Bayesian Forecasting and Decision Analysis in Macroeconomic Time Series Using Dynamic Dependence Networks

Liaowang (Zoey) Zou *

April 1, 2017

Abstract

Macroeconomics and monetary policies are constantly evolving. Time-varying vector autoregressive models (TV-VARs) in forecasting multivariate time series aims to capture these variations and allow model dynamics to change over-time. We utilize dynamic dependence networks as an extension to TV-VARs to develop models that capture the dynamic relationships among macroeconomic time series. Dynamic dependence network models allow univariate series to be decoupled for fast, parallel processing, and recoupled for forecasting and decision analysis. We focus on changes in policy-related variables such as interest rate and the resulting effects on the economy, paying special attention to an identified policy target such as inflation. We compare a 8-variable and 3-variable model and perform decision analysis using both at two specific time points with a specified loss function. The insights from these models shed light on the central banks' decision making process.

KEY WORDS: Bayesian forecasting; time-varying vector autoregressive models; macroeconomic time series; decision analysis.

1 Introduction

Macroeconomic forecasting is an important element in central banks' monetary policymaking processes, especially when central banks anchor certain macroeconomic variables as their policy targets. For instance, since the 1980s, the U.S. Federal Reserve has placed emphasis on targeting inflation to stabilize the economy [Boivin and Giannoni, 2003]. To achieve policy targets, central banks use many monetary policy instruments, including nominal interest rate targets. As a result, it is critical to have a forecasting model that forecasts well how a change in a policy control variable could impact policy targets over a defined time period. This paper is interested in the dynamic relationship between policy instruments and policy targets, correlated with other confounding macroeconomic variables which could help with central banks' decision making processes by choosing the values of policy control variables to best guide the economy towards some specified, desirable policy targets.

To achieve this, we apply a stochastic time series model – a dynamic dependence network – to model and forecast a set of macroeconomic variables multiple-steps ahead. Section 2 introduces

*This work was developed as my undergraduate honors thesis in Statistical Science during 2016-17. I would like to thank Professor Mike West and Amy Meng Xie (PhD student in Statistical Science) for their relentless support and guidance throughout the research process and edits on the paper. I would also like to thank Jouchi Nakajima (visiting senior economist at Bank of International Settlements) for offering insights and constructive comments.

the setting and notation in a macroeconomic context. Section 3 explains the setup for DDNMs, discussing the time-varying autoregressive (TVAR) features and conditional independence structure. It also notes the evolution of state vectors through random walks and priors for forward-filtering and 1-step forecasting. Then, we bring in decision analysis in Section 5 and develop a loss function that best reflects the central bank's economic goal while taking into account other considerations like smoothness of the projected path. We derive the Bayesian risk from posterior predictive distributions with technical details in Appendix.

Building on the context of the US economy from 1965 to 2016, Sections 4 and 6 compare an 8-variable model with a 3-variable model and summarize key aspects of their DDNM specifications, predictions and decision analysis. Section 7 concludes the paper with some remarks and suggests directions for future work.

2 General Setting

The set of q macroeconomic variables forms a q-vector time series $\mathbf{y}_t = (y_{1t}, \cdots, y_{qt})\prime$ over $t = 0, 1, \cdots$, discrete time. $\mathbf{y}_t$ can be partitioned into $\mathbf{x}_t,\ \mathbf{u}_t,\ \mathbf{r}_t$, where:

- $\mathbf{x}_t$ is q_x-vector of response variables whose levels into the future are of main concern to policymakers (e.g., inflation and/or unemployment), and are selected as defining economic goals in the decision analysis;
- $\mathbf{u}_t$ is a q_z-vector of response variables of secondary interest (e.g., consumption, investment, etc);
- $\mathbf{r}_t$ is a q_r-vector of indices that are potential targets to control (e.g., central bank interest rates, or other putative policy instruments).

At time point t, we model $\mathbf{y}_t$ and predict h-steps ahead. We denote the posterior density function of predictive distributions over the next h steps as $p(\mathbf{y}_{t+1:t+h}|\mathcal{D}_t)$ where $\mathcal{D}_t$ represents the prior information available at time t, including all past observations $\mathbf{y}_{1:t}$ and any additional information used to modify the model at any past time up to the current time t. Decisions are made at time t by setting the control variables $\mathbf{r}_t = \mathbf{r}$ to guide $\mathbf{x}_t$ towards specific targets over the next h time periods.

3 Dynamic Dependence Network Models

To address the complex dynamics arising from cross-series relationships of multiple time-varying indices, we use dynamic dependence network models (DDNMs) developed by Zhao et al. [2016]. They are extensions of multiregression dynamic models (MDMs) by allowing for TVAR components in each series.

The setup for DDNMs is, in a q-variable model over times $t = 1, 2, ...$,

$$y_{jt} = c_{jt} + \sum_{i=1:p_{j\lambda}} \mathbf{y}'_{t-i}\boldsymbol{\phi}_{jit} + \mathbf{y}'_{pa(j),t}\boldsymbol{\gamma}_{jt} + v_{jt}, \quad j = 1:q, \tag{1}$$

with components as follows:

- y_{jt} is a univariate time series observed over time $t = 1, 2, ...$
- $\mathbf{y}_t = (y_{1t}, ..., y_{qt})'$ is $q \times 1$-vector time series.

- c_{jt} is a time-varying intercept.
- Each $\boldsymbol{\phi}_{jit}$ is a q-vector of TV-VAR coefficients for lag $i = 1, \cdots, p_{j\lambda}$ where $p_{j\lambda}$ is the maximum lag.
- $pa(j)$ is a parental set. It is a subset of indices of those series higher than j in the hierarchical structure.
- $\mathbf{y}_{pa(j),t}$ is the $|pa(j)| \times 1$-vector of time t values on the series in the parental set $pa(j)$.
- $\boldsymbol{\gamma}_{jt}$ is a vector of dynamic regression coefficients, linking contemporaneous values of some other series to series j. The number of parents and dimension of $\boldsymbol{\gamma}_{jt}$ is $|pa(j)|$.
- v_{jt} is the observation noise; these are conditionally independent over j and t with $v_{jt} \sim N(0, 1/\lambda_{jt})$ where λ_{jt} is the time-varying precision.

One key feature of DDNMs is the parental set for each univariate dynamic linear model (DLM). For $j = 1 : q-1$, $pa(j) \subseteq j+1 : q$ is a subset of indices of those series higher than j in the selected order, and we set $pa(q) = \emptyset$ as it is at the bottom of the hierarchical structure. The parental set $pa(j)$ contains contemporaneous values of some other univariate series, which are unknown at time t. These contemporaneous values, however, could be predicted using lagged predictors, and the results are fed into the model for forecasting. This reflects the contemporaneous conditional dependence structure: for any $i > j$, $y_{it} \perp\!\!\!\perp y_{jt} | \mathbf{y}_{pa(j),t}$ if $i \notin pa(j)$. The hierarchical structure of q independent, univariate DLMs defines a full multivariate model for $\mathbf{y}_t$.

The conditional independence structure of DDNMs allow models to be decoupled into the set of univariate DLMs for fast, efficient forward-filtering and 1-step forecasting, and then recoupled for multivariate forecasting and decision analysis. For each DLM, we define a state vector $\boldsymbol{\theta}_{jt} = (\boldsymbol{\phi}'_{jt}, \boldsymbol{\gamma}'_{jt})'$ and specify its dynamic model forms and precision λ_{jt}. As in Zhao et al. [2016], we adopt random walk models with two discount factors $\delta_j \in (0, 1]$ and $\beta_j \in (0, 1]$. δ_j defines the error variance matrix in random walk evolutions for $\boldsymbol{\theta}_{jt}$ and β_j defines the beta random variate in the evolution of the time-varying precision λ_{jt}. At time $t-1$, $\boldsymbol{\theta}_{j,t-1}$ and $\lambda_{j,t-1}$ have the following conditional distribution:

$$\begin{aligned} \boldsymbol{\theta}_{j,t-1} | \lambda_{j,t-1}, \mathcal{D}_{t-1} &\sim N(\mathbf{m}_{j,t-1}, \mathbf{C}_{j,t-1}/(s_{j,t-1}\lambda_{j,t-1})) \\ \lambda_{j,t-1} | \mathcal{D}_{t-1} &\sim G(n_{j,t-1}/2, n_{j,t-1}s_{j,t-1})/2) \end{aligned} \tag{2}$$

where $\mathbf{m}_{j,t-1}$, $\mathbf{C}_{j,t-1}$, $s_{j,t-1}$ and $n_{j,t-1}$ are parameters of the normal and gamma distribution, and $\mathcal{D}_{t-1}$ summarizes information at $t-1$. The posterior calculation, 1-step forecast at time $t-1$ and multivariate predictive distribution 1-step ahead are detailed in Zhao et al. [2016, Section 2].

For any $h > 1$, forecasting $h-$step ahead could lead to unknown lagged predictors at the time of forecasting. DDNMs use a simulation method to sample $\mathbf{y}_{t+2}$ conditional on the simulated samples of $\mathbf{y}_{t+1}$. This is done recursively up to h-step. The independent samples then form a Monte Carlo sample from the full predictive distribution $p(\mathbf{y}_{t+1}, ..., \mathbf{y}_{t+h} | \mathcal{D}_t)$ which can be summarized.

4 A Study of US Macroeconomic Time Series

Bayesian modeling is increasingly popular in macroeconomic studies. Some of the common methods include dynamic factor models [Stock and Watson, 2002, Forni et al., 2005, Amisano and Geweke, 2013], vector autoregressive (VAR) models [Sims, 1980], and TV-VAR models [Primiceri, 2005, Koop and Korobilis, 2009]. Nakajima and West [2013] introduced latent threshold process to TV-VAR modeling to ensure dynamic sparsity by shrinking some time-varying parameters to zero for certain periods of time. This method, however, is computationally expensive, making it less suitable to analyze a multivariate time series with a large number of variables. DDNMs, on the other hand, are more efficient to implement due to decoupling and recoupling features discussed in Section 3. One main goal of the study is to utilize DDNMs to explore models in a macroeconomic context. DDNMs with extended model uncertainty analysis have proven forecast accuracy and model stability in the context of a 13-series, high-frequency financial portfolio analysis [Zhao et al., 2016]. We are interested in examining if conditional dependence structures in macroeconomic series can be captured by sparse DDNMs which could feed into decision analysis using multiple-steps ahead forecasting.

4.1 Context and Data

We analyze US macroeconomic time series with $\mathbf{y}_t$ being monthly values of these series. The dual mandate of Federal Reserve sets price stability and maximum sustainable employment as two goals of US monetary policy. Multiple studies suggest that the Federal Reserve has placed heavy emphasis on price stability (in practice, low and stable inflation), especially after 1980s [Poole and Wheelock, 2008]. As such, we pick inflation as the policy target x_t.

Adjusting the nominal interest rate is a major tool of monetary policy. The Federal Reserve changes the target Federal funds rate which is the best proxy to the nominal interest rate. Here, we assume that interest rate could be directly controlled, that is, a decision to set $r_t = r$ actually achieves that target value. In the attempt to capture the complex interaction among more macroeconomic variables, we include 6 other series of secondary interest $\mathbf{u}_t$. They are wages, unemployment, consumption, investment, M1 and M2 money supply.

The data we use comes from the Federal Reserve Economic Data and Bloomberg terminal. The inflation rate is the annual percentage change in a chain-weighted consumer price index excluding food and energy, and the interest rate is the Federal funds rate. The detailed description of all series is presented in Appendix 8.1. The data cover a time period of 612 months from January 1965 to September 2016 which includes the Great Inflation, the Great Moderation, the financial crisis and recovery, each with different characteristics; see Figure 1.

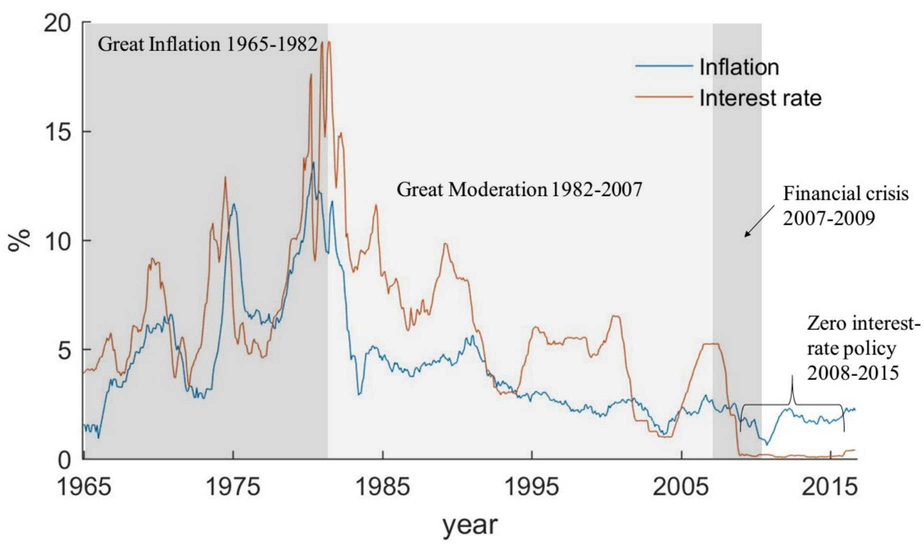

Figure 1: US inflation and interest rate from 1965 to 2016

4.2 8-Variable Model Setup

For DDNMs, appropriate ordering and structuring is crucial. We perform preliminary model selection using training data and a linear model. The training period is from January 1984 to December 2003 (240 observations) which excludes the volatility in the 1970s and provides ample observations for variable selection. For $q = 8$ identified macroeconomic series, consider each univariate series y_{jt} at time t:

$$y_{jt} = \mu_j + \alpha_1 y_{j(t-1)} + \cdots + \alpha_{12} y_{j(t-12)} + \boldsymbol{\beta}_1^T \mathbf{z}_t + \boldsymbol{\beta}_2^T \mathbf{z}_{t-1} + \cdots + \boldsymbol{\beta}_{13}^T \mathbf{z}_{t-12}, \quad j = 1:q \tag{3}$$

Here, μ_j is the intercept and $\mathbf{z}_t$ are contemporaneous variables at time t. 12 months of lagged values of all variables from $t-1$ to $t-12$ are also included in the model. Using both forward and backward stepwise regression on each of the 8 variables with a threshold of p <0.05, we constrain the model to have at most 1 predictor in $pa(j)$ and 4 lagged predictors for each series to ensure sparsity. Another way is applying DDNM with model uncertainty [Zhao et al., 2016]. Here model probabilities and hyperparameter uncertainty is not the key focus. Hence, we do not use this method in our case study. Based on the preliminary model selection results, we decide on the parental set $pa(j)$ and lagged predictors for each series detailed in Table 1.

j	Name	Parental Set $pa(j)$	Lagged Predictors
1	Interest rate(r)	p	AR(1), AR(2), w(3), $m2$(12)
2	Inflation(p)	$\emptyset$	AR(1),AR(3), r(1), $m2$(12)
3	Wage(w)	u	AR(1), AR(3), u(8)
4	Unemployment(u)	c	AR(1), AR(3) w(4), i(1)
5	Consumption(c)	$m2$	AR(1), AR(12), i(3)
6	Investment(i)	$m2$	AR(1), AR(12), $m2$(3)
7	M2 money supply($m2$)	$m1$	AR(1), r(1), i(12), c(12)
8	M1 money supply($m1$)	$\emptyset$	AR(1), $m1$(1), r(3), i(12)

Table 1: Ordering and specification for 8-variable model

It is natural that the response variable with lag 1-month AR(1) is the most significant factor in predicting current response variable. In addition, various variables with lag 3-month and lag 12-month are often selected, particularly M2 money supply, investment and consumption, suggesting the presence of seasonality both quarterly and annually. M1 money supply is at the bottom of the hierarchical structure with an empty parental set. On the other hand, interest rate as the monetary policy instrument is at the top, reflecting the need of contemporaneous variables, aka current economic condition in its forecast. Decisions made at time t of setting $r_t = r$ is fed into the model starting from time $t+1$.

Using data from January 1965 to December 1969 (60 observations), we get a rough estimate of the prior values. The prior and discount factor specifications for each series y_{jt} are listed as follows:

- $\mathbf{m}_0$ is $(|\boldsymbol{\gamma}_{jt}| + |\boldsymbol{\phi}_{jt}|) \times 1$ vector with all zeros, except for the AR(1) coefficient set to 0.95.
- $\mathbf{C}_0$ is $(|\boldsymbol{\gamma}_{jt}| + |\boldsymbol{\phi}_{jt}|) \times (|\boldsymbol{\gamma}_{jt}| + |\boldsymbol{\phi}_{jt}|)$ diagonal matrix with $\mathbf{C}_0(1,1) = 0.01$ and $\mathbf{C}_0(d,d) = 1000$ for $d = 2 : |\boldsymbol{\gamma}_{jt}| + |\boldsymbol{\phi}_{jt}|$.
- for $j = 1$ (interest rate), $s_0 = 0.75$; for $j = 4$ (unemployment), $s_0 = 0.05$; for other js, $s_0 = 0.15$.
- $n_0 = 20$.

- $\delta = 0.99$.
- $\beta = 0.99$.

We set δ and β at relatively high values to encourage stability in inferred parameter trajectories.

With the model and priors as discussed above, we run the DDNM analysis for the time period from January 1970 to September 2016. At each time step, we perform forward-filtering, update posteriors, and evaluate 12- to 24-step ahead forecast distributions at each time point. Each forecast generates a Monte Carlo sample with size 100,000 from the full, multiple-steps ahead predictive distribution.

4.3 Point Forecasts

Figure 2 displays the predictive means of 1- and 2-year ahead forecast of three major series: interest rate, unemployment and inflation. This is done recursively: starting at January 1970 up to any month t, we perform forward-filtering to generate predictions based on $\mathbf{y}_{1:t}$. Then, we simulate out-of-sample predictive distributions over time $t+1 : t+24$. The means at time $t+12$ and $t+24$ are recorded and are compared to the actual values. We then move to the next month $t+1$, generate predictions based on the updated series $\mathbf{y}_{1:t+1}$ and forecast the next 24 months $t+2 : t+25$. This is repeated to September 2014, generating a series of out-of-sample forecasts for each month over the entire time frame. From Figure 2, we can see that both 1- and 2-year ahead forecast exhibit similar trends as the actual series, with 2-year forecasts having a larger volatility, especially in the period of 1980 to 1985. This could be attributed to the small number of observations to train the model since the learning period starts on January 1970. The great inflation from 1965 to 1982 also contributes to the heightened volatility.

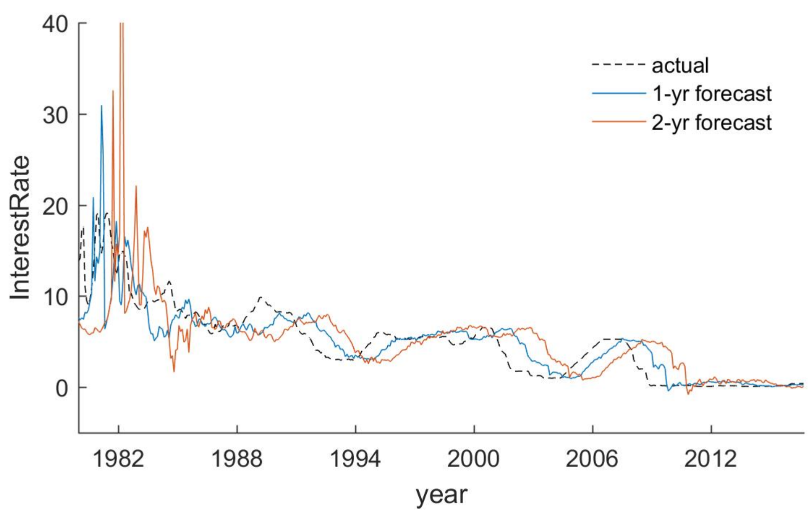

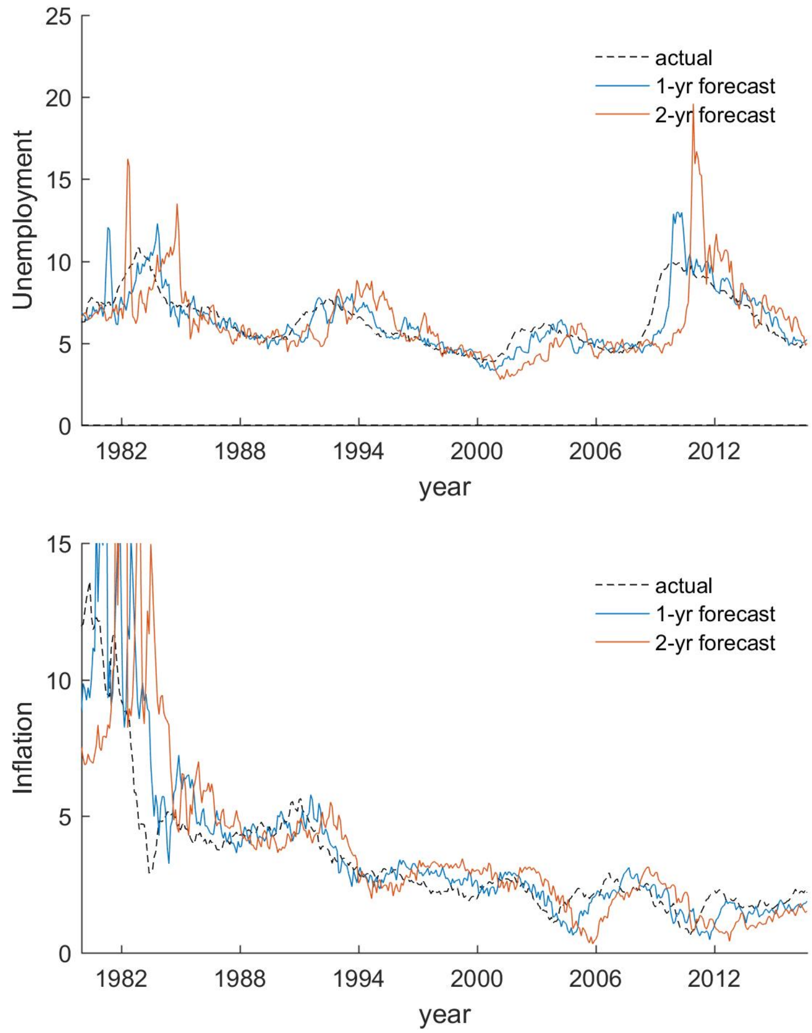

Figure 2: 1- and 2-year ahead forecast vs. actual value for US interest rate, unemployment and inflation from 1980 to 2016

4.4 Impulse Response Analysis

Another key aspect of out-of-sample prediction is impulse response analysis. It describes the reaction of the economy given a shock in the system. In the context of the Federal Reserve adjusting interest rate, we set the shock to be a 2% increase in interest rate at time t, i.e. $r_t^* = r_t + 2$. This is a rather radical increase given that the average monthly change in interest rate is 0.005. We hope this drastic change would give a clearer picture of the direction of movement of various macroeconomic series in the economy. From the recursive forecasting analysis in Section 4.3, at each month we sample the predictive distribution for each variable given the shock. The difference between the predictive means with and without the shock is plotted in Figure 3.

When analyzing Figure 3, we pay particular attention to the following features: (i) whether there is a positive or negative relationship between the series and the shock on interest rate; (ii) the difference in impulse response trajectories between short-run (1-year ahead) and longer-run (2-year ahead) forecasting horizons; (iii) the variation in volatility over the time period.

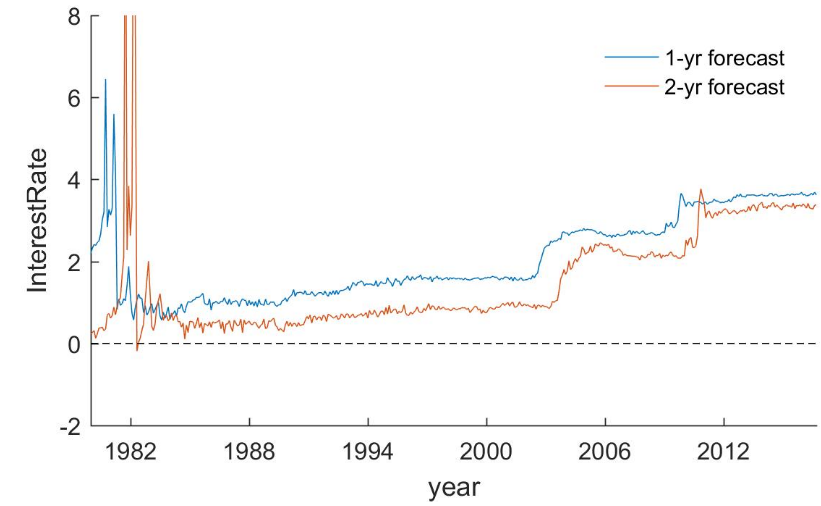

1-yr forecast
2-yr forecast
InterestRate
8
6
4
2
0
-2
1982
1988
1994
2000
2006
2012
year

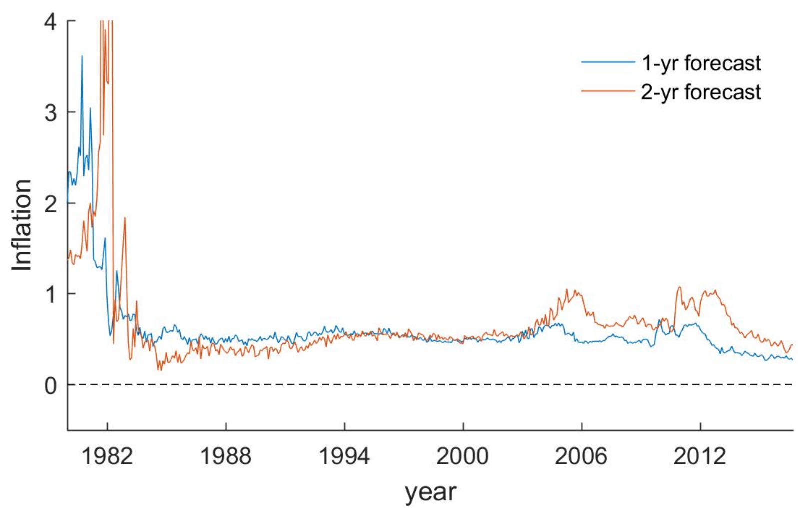

1-yr forecast
2-yr forecast
Inflation
4
3
2
1
0
1982
1988
1994
2000
2006
2012
year

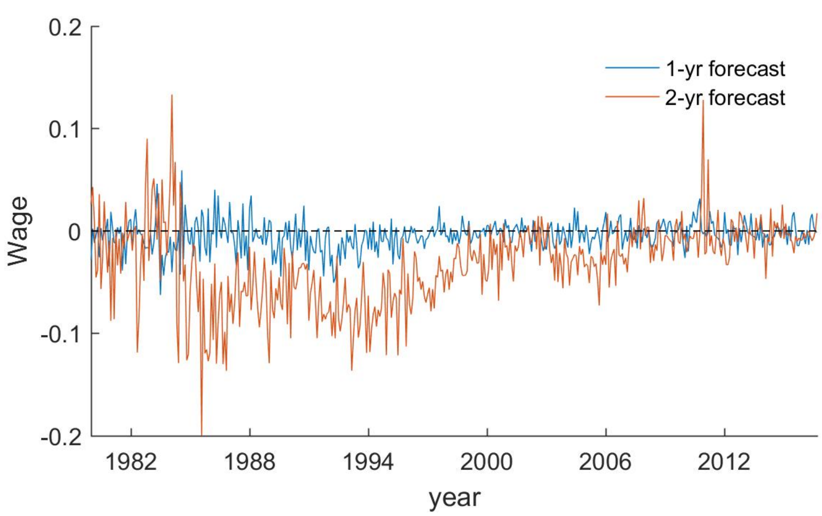

1-yr forecast
2-yr forecast
Wage
0.2
0.1
0
-0.1
-0.2
1982
1988
1994
2000
2006
2012
year

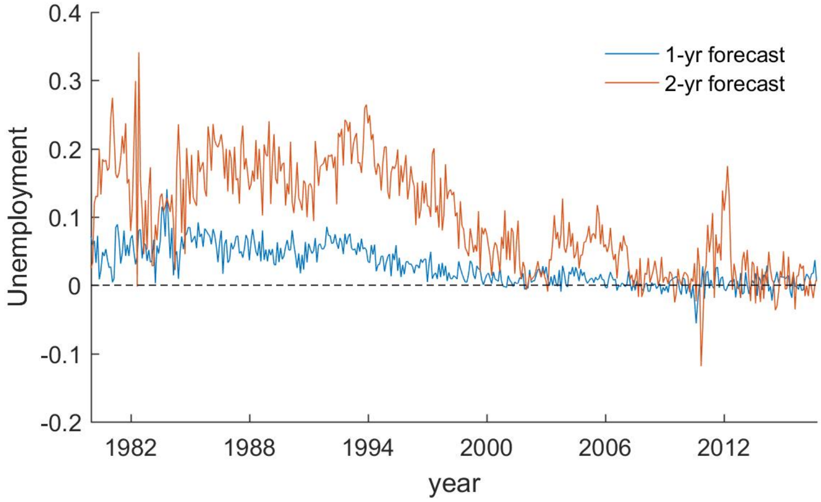
Unemployment
year
1-yr forecast
2-yr forecast
0.4
0.3
0.2
0.1
0
-0.1
-0.2
1982
1988
1994
2000
2006
2012

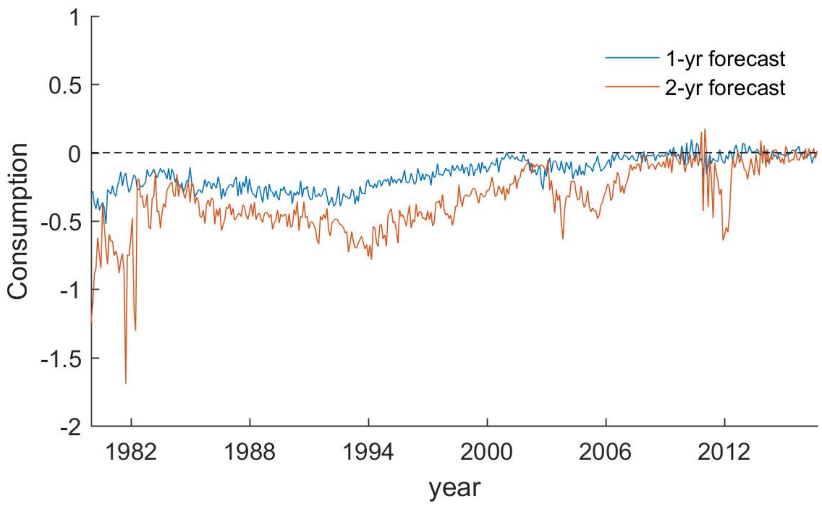
Consumption
year
1-yr forecast
2-yr forecast
1
0.5
0
-0.5
-1
-1.5
-2
1982
1988
1994
2000
2006
2012

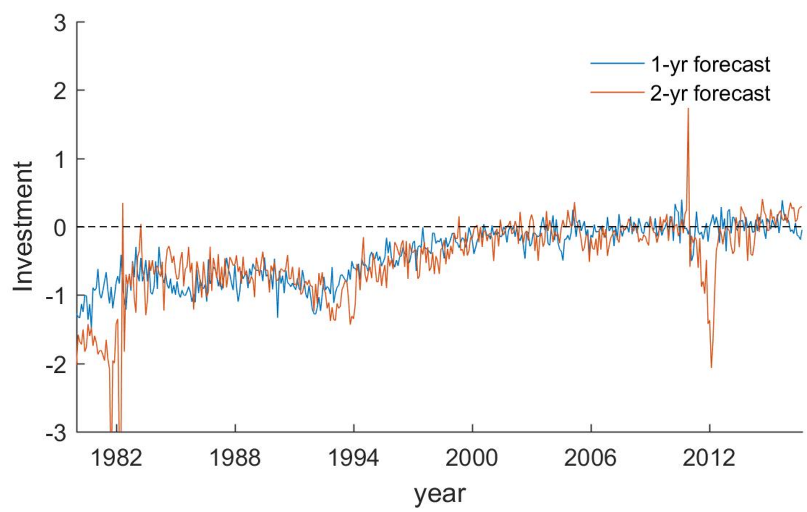
Investment
year
1-yr forecast
2-yr forecast
3
2
1
0
-1
-2
-3
1982
1988
1994
2000
2006
2012

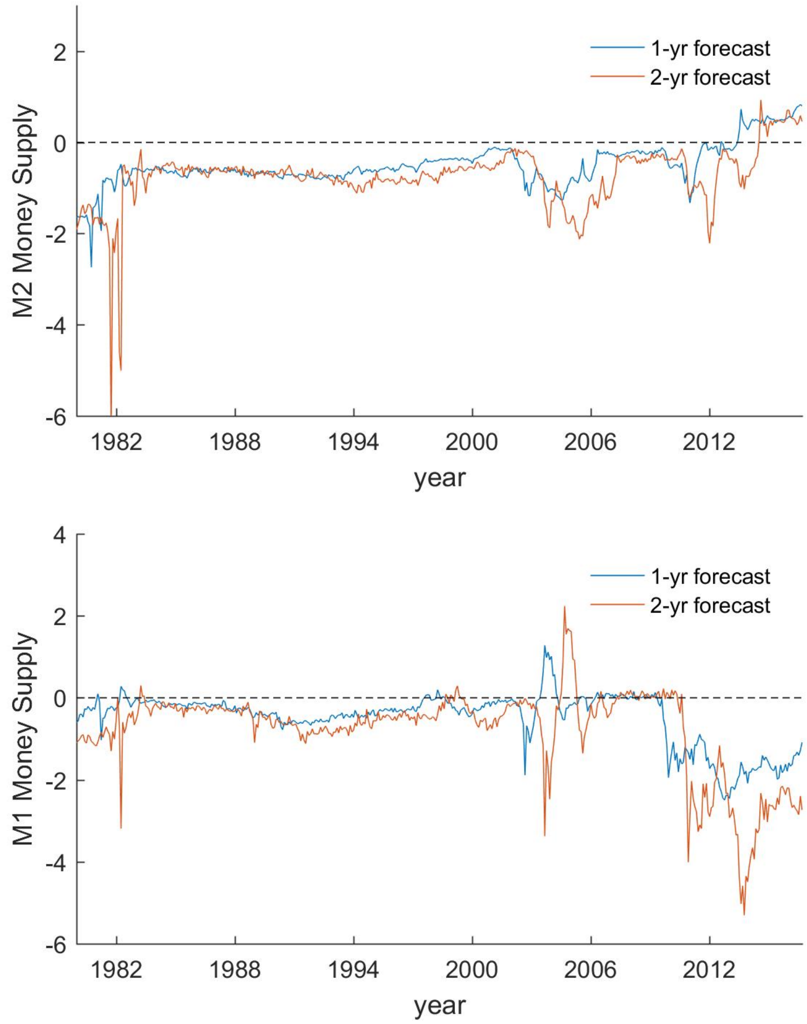

Figure 3: Impulse response trajectories for 1- and 2-year ahead horizons for US macroeconomic series from 1980 to 2016. The shock introduced is a 2% increase in interest rate.

From Figure 3, we can see that (i) interest rate, inflation, and unemployment are generally predicted to respond positively to a shock in interest rate, while consumption and M2 money supply respond negatively. Wage exhibits minimal change. Investment responds negatively until 1997, then, the response becomes ambiguous. The initial response of M1 money supply is also unclear, but after 2008, it shows a clear negative response. (ii) For most series, the response is more prominent with a longer forecasting horizon, except for interest rate and inflation. Excluding the volatile period in early 1980s, interest rate exhibits a larger increase in the short-run given a fixed shock level. Inflation, on the other hand, first exhibits similar trend as interest rate before 2004, then, long-run response becomes more prominent. (iii) In early 1980s, the response trajectories have high volatility. This coincides with point forecasts in that period from Figure 2. After the 1980s, most series are stable. However, the trajectories for wage, unemployment and investment are erratic throughout the forecasting period. This suggests that their DDNM parameters lie outside the stationary region of the implied "local" VAR models at specific time points. While it is difficult to impose constraints on these parameters, one possible solution is only accepting simulated parameters that fall inside the stationary region. Future research is needed to address this problem.

To further examine the response of each series from an interest rate hike, we plot the impulse response function at specific time points. We choose two baselines in different economic contexts: June 2005 and October 2011. June 2005 was at the tail of the Great Moderation, with significant decline in macroeconomic volatility and steady real GDP growth. The economy in October 2011, on the other hand, was in the recovery period after the financial crisis. The Federal Reserve has been practicing zero interest rate policy since 2008 and second round of quantitative easing has just ended. As such, even though they both have around 2% inflation, the interest rates are very different. In June 2005 the interest rate is 3.04% while in October 2011 it is 0.07%. We are interested in understanding how the economy with different interest rate levels respond differently when it is subjected to a same shock. As in the previous analysis, we introduce a 2% increase in interest rate. Figures 4,5 show the prediction trajectories from time $t+1$ to $t+24$ given the baseline.

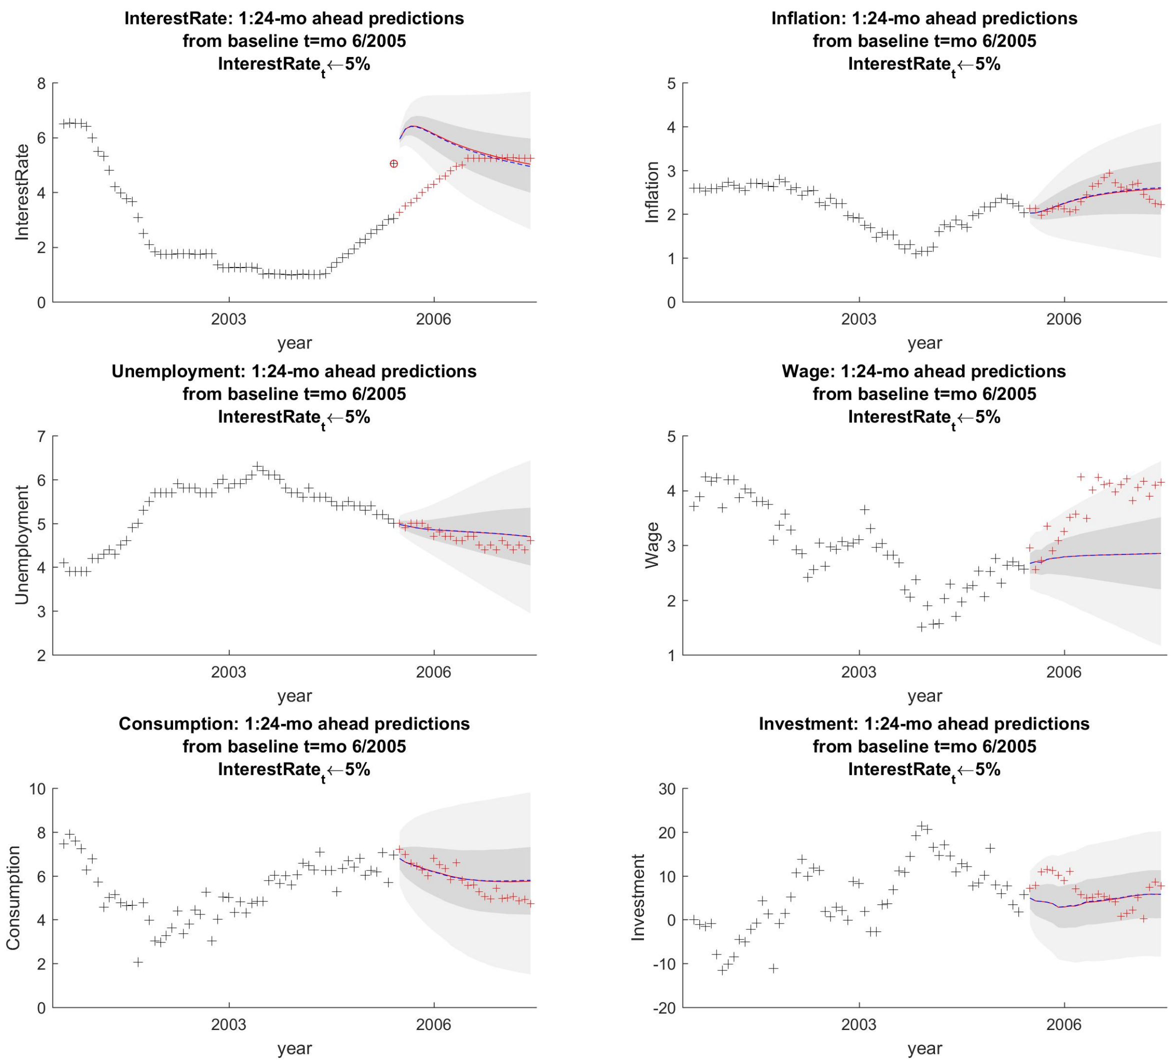

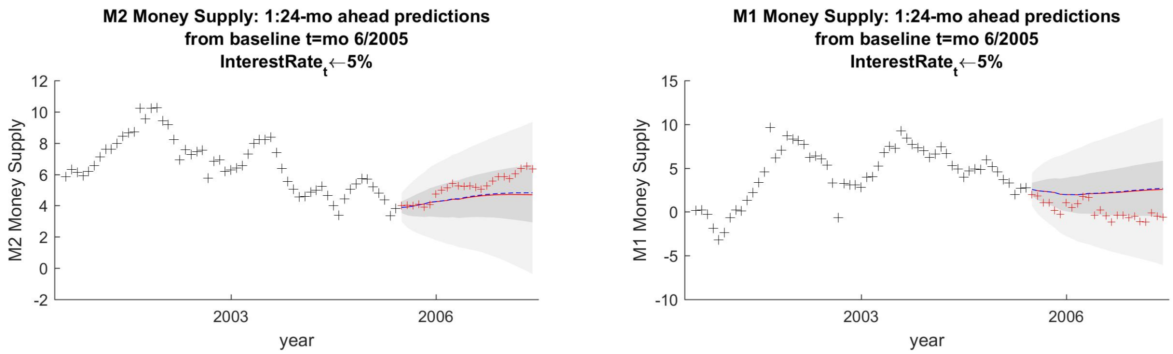

Figure 4: 24-month ahead forecasting given a 2% increase in interest rate for US macroeconomic series from baseline June 2005. Red and blue dashed lines represent means and medians of the posterior distribution respectively. Grey regions mark 50% and 90% credible interval. The actual values of the series are plotted using red crosses.

As illustrated in Figure 4, in response to a 2% interest rate hike at baseline, interest rate first responds by increasing further, then drops gently. Inflation exhibits a gradual increase while unemployment a slight decrease. The rest of the variables show some changes in either directions, but level off at the end. This shows that the economy is fairly stable, even subject to a radical change. However, the positive response of inflation and negative response of unemployment contradict the common belief that an increase in interest rate would bring down inflation and lead to higher unemployment.

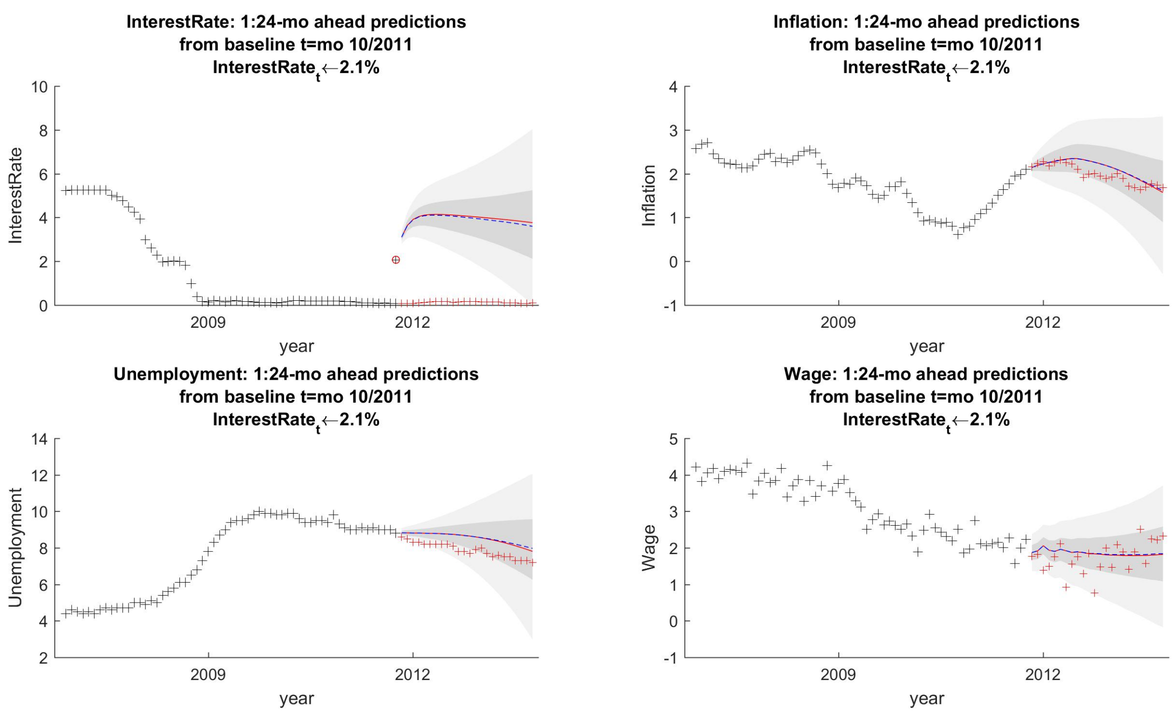

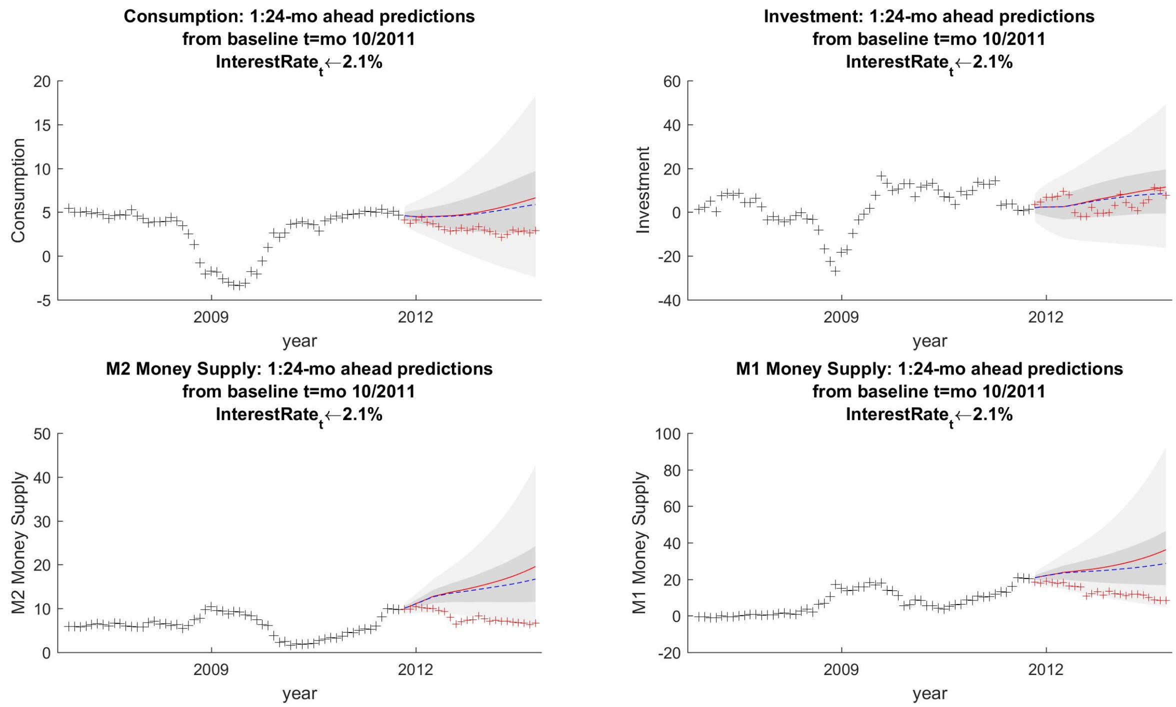

Figure 5: Impulse response trajectories 24-month ahead for US macroeconomic series from baseline October 2011

Figure 5 tells a different story. With the same 2% interest rate increase, interest rate follows a similar trajectory. Both inflation and unemployment experience declines, while consumption, investment, M1 and M2 money supply all experience increases. This is peculiar as we expect heightened interest rate to suppress consumption, investment and money supply, hence, bring down inflation. This shows that the relationship between inflation and other variables is more intricate. The economy at this period is more susceptible to external shocks, though the direction of movement is not as expected.

5 Decision Analysis

5.1 Loss Function and Risks

We specify the loss function to have a quadratic form:

$$L_t(\mathbf{y}_{t+1:t+h}) = \sum_{k=1:h} \delta_k(\mathbf{x}_{t+k} - \mathbf{x}_{t+k-1})'(\mathbf{x}_{t+k} - \mathbf{x}_{t+k-1}) \; + \; \rho(\mathbf{x}_{t+h} - \boldsymbol{\tau})'(\mathbf{x}_{t+h} - \boldsymbol{\tau}) \tag{4}$$

that accounts for reaching specified target values $\boldsymbol{\tau}$ over the next h periods. $\boldsymbol{\tau}$ is a q_x-vector targets for response variable $\mathbf{x}_{t+h}$. $\delta = (\delta_1, ...\delta_h)'$ is a vector of non-negative weights that penalizes large changes in the $\mathbf{x}_{t+k}$ variables over each time period and aims to reward a smooth trajectory towards the target. The non-negative ρ penalizes large deviations of the terminal values $\mathbf{x}_{t+h}$ from the specified target $\boldsymbol{\tau}$. In this case study, we treat inflation as the only target variable, thought it can be extended to include more variables such as unemployment.

With the loss function specified, we can derive risks from predictive distributions generated from the forecasting model which directly sets $\mathbf{r}_t = \mathbf{r}$. The Bayesian risk is the expected loss. Thus, the resulting risk function is

$$
\begin{aligned}
R_t(\mathbf{r}) &= E[L_t(\mathbf{y}_{t+1:t+h})|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t] \\
&= \sum_{k=1:h} \delta_k E[(\mathbf{x}_{t+k} - \mathbf{x}_{t+k-1})'(\mathbf{x}_{t+k} - \mathbf{x}_{t+k-1})|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t] \\
&\qquad + \rho E[(\mathbf{x}_{t+h} - \boldsymbol{\tau})'(\mathbf{x}_{t+h} - \boldsymbol{\tau})|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t]
\end{aligned}
\tag{5}
$$

Since we use a quadratic loss, $R_t(\mathbf{r})$ depends only on the means, variances and covariances of $\mathbf{x}_{t_1:t+h}$ conditional on $\mathbf{r}_t = \mathbf{r}$ and $\mathcal{D}_t$. The detailed calculation appears in Appendix 8.2. Here, $R_t(\mathbf{r})$ is only the mean of the loss distribution. We are interested in other aspects of that distribution, for instance, the median values. The solution to this is simulation: As we simulate the predictive distribution $p_t(\mathbf{y}_{t+1:t+h}|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t)$ from the Monte Carlo sample, we can directly calculate values of loss in eqn. 4 and take the median.

5.2 Decision Examples

In the context of the case study, we treat interest rate as the policy control variable r and specify a range of values of r. For each inflation target, the risk is evaluated at each value of r. By setting $\rho = 1$ and $\delta_k = 0.9^k$, we plot the risk function $R_t(r)$ and choose r to minimize the risk.

Figures 6, 7 show the risk function of the mean and median loss for different levels of interest rate by setting inflation target at 1-, 1.5-, 2-, 2.5- and 3%. We constrain interest rate to be non-negative as the US has yet to adopt the negative interest rate policy (NIRP), even though other central banks like European Central Bank, Bank of Japan have shown that NIRP is definitely a possibility.

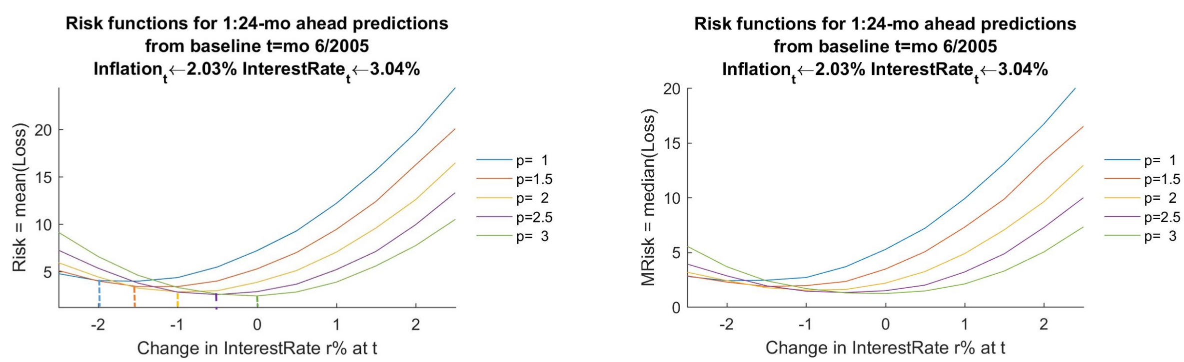

Figure 6: Mean and median loss for p=1:3 over 24-month ahead forecasting horizons from baseline June 2005 with the local minima marked by dotted lines.

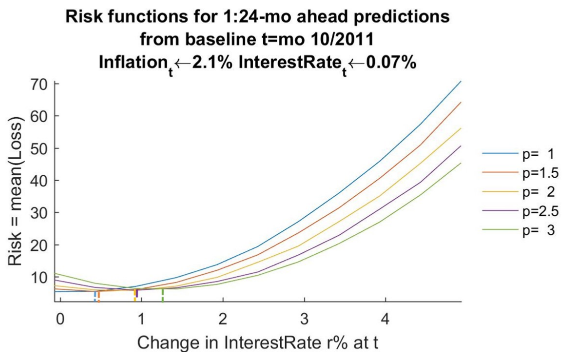

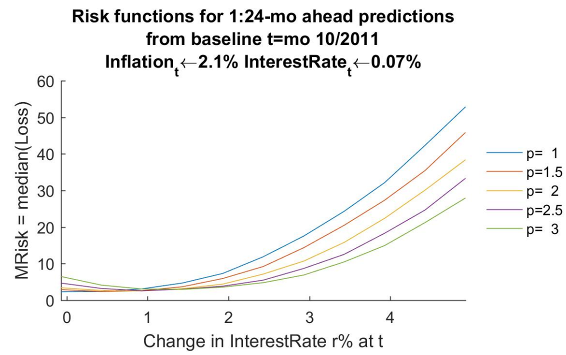

Figure 7: Mean and median loss for p=1:3 over 24-month ahead forecasting horizons from baseline October 2011 with the local minima marked by dotted lines.

With similar inflation but different interest rates at the baseline, both Figures 6 and 7 illustrate that as the target value of inflation increases, the local minimum of risk functions shifts to the right. In June 2005, because interest rate is already at a relatively high level, the model suggests a decrease with respect to baseline interest rate to both increase inflation to 3% and decrease inflation to 1%. On the other hand, the nearly zero interest rate in October 2010 prevents it dropping further. Hence, the model suggests an increase in interest rate to achieve inflation adjustment in both directions and a greater interest rate increase for higher inflation.

In addition, both figures depict that mean loss is higher than median loss for a specified policy target and control variable, suggesting a right-tailed distribution of loss. From impulse response analysis and decision analysis, our 8-variable model suggests that an increase in interest rate tends to result in higher inflation levels.

6 Comparison with a 3-Variable Model

To test if the insights derived from the 8-variable model are supported by other models, we reduce the model complexity by selecting 3 key variables of interest: interest rate, inflation and unemployment. As in the previous analysis, interest rate is the policy control variable, inflation is the target variable, and unemployment is the response variable of secondary interest.

6.1 3-Variable Model Setup

The DDNM is adapted from Nakajima and West [2013]. In that paper, the authors use latent threshold in TV-VAR modeling of US macroeconomic data. We select predictors that their analysis indicated based on quarterly data.We convert lagged predictors from quarterly scale to monthly scale through a multiplication of 3. The model specification is shown in Table 2:

j	Name	Parental Set $pa(j)$	Lagged Predictors
1	Interest rate(r)	u, p	AR(3), p(9), u(9)
2	Unemployment(u)	$\emptyset$	AR(3), r(6), r(9)
3	Inflation(p)	$\emptyset$	AR(3), AR(6), u(3)

Table 2: Ordering and specification for 3-variable model

As in the 8-variable model, interest rate is placed at the top of the hierarchical structure. Priors and discount factor specification are the same as the previous analysis. They are:

- $\mathbf{m}_0$ is $(|\gamma_{jt}| + |\phi_{jt}|) \times 1$ vector with all zeros, except for the AR(1) coefficient set to 0.95.
- $\mathbf{C}_0$ is $(|\gamma_{jt}| + |\phi_{jt}|) \times (|\gamma_{jt}| + |\phi_{jt}|)$ diagonal matrix with $\mathbf{C}_0(1,1) = 0.01$ and $\mathbf{C}_0(d,d) = 1000$ for $d = 2 : |\gamma_{jt}| + |\phi_{jt}|$.
- for $j = 1$ (interest rate), $s_0 = 0.75$; for $j = 2$ (unemployment), $s_0 = 0.05$; for $j = 3$ (inflation), $s_0 = 0.15$.
- $n_0 = 20$.
- $\delta = 0.99$.
- $\beta = 0.99$.

We rerun the DDNM analysis over the same time period – January 1970 to September 2016.

6.2 On-line Trajectories of State Variables

Figure 8 displays the evolution of state variables over time for each univariate DLM for interest rate, unemployment and inflation series. With mean and 95% credible interval clearly marked, this shows how coefficients of these predictors change over time, which could shed light on their interaction when we combine the series for multivariate forecasting.

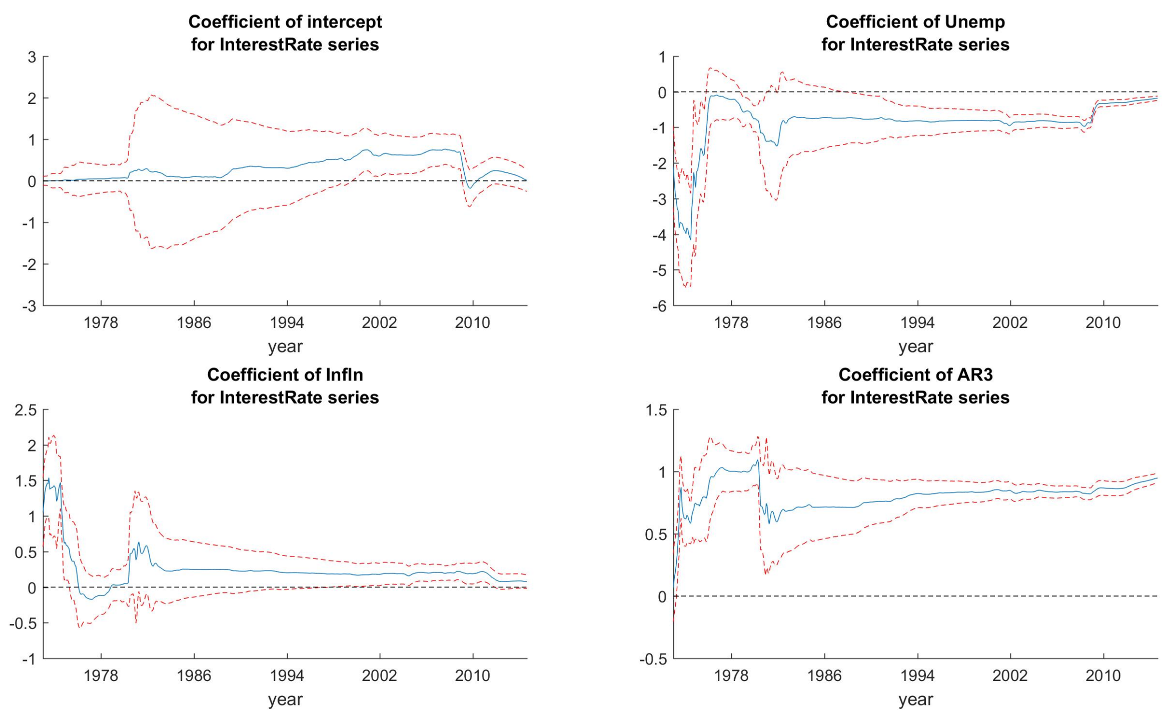

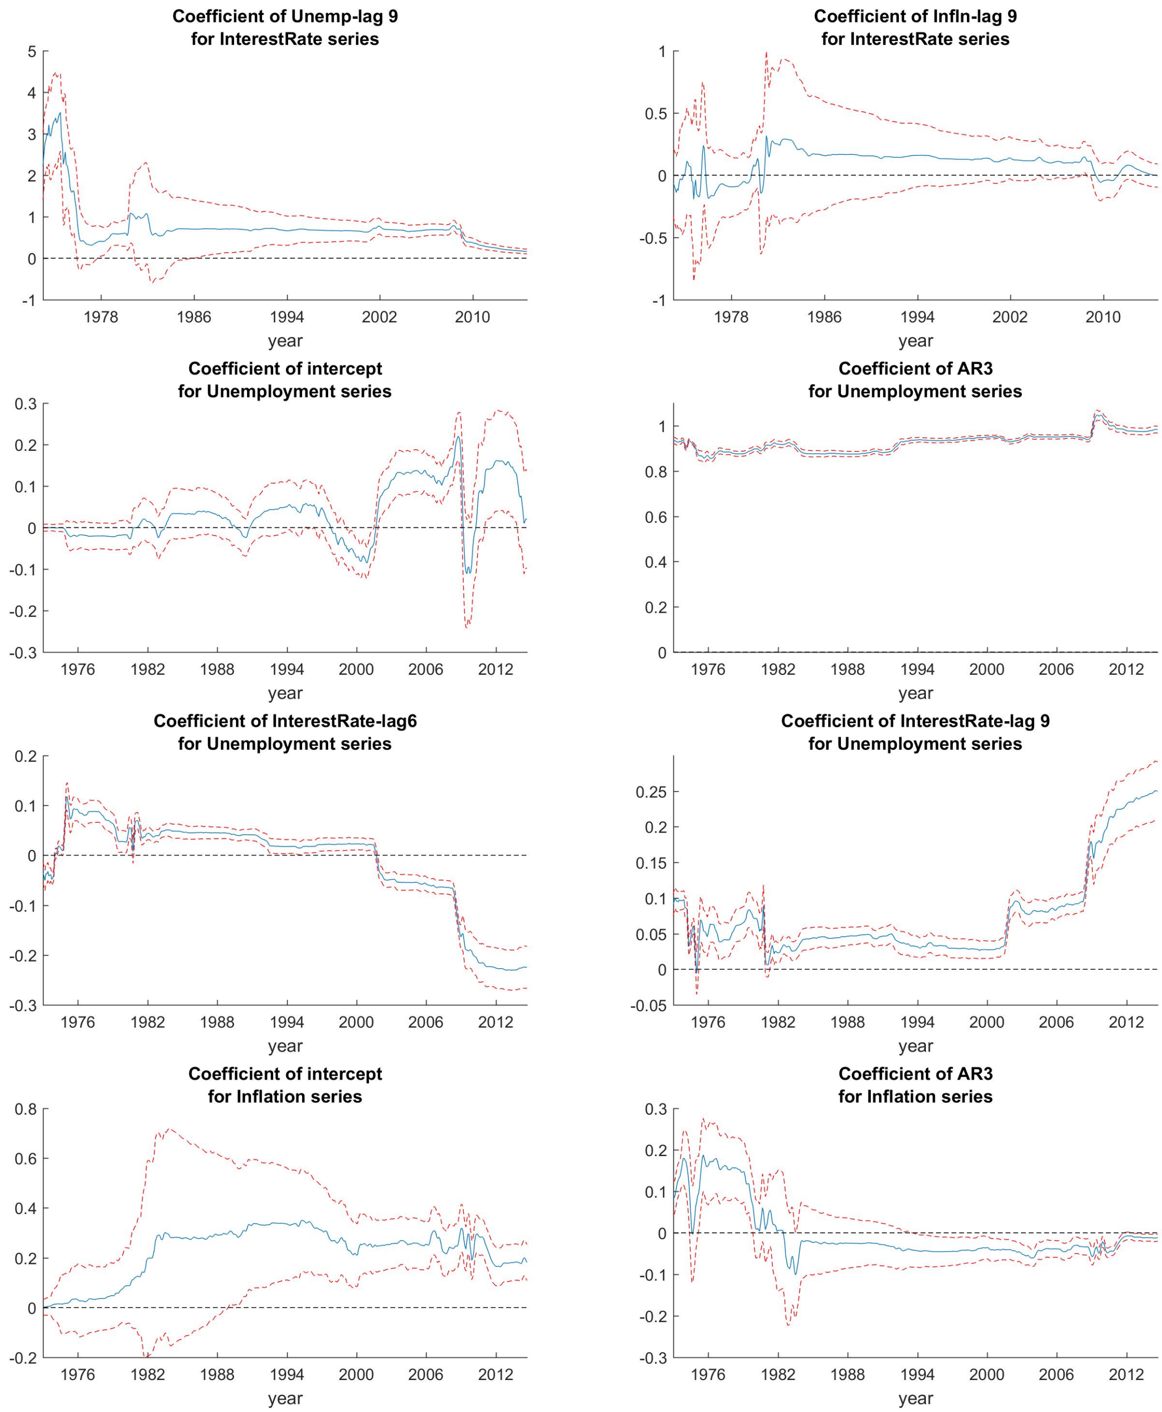
Coefficient of Unemp-lag 9
for InterestRate series
Coefficient of Infln-lag 9
for InterestRate series
Coefficient of intercept
for Unemployment series
Coefficient of AR3
for Unemployment series
Coefficient of InterestRate-lag6
for Unemployment series
Coefficient of InterestRate-lag 9
for Unemployment series
Coefficient of intercept
for Inflation series
Coefficient of AR3
for Inflation series
year

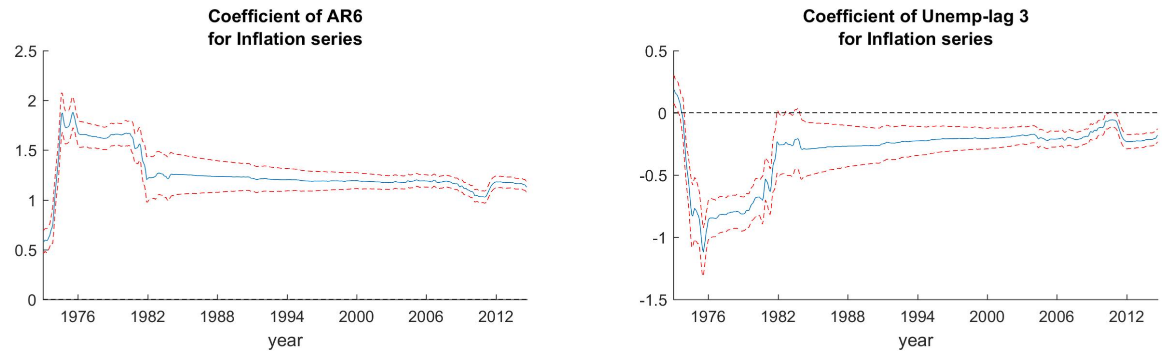

Figure 8: On-line trajectories of state variables for each series. Posterior mean is represented by the blue line and 95% credible interval is marked by red dotted lines.

Figure 8 illustrates some key features: (i) there is clear regime switching. In early 1980s, there is a sudden increase in volatility as experienced by many predictors. From late 1980s through early 2000s, the economy is relatively stable. After the 2008 financial crisis, many predictors exhibit changes in coefficient in either directions. However, except for a few, namely the intercept, interest rate with lag 6 and lag 9 for unemployment series, there is no significant increase in volatility post-2008. The identified regime switching demonstrates adaptivity of DDNMs. (ii) There are some predictors (AR(3) for interest rate, AR(3) and interest rate with lag 9 for unemployment, AR(6) for inflation) with entirely positive coefficients and some (unemployment with lag 3 for inflation, unemployment for interest rate) with negative coefficients. The rest change signs as time evolves, thus, affecting the response variables differently. (iii) There is no prolonged period for any predictor which the 95% credible interval covers zero. Except for interest rate with lag 6 and lag 9 for unemployment series which appear to move concurrently in different directions, there is no clear collinearity. This suggests that most predictors are significant in the context of this DDNM.

6.3 Impulse Response Analysis

As in the 8-variable model, Figure 9 shows impulse response trajectories for 1- and 2-year ahead horizons for each series. The shock remains as a 2% increase in interest rate. The difference between the predictive means with and without the shock is plotted.

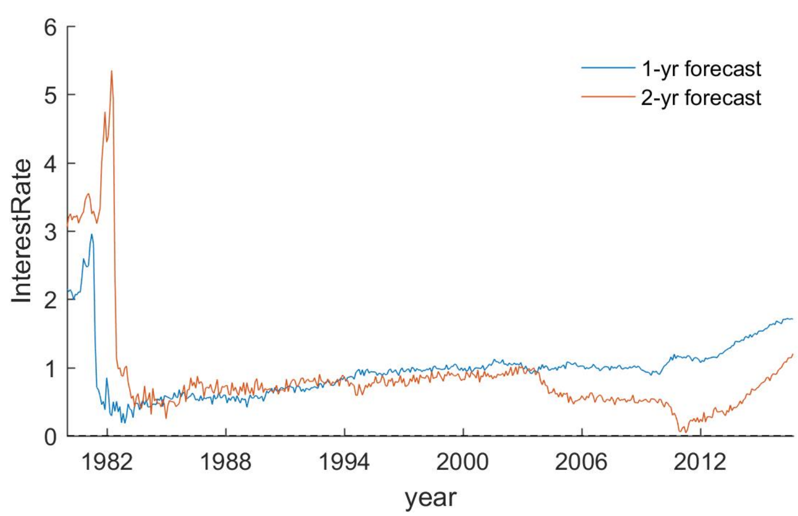

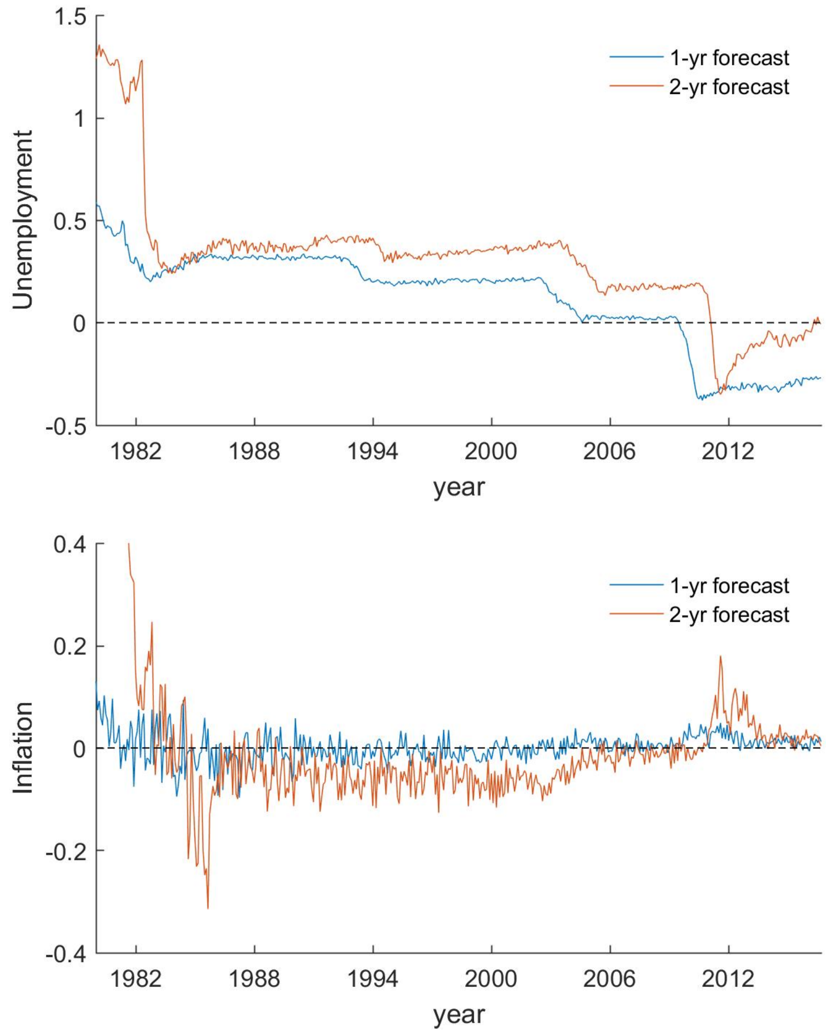

Figure 9: Impulse response trajectories for 1- and 2-year ahead horizons for US interest rate, unemployment and inflation from 1980 to 2016. The shock introduced is a 2% increase in interest rate.

As in the 8-variable model case, interest rate is predicted to respond positively to a 2% shock, although the extend of increase is smaller. Unemployment and inflation, however, show different trends. Before 2009, unemployment responds positively to an interest rate shock. Between 2009 and 2012, there is a sharp decrease, making the response negative. After 2012, 1-year forecast remains in the negative region while 2-year forecast bounce back to around zero. On the contrary, before 2009, inflation exhibits a long-run negative response to interest rate. This is reserved post-2009, as inflation starts to respond positively with interest rate. Again, the erratic forecast path of inflation reflects that its DDNM parameters lie outside the stationary region. Using accept-reject sampling method mentioned in Section 4.4 on the simulated parameters into the future is needed to address this issue.

The general trends of these variables resemble those under Nakajima and West [2013] latent threshold VAR model, but the sizes of the responses from DDNM are smaller than those from the latent threshold-VAR analysis. This arises from the different shock structures used. Both suggest that before 2009, an increase in interest rate would suppress inflation and raise unemployment, which aligns with macroeconomic theories. After 2009, however, the effect is ambiguous, sometimes even reversed.

6.4 Decision Analysis

The previous section has discovered a clear regime switch after the financial crisis. June 2005 and October 2011 as the two baselines fall into two regimes. We are interested in finding out if the model recommends different decisions at different baselines. The decision context and setting is the same as Section 5.

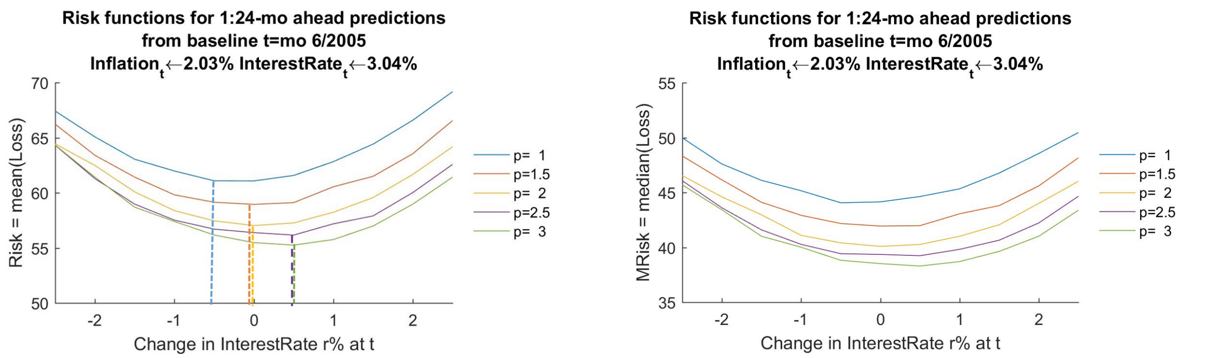

Figure 10: Mean and median loss for p=1:3 over 24-month ahead forecasting horizons from baseline June 2005 with the local minima marked by dotted lines.

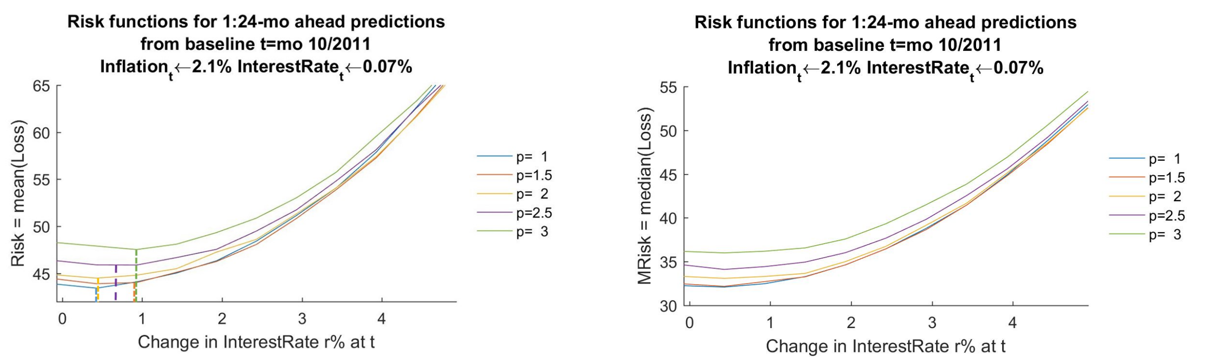

Figure 11: Mean and median loss for p=1:3 over 24-month ahead forecasting horizons from baseline October 2011 with the local minima marked by dotted lines.

Figures 10 and 11 again illustrate that as the inflation target increases, the local minimum of risk functions shifts to the right, suggesting an increase in interest rate. This is especially apparent in Figure 10: with baseline inflation at 2.03%, to reach a 3% target inflation, the risk function suggests a 0.5% increase in interest rate. To reach a 1% target inflation, the risk function suggests a 0.5% decrease in interest rate. Thus, the regime change does not impact the decision analysis and both 3-variable and 8-variable models recommend to increase interest rate to reach a higher inflation target.

As in Section 5.2, both figures illustrate a smaller median loss than mean loss, suggesting a right-tailed loss distribution. Another interesting point to note is the relative risk of different targets. In Figure 10, the risk function of a lower inflation target lies entirely above that of a higher target, while the reverse is true in Figure 11. This shows that at baseline June 2005, it is riskier to decrease interest rate in order to reach a lower inflation target than to increase interest rate for a higher target. On the other hand, with a nearly zero interest rate at baseline October 2011, increasing the interest rate slightly to achieve 1% inflation in 2 years is less risky as increasing it drastically to achieve higher inflation levels.

7 Summary Comments and Future Steps

DDNMs with the flexibility to customize model specifications for each univariate series, overlaying the ability to adapt the model over time, has natural application in macroeconomic time series analysis. By applying DDNMs to US macroeconomic data, we have explored the relationship between inflation as the policy target and interest rate as the control variable over the period of 1970 to 2016. The 8-variable model chosen by forward and backward variable selection to specify each univariate series captures richer information about the economic conditions. Even though the DDNM is sparse, the parental predictors and lagged predictors in a hierarchical structure give rise to complex interactions among variables. The 3-variable model adapted from Nakajima and West [2013] includes only interest rate, inflation and unemployment, hence reducing the dimensionality and model complexity drastically. This directly translates to smaller volatility in impulse response trajectories for the same forecast horizon. It also captures regime shifts better, clearly identifies three periods: (i) late 1970s to early 1980s with high volatility (ii) 1990s to early 2000s with stable economy and (iii) post-2009 with predictor coefficients moving in opposite directions. These regimes can also be spotted in the 8-variable model, but are often masked by confounding interactions among variables.

Another major focus of the study is decision analysis for the Federal Reserve's monetary policy. Using June 2005 and October 2011 as baselines of decision making, we find out that both models recommend a higher interest rate for a higher inflation target. This suggests that traditional monetary policy tools (such as a change in the Federal funds rate) may not behave exactly as suggested by conventional economy theory on impacting inflation. Multiple papers [Silvia et al., 2014, Balke and Emery, 1994a] find similar results, highlighting a change in relationship of the Federal funds rate with unemployment and inflation. In fact, the positive relationship between the Federal funds rate and inflation has become known as the ”price puzzle” [Sims, 1992, Balke and Emery, 1994b]. One possible explanation is the Federal Reserve systematically responding to signals of higher future inflation by raising the Federal funds rate, but not by enough to fully offset the subsequent inflation [Balke and Emery, 1994b]. This is especially apparent during the Great Inflation, as inflation and interest rate move together. To capture the response of the Federal Reserve to anticipated inflation, Christiano and Evansi [1994] propose adding commodity prices into the analysis to solve the price puzzle as they provide information about future inflation. This method can be included in future research.

Moreover, these discussions on the price puzzle focus on the 1960s and 1980s while our analysis suggest that the price puzzle is also prominent in the 2000s. In the 2000s, the Federal Reserve has practiced unconventional monetary policies. For instance, before increasing its target for the federal funds rate in June 2004, the Federal Open Market Committee (FOMC) used forward guidance – a sequence of changes in its statement language to signal that it was approaching the time at which a tightening of monetary policy was warranted [FRB, 2015]. As a result, expectation change caused by forward guidance leads to a structural change in the economy, but is not captured by the models. Hence, the question arises: How to capture interest rate expectation in the economy?

Moving forward, we can include more series into the analysis: commodity prices as suggested by Christiano and Evansi [1994] and long-term government bond yields to capture expectations about future inflation and interest rate. In view of the erratic projection paths for multiple series, we can consider methods such as accept-reject sampling to constrain parameters to ensure local stationarity in forecasting. There are also many aspects of the decision analysis that can be explored further. In this paper, we assume that a decision to set $\mathbf{r}_t = \mathbf{r}$ actually achieves that target value. The question is: whether setting the Federal funds target rate effective achieves the desired target $\mathbf{r}$. We could introduce a uncertainty to the intervention. In addition, in our analysis, we constrain interest rate to be non-negative. The past recession caused a plunge in interest rate. Does the relationship among inflation, unemployment and interest rate change around zero lower bound (ZLB) or even pass ZLB? With the recent lifting from zero interest-rate policy, it would be interesting to analyze the impact of this policy and potentially broaden the analysis to consider negative interest rate.

8 Appendix

8.1 Data Description

The data we used comes from the Federal Reserve Economic Data and Bloomberg terminal. We collected monthly data starting from January 1965 to September 2016. We have identified 8 relevant macroeconomic series. They are detailed in Table 3 below.

Notes on some of the series:

Inflation: This is the main response variable $\mathbf{x}_t$, in another words, a policy target of the Federal Reserve. We used the annual percentage change in the CPI index.

Interest Rate: This is $\mathbf{r}_t$, a policy instrument of central banks as a potential target to control. The effective Federal funds rate we used is the best proxy to the interest rate the Federal Reserve aims to control. It refers to the overnight interest rate at which depository institutions trade federal funds with each other overnight. The Federal Open Market Committee (FOMC) meets eight times a year to determine the federal funds target rate. Here, we assume that the effective Federal funds rate is directly adjustable and set $\mathbf{r}_t = \mathbf{r}$ actually achieves that target value.

Investment: We used ISM Manufacturing Index on business new orders to represent investment level. The index ranges in value from 0 to 100 with any number above 50 considered as expansionary. Thus, we adjusted the data to center around the critical 50 level.

Money Supply: The Federal Reserve releases two types of money supply data: M1 Money Stock (M1NS) and M2 Money Stock (M2NS). M1 is defined as the sum of currency held by the public and transaction deposits at depository institutions (which are financial institutions that obtain their funds mainly through deposits from the public, such as commercial banks, savings and loan associations, savings banks, and credit unions). M2 is defined as M1 plus savings deposits, small-denomination time deposits (those issued in amounts of less than \$100,000), and retail money market mutual fund shares. We include both in our analysis.

Some other series we have considered are GDP, the Standard & Poor's 500 index, and money velocity. For GDP and money velocity, only quarterly data are available, which give us insufficient data points for the analysis. S&P 500, on the other hand, introduces high volatility. Hence, we do not include these series in our model.

Variable Name	Index Code	Description	Seasonal Adjustment?	Treatment
Inflation	CPILFENS	Consumer Price Index for All Urban Consumers: All Items Less Food and Energy	No	Annual percentage change
Interest Rate	FEDFUNDS	Effective Federal Funds Rate	No	No treatment
Wage	CEU0500000008	Average Hourly Earnings of Production and Nonsupervisory Employees: Dollars per Hour	No	Annual percentage change
Unemployment	UNRATE	Civilian Unemployment Rate	Yes	No treatment
Consumption	PCE	Personal Consumption Expenditures: Billions of Dollars	Yes	Annual percentage change
Investment	NAPMNEWO	ISM Manufacturing Index on Business New Orders	No	Index minus 50
M1 Money Supply	M1NS	M1 Money Stock: Billions of dollars	No	Annual percentage change
M2 Money Supply	M2NS	M2 Money Stock: Billions of dollars	No	Annual percentage change

Table 3: Data description

8.2 Risk Function Calculation

The Bayes risk function is

$$\begin{aligned} R_t(\mathbf{r}) &= E[L_t(\mathbf{y}_{t+1:t+h})|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t] \\ &= \sum_{k=1:h} \delta_k E[(\mathbf{x}_{t+k} - \mathbf{x}_{t+k-1})'(\mathbf{x}_{t+k} - \mathbf{x}_{t+k-1})|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t] \\ &\quad + \rho E[(\mathbf{x}_{t+h} - \boldsymbol{\tau})'(\mathbf{x}_{t+h} - \boldsymbol{\tau})|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t]. \end{aligned} \tag{6}$$

Note that

$$\begin{pmatrix} \mathbf{x}_{t+1}-\mathbf{x}_t \\ \mathbf{x}_{t+2}-\mathbf{x}_{t+1} \\ \mathbf{x}_{t+3}-\mathbf{x}_{t+2} \\ \vdots \\ \mathbf{x}_{t+h-1}-\mathbf{x}_{t+h-2} \\ \mathbf{x}_{t+h}-\mathbf{x}_{t+h-1} \\ \mathbf{x}_{t+h}-\boldsymbol{\tau} \end{pmatrix} = \begin{pmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ -\mathbf{I} & \mathbf{I} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I} & \mathbf{I} & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & -\mathbf{I} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & -\mathbf{I} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{x}_{t+1} \\ \mathbf{x}_{t+2} \\ \mathbf{x}_{t+3} \\ \vdots \\ \mathbf{x}_{t+h-2} \\ \mathbf{x}_{t+h-1} \\ \mathbf{x}_{t+h} \end{pmatrix} - \begin{pmatrix} \mathbf{x}_t \\ \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{0} \\ \boldsymbol{\tau} \end{pmatrix} = \mathbf{E}\mathbf{z}_{th} - \mathbf{e}_t$$

where $\mathbf{I}$ is the $q_x \times q_x$ identity matrix and $\mathbf{0}$ the $q_x \times q_x$ zero matrix, and with

$$\mathbf{z}_{th} = \begin{pmatrix} \mathbf{x}_{t+1} \\ \mathbf{x}_{t+2} \\ \mathbf{x}_{t+3} \\ \vdots \\ \mathbf{x}_{t+h-2} \\ \mathbf{x}_{t+h-1} \\ \mathbf{x}_{t+h} \end{pmatrix}, \quad \mathbf{E} = \begin{pmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ -\mathbf{I} & \mathbf{I} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I} & \mathbf{I} & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & -\mathbf{I} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & -\mathbf{I} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{I} \end{pmatrix} \quad \text{and} \quad \mathbf{e}_t = \begin{pmatrix} \mathbf{x}_t \\ \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{0} \\ \boldsymbol{\tau} \end{pmatrix}.$$

Given $\mathcal{D}_t$, we of course know $\mathbf{x}_t$ and have fixed the value $\boldsymbol{\tau}$ at target, so that $\mathbf{e}_t$ is a known vector depending on these values. Now introduce the $(h+1) \times (h+1)-$ diagonal matrix $\mathbf{G} = \text{diag}(\delta_1, \delta_2, \ldots, \delta_h, \rho)$. It follows that the quadratic loss function in eqn. 6 is given by

$$L_t(\mathbf{y}_{t+1:t+h}) = (\mathbf{E}\mathbf{z}_{th} - \mathbf{e}_t)'\mathbf{G}(\mathbf{E}\mathbf{z}_{th} - \mathbf{e}_t). \tag{7}$$

Under the forecast distribution $p_t(\mathbf{y}_{t+1:t+h}|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t)$, suppose that the implied hq_x-mean vector and $hq_x \times hq_x-$variance matrix of $\mathbf{z}_{th}$ are denoted by

$$\mathbf{f}_{th}(\mathbf{r}) = E(\mathbf{z}_{th}|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t) \quad \text{and} \quad \mathbf{Q}_{th}(\mathbf{r}) = E(\mathbf{z}_{th}|\mathbf{r}_t = \mathbf{r}, \mathcal{D}_t).$$

After taking expectations of eqn. 7, the resulting risk function of eqn. 6 becomes

$$R_t(\mathbf{r}) = [\mathbf{E}\mathbf{f}_{th}(\mathbf{r}) - \mathbf{e}_t]'\mathbf{G}[\mathbf{E}\mathbf{f}_{th}(\mathbf{r}) - \mathbf{e}_t] \; + \; \text{trace}[\mathbf{G}\mathbf{E}\mathbf{Q}_{th}(\mathbf{r})\mathbf{E}']. \tag{8}$$

This simplified form is then used to calculate $R_t(\mathbf{r})$ and find $\mathbf{r}$ that minimizes the risk. In models that involve $\mathbf{r}$ in complicated ways such that we are not able to solve it analytically, we can simulate specify a range of values of $\mathbf{r}$ and simulate predictions to get a Monte Carlo sample of the loss distribution in eqn. 7. $R_t(\mathbf{r})$ is the mean of that distribution. We can also look at other values of interest of the distribution, such as median or tail values.

References

G. Amisano and J. Geweke. Prediction using several macroeconomic models. *European Central Bank*, 1537, 2013.

N. S. Balke and K. M. Emery. The Federal funds rate as an indicator of monetary policy: Evidence from the 1980s. *Economic Review*, Q I:1–15, 1994a.

N. S. Balke and K. M. Emery. Understanding the price puzzle. *Economic Review*, Q IV:15–26, 1994b.

J. Boivin and M. P. Giannoni. Has monetary policy become more effective? Working Paper 9459, National Bureau of Economic Research, January 2003.

M. E. Christiano, Lawrence J. and C. Evansi. The effects of monetary policy shocks: Evidence from the flow of funds. Working Paper 4699, National Bureau of Economic Research, April 1994.

M. Forni, M. Hallin, M. Lippi, and L. Reichlin. The generalized dynamic factor model: One-sided estimation and forecasting. *Journal of the American Statistical Association*, 100:830–840, 2005.

FRB. What is forward guidance and how is it used in the Federal reserve's monetary policy? *Federal Reserve Board*, December 2015.

G. Koop and D. Korobilis. Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics*, 3:267–358, 2009.

J. Nakajima and M. West. Bayesian analysis of latent threshold dynamic models. *Journal of Business and Economic Statistics*, 31(2):151–164, 2013.

W. Poole and D. C. Wheelock. Stable prices, stable economy: Keeping inflation in check must be no. 1 goal of monetary policymakers. *The Regional Economist*, January 2008.

G. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72:821–852, 2005.

J. E. Silvia, A. Iqbal, and B. Zachary. Is the Federal funds target rate effective. Working paper, Wells Fargo, March 2014.

C. Sims. Macroeconomics and reality. *Econometrica*, 48:1–48, 1980.

C. A. Sims. Interpreting the macroeconomic time series facts: The effects of monetary policy. *European Economic Review*, 36:975–1000, June 1992.

J. Stock and M. Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97:147–162, 2002.

Z. Zhao, M. Xie, and M. West. Dynamic dependence networks: Financial time series forecasting and portfolio decisions (with invited discussion). *Applied Stochastic Models in Business and Industry*, 32:311–332, 2016.