t, \quad (3)$$

where  $u_t$  is a zero mean i.i.d. Gaussian process. We model the process for stochastic volatility by assuming that the covariance matrix of the one-step-ahead forecast error  $e_t$  can be decomposed as follows:

$$e_t = \Lambda_t^{-1} \Sigma_t \varepsilon_t, \quad (4)$$

where the standardized residuals are distributed as  $\varepsilon_t \sim N(0, I)$ .  $\Lambda_t$  is a lower triangular matrix with ones on the main diagonal and representative nonfixed element  $\lambda_t^i$ .  $\Sigma_t$  is a diagonal matrix with representative nonfixed element  $\sigma_t^j$ . The dynamics of the nonfixed elements of  $\Lambda_t$  and  $\Sigma_t$  are given by:

$$\lambda_t^i = \lambda_{t-1}^i + \zeta_t^i. \quad (5)$$

$$\log \sigma_t^j = \log \sigma_{t-1}^j + \eta_t^j. \quad (6)$$

We assume that all these innovations are normally distributed with covariance matrix  $V$ . In order to provide some structure for the estimation, we restrict the joint behavior of the innovations as follows (following Primiceri 2005):

$$V = Var \left[ \begin{pmatrix} \varepsilon_t \\ u_t \\ \zeta_t \\ \eta_t \end{pmatrix} \right] = \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix}. \quad (7)$$

$S$  is further restricted to be block diagonal, which simplifies inference. We use a Gibbs-sampling algorithm to generate draws from the posterior. The implementation of the Gibbs-sampling approach used for Bayesian inference follows Del Negro and Primiceri (2015) and is also described in more detail in Lubik and Matthes (2015a).

A key choice for TVP-VAR modeling is how to set the prior. In order to achieve sharp inference, given the multiple sources of variation in TVP-VAR models, a researcher needs to impose restrictions on the relationship between the covariance matrices of the parameters. The trade-off, however, is that a too-restrictive prior may not leave room for the time-variation to appear. In our benchmark, we impose a typical choice of prior as recommended in Primiceri (2005). Specifically, we assume the following:

$$Q \sim IW(\kappa_Q^2 * 40 * V(\theta_{OLS}), 40), \quad (8)$$

$$W \sim IW(\kappa_W^2 * 2 * I, 2), \quad (9)$$

$$S \sim IW(\kappa_S^2 * 2 * V(\Lambda_{OLS}), 2), \quad (10)$$

where  $IW$  denotes the Inverted Wishart distribution. Priors for all other parameters are the same as in Primiceri (2005). For the prior hyperparameters  $\kappa_Q, \kappa_W$ , and  $\kappa_S$ , we use the values  $\kappa_Q = 0.01$ ,  $\kappa_W = 0.01$ , and  $\kappa_S = 0.1$ .

### Computing ZLB Probabilities

The probability that the economy will reach or fall below the ZLB in the future is based on a forecast of the joint evolution of the variables in the

statistical model. This is not a point forecast but rather a collection of forecasts that detail all likely paths the economy will take given where it is now. The ZLB probability then simply captures how many times the interest rate will be at or below zero. In order to operationalize this idea, we proceed as follows. In the first step, we estimate the TVP-VAR over the entire sample period. Our posterior sampler delivers the posterior distribution of parameters for any point in the sample, which we will exploit later. We then fix the coefficients at their last estimated posterior mean and keep them fixed over the forecast horizon. This assumption is made for computational expediency as it does not require simulating paths of parameters. Recall that all coefficients in the model are varying over time, including the trends, the lag coefficients, and the parameters governing the volatility of the shocks. In the simulation exercise we do not draw from the innovation distributions of the TVP-VAR parameters as this would add an additional layer of uncertainty.

This approach is consistent with the idea of an unchanged forecast where the structure of the economy is not expected to change. This assumption seems reasonable as a baseline, especially in light of the fact that we model the evolution of parameters as random walks. In addition, it is a well-known drawback of TVP-VARs that because of all the moving parts uncertainty about forecasts is generally higher. Forecasting the paths of parameters as well would thus just compound this uncertainty. In that sense, there would be too much parameter variation to make the forecasts meaningful, or alternatively, current conditions would be uninformative about the future. We therefore choose to err on the side of sharper predictions. What we might miss are, at lower frequencies, changes in trend growth, as we have seen over the last decades for real GDP and the real rate of interest, and periods of excess volatility, such as we have seen during deep recessions and financial crises. At the same time, such events are notoriously difficult to forecast. We show an example of this and its implications for our exercise below.

Given this structure of the TVP-VAR, we produce forecasts over a ten-year horizon. The forecasts are such that for each future date we generate realizations of the shocks hitting the economy; that is, we draw from the estimated distribution of the innovations to the exogenous processes and record how they propagate through the economy. This generates sample trajectories of the model's endogenous variables that can be collected at every point in time as a distribution of likely outcomes. From this collection of sample paths, we can then compute the probability that the interest rate will be at or below the ZLB.

In principle, one can think of two alternative measures. The first gives an answer to the question: What is the probability that at a given

point in time the FFR is forecast to be at or below zero? We measure this by counting how many times the interest rate is subject to the ZLB at a given date under all simulated trajectories. The ZLB probability is then found by dividing this count by the total number of simulations at the specific point in time. We label this measure ‘unconditional’ as it represents the marginal probability of being at the ZLB at a given time period. In terms of our forecasting exercise, it is a simple count of ZLB events at every point in time, normalized by the total number of forecast paths.

We also consider an alternative measure that we label ‘conditional.’ This measure represents the probability that the economy has been at the ZLB at least once up to and including the current period. It thereby takes into account the dependency of the ZLB episodes. As the forecast horizon increases, the count is accumulated. For instance, consider a trajectory of the FFR that is below zero in period  $t + 1$  and period  $t + 2$  and is above zero in period  $t + 3$ . For the first, unconditional measure, we record a count of one, one, and zero, since the measure focuses on the ZLB episodes in any period. For the second, conditional measure, the count is one, one, one, since this hypothetical path features two incidences of a ZLB episode. In this case, the trajectory still enters the ZLB count in the last period since the given incidence of shocks resulted in a ZLB episode in prior periods and thus contributed to the overall “risk” of the ZLB. In this sense, it is a cumulative probability measure for the question at hand. From a policymaker’s point of view, it conveys the information that even if a trajectory is not subject to the ZLB at a particular date, it may have been so in prior periods and may therefore have to be avoided. Naturally, the unconditional measure is bounded from above by the conditional measure.

## **2. THE PROBABILITY OF BEING AT THE ZLB**

We report the key findings of this article in Figures 1 and 2, which report the two measures of the ZLB probabilities discussed above. We consider a forecast horizon of ten years at a quarterly frequency. The figures show the ZLB probabilities for three sample periods each, namely ending in 2018Q1, 2018Q2, and 2018Q3, respectively, but with the same start date, 1963Q1. The samples are real time in that we have used the data actually available to policymakers at that time. For the sample with the most recent data, up to and including 2018Q3, the figures show that for both measures the probability is essentially zero for one year out. It rises gradually toward a long-run level of 7 percent in the case of the unconditional measure in Figure 1 and around 13 percent for the alternative measure in Figure 2. As can be seen from



the figures, the ZLB probabilities are rising over time. This stems from the fact that uncertainty is expanding as we forecast further out into the future and that the conditional probability is cumulative.

As discussed above, the probabilities shown in Figure 1 give policymakers an unconditional view that the ZLB may occur again in the future. From this perspective it indicates that there is less than a one in ten chance that in 2028 the economy will be in a situation where the FFR is again constrained at zero. Since the ZLB has been observed in the dataset, the model deems it likely to happen again, given the estimated historical patterns of shocks when extrapolated forward. This is irrespective of whether any trajectory has been at the ZLB before or not. The estimates in Figure 2 show a similar pattern with virtually zero probability for one year from now. It is then rising to a long-run level of close to 15 percent. The interpretation of this conditional measure is that in 2028 roughly one-eighth of all forecast trajectories of the federal funds rate will have hit the ZLB at some point, either only once or repeatedly.

The ZLB probabilities for the sample ending in 2018Q2 are essentially identical to the most recent sample, with the latter's unconditional probability slightly higher in the first half of the sample but slightly lower in the second half. In contrast, the ZLB probabilities for the 2018Q1 sample are considerably higher, rising to a long-run level of 15 percent in the case of the unconditional measure and around 30 percent for the other measure. Still, for one year out, the probability is effectively zero. This shift in the estimated ZLB probabilities is driven by the strong GDP growth data in the second and third quarters of 2018, which imply a higher forecast FFR path and thus a larger distance from the ZLB for all trajectories. Moreover, the stronger growth data may lead the model to reevaluate the underlying properties of GDP, which also support a higher FFR path. In addition, a stronger economy in 2018 reduces the likelihood of a recession in the near term, thus reducing the ZLB probabilities.

This discussion raises the question of how reliable the estimates of the ZLB probabilities are. The quality of the estimates rests crucially on how well the model captures past experiences, including the ZLB period during 2009–14. The future is, of course, uncertain, but we can get a sense of how well the TVP-VAR has performed in the past by conducting a pseudo out-of-sample forecasting exercise. We proceed as follows. As a starting point, we use the baseline estimated posterior distribution of parameters at different points in time. The posterior estimates of parameters at any point in time are a function of all available data—data at time  $t + j, j \geq 1$  are generally informative about the parameter values in place at time  $t$ . Using this approach instead of

a true out-of-sample exercise, where we would have to reestimate the model period by period, is computationally much more tractable but is naturally subject to the caveat that it is based on information that could not have been known at that time.<sup>7</sup>

We then perform the same forecasting exercise as discussed above. At each date we simulate the model forward using the posterior estimates of the coefficients, which are held fixed over the forecast horizon. We produce the same counts as in the prior exercise, namely how many times the interest rate is at or below the ZLB for each quarter. Figure 3 shows results from this exercise. We focus on four forecast horizons: one quarter ahead, four, eight, and then twenty. The horizontal axis in the figure denotes the period in which the forecast is made, while each panel reports results for a specific forecast horizon. The four panels in the figure depict the unconditional ZLB probabilities at the respective forecast horizons as they change over time.

The upper left-hand panel shows the one-step-ahead ZLB probability. For almost the entire sample period the probability is zero on account of high interest rates (and high inflation during the 1970s). This changes in early 2009 as this probability shoots up to 80 percent with the onset of the Great Recession. This is driven by the sharp decline in real GDP growth, which, given historical patterns embedded in the estimated model, prompts a sharp drop in the interest rate. The one-step-ahead forecast at the next data point drops to below 50 percent and hovers around 20 percent until 2015 with the start of the exit from the quantitative easing period. The ZLB probabilities during this period are punctuated by occasional spikes that line up with weak data on GDP growth and low inflation numbers.

Nevertheless, these estimates indicate one weak point of the TVP-VAR, namely that it exhibits something akin to mean reversion.<sup>8</sup> Despite having observed policy rates at zero for several years, the TVP-VAR continues to predict an immediate rise in rates and hence a low

---

<sup>7</sup> Another caveat associated with this exercise is that we use final data for the estimates. This presumes knowledge that policymakers at that time could not have had, as initial data releases are typically subject to measurement errors and later revisions. The TVP-VAR estimates thereby do not reflect the actual decision-making environment that policymakers faced, which can result in biased estimates of the implicit policy rule. Lubik and Matthes (2016) show that policymaking under this type of data mismeasurement can considerably affect macroeconomic outcomes.

<sup>8</sup> To be clear, the TVP-VAR does not per se exhibit mean reversion as the coefficients are modeled as random walk processes and the underlying data are allowed to be nonstationary. In that sense, there is no ergodic distribution, but this does not rule out a proper posterior distribution. In addition, in our forward simulation exercise, we are holding the coefficients fixed at their last estimated value. It can be seen from Figures 1 and 2 that the ZLB probabilities seem to stabilize. It is in that sense that we apply the moniker mean reversion.

ZLB probability. As outside observers, the persistence of a ZLB policy is apparent, not least from Federal Reserve communication; yet it is not straightforward to capture this feature in a statistical model.<sup>9</sup> Overall, this serves as a caveat for the findings above, namely that the TVP-VAR does not fully capture the underlying dynamics in the data.

The other three panels in Figure 3 show similar patterns. There is a spike of the ZLB probabilities at the onset of the Great Recession, which then settle at a lower level before dropping to almost zero with the beginning of the tightening period. What differs across panels is the time horizon along which the probabilities are estimated. As the horizon expands from one quarter ahead to twenty quarters, the initial ZLB probability drops. For instance, at five years out, the model implies a 50 percent chance that the economy will be at the ZLB. In that sense, the TVP-VAR captures the underlying ZLB dynamics reasonably well as it incorporates the sharp interest rate drop to fight the downturn and ascribes persistence to it.

Moreover, the same panel shows that the slowdown in growth observed around the middle of the 2000s translates into an increased ZLB probability. In a sense, the TVP-VAR shows that there was information available before the actual onset of the Great Recession that put increased likelihood on very low interest rates at a five-year mark. Finally, the results for longer time horizons also show higher ZLB probabilities in the mid-1970s, the early 1980s, and the early 2000s. All three periods are characterized by large interest rate movements and higher macroeconomic volatility around and during recessions. This translates into more forecast uncertainty, which is reflected in higher ZLB probabilities.

### 3. CONCLUSION

This article discusses an approach to how policymakers can think about the risk of having to face the zero lower bound on the nominal interest rates. It is based on a forecasting model for the US economy that is flexible enough to capture nonlinearities such as the ZLB. Our main finding is that the probability of being at the ZLB is small but not insignificant for the US economy over a ten-year time horizon. Depending on how one interprets the notion of being at the ZLB, either over the course of a time path (our conditional measure) or as a pure point-in-time forecast (our unconditional measure), this likelihood is close to

---

<sup>9</sup> DSGE models have struggled with this feature of the data, too. See, for instance, Kiley and Roberts (2017). Jones (2017) does better than most by incorporating forward guidance explicitly into the policy rule.

one-fifth in the longer run, albeit it is close to zero over a shorter time horizon. A robustness exercise shows that our methodology is reasonably successful in capturing the probability of being at the ZLB before and during the Great Recession.

Our findings can inform the discussion on the pace of interest rate increases since the probability of being at the ZLB is inversely related to its distance from the interest rate, other things being equal. Whether a more aggressive path of hikes has the potential to tip the economy into recession, thereby stimulating interest rate cuts and an increased risk of the ZLB, depends on the underlying structure of the economy and its monetary transmission mechanism. It goes much beyond the scope of this article to assess whether the employed TVP-VAR is a good descriptor of this mechanism. Nevertheless, our estimated ZLB probabilities are also a convenient way of summarizing forecasts by condensing a lot of information into a single statistic. Overall, we regard the ZLB probabilities discussed in this article as a useful tool for policymakers to assess the current stance of monetary policy in light of the estimated likelihood that the policy will be constrained by the ZLB.

---

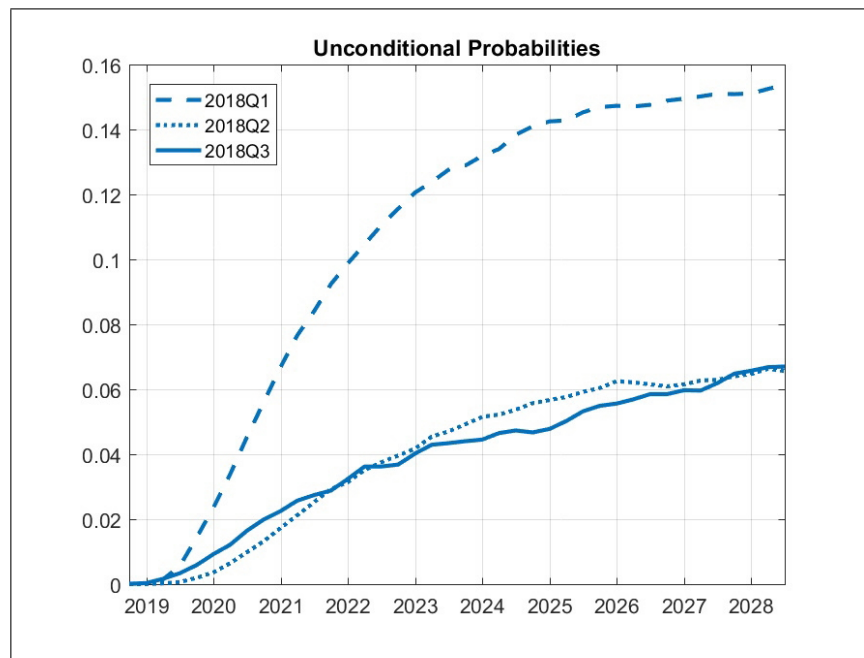
---

**REFERENCES**


- Canova, Fabio, and Luca Gambetti. 2009. "Structural Changes in the US Economy: Is There a Role for Monetary Policy?" *Journal of Economic Dynamics and Control* 33 (February): 477–90.
- Chung, Hess, Jean-Philippe Laforte, David Reifschneider, and John C. Williams. 2012. "Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events?" *Journal of Money, Credit and Banking* 44 (February): 47–82.
- Clark, Todd E., and Francesco Ravazzolo. 2015. "Macroeconomic Forecasting Performance under Alternative Specifications of Time-Varying Volatility." *Journal of Applied Econometrics* 30 (June/July): 551–75.
- Del Negro, Marco, and Giorgio E. Primiceri. 2015. "Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum." *Review of Economic Studies* 82 (October): 1342–45.
- Doh, Taeyoung, and Michael Connolly. 2012. "The State Space Representation and Estimation of a Time-Varying Parameter VAR with Stochastic Volatility." Federal Reserve Bank of Kansas City Working Paper 12-04 (July).
- Jones, Callum. 2017. "Unanticipated Shocks and Forward Guidance at the ZLB." Manuscript (October).
- Kiley, Michael, and John M. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity* (Spring): 317–72.
- Laubach, Thomas, and John C. Williams. 2016. "Measuring the Natural Rate of Interest Redux." *Business Economics* 51 (April): 57–67.
- Lubik, Thomas A., and Christian Matthes. 2015a. "Time-Varying Parameter Vector Autoregressions: Specification, Estimation, and an Application." Federal Reserve Bank of Richmond *Economic Quarterly* 101 (Fourth Quarter): 323–52.
- Lubik, Thomas A., and Christian Matthes. 2015b. "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches." Federal Reserve Bank of Richmond *Economic Brief* 15-10 (October).

- Lubik, Thomas A., and Christian Matthes. 2016. "Indeterminacy and Learning: An Analysis of Monetary Policy in the Great Inflation." *Journal of Monetary Economics* 82 (September): 85–106.
- Lubik, Thomas A., Christian Matthes, and Andrew Owens. 2016. "Beveridge Curve Shifts and Time-Varying Parameter VARs." Federal Reserve Bank of Richmond *Economic Quarterly* 102 (Third Quarter): 197–223.
- Primiceri, Giorgio E. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (July): 821–52.

**Figure 1** Estimated Probabilities of Being at the Zero Lower Bound for Three Sample Periods



**Figure 2 Estimated Probabilities of Being at the Zero Lower Bound for Three Sample Periods**





**Figure 3 Probabilities of the Zero Lower Bound at Different Time Horizons**

