Working Paper Series

Drifts, Volatilities, and Impulse Responses Over the Last Century

WP 14-10

Pooyan Amir-Ahmadi
Goethe University Frankfurt

Christian Matthes
Federal Reserve Bank of Richmond

Mu-Chun Wang
University of Hamburg

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

Drifts, Volatilities, and Impulse Responses Over the Last Century

Pooyan Amir-Ahmadi Christian Matthes
Mu-Chun Wang

Working Paper No. 14-10*

April 7, 2014

Abstract

How much have the dynamics of U.S. time series and in particular the transmission of innovations to monetary policy instruments changed over the last century? The answers to these questions that this paper gives are "a lot" and "probably less than you think," respectively. We use vector autoregressions with time-varying parameters and stochastic volatility to tackle these questions. In our analysis we use variables that both influenced monetary policy and in turn were influenced by monetary policy itself, including bond market data (the difference between long-term and short-term nominal interest rates) and the growth rate of money.

JEL Classification: C50, E31, N12
Keywords: Bayesian VAR, Time variation, U.S. monetary policy

*Affiliations: Goethe University Frankfurt (Amir-Ahmadi), Federal Reserve Bank of Richmond (Matthes), and University of Hamburg (Wang). We would like to thank Luca Benati, Jochen Günther, Andrew Owens, Pierre Sarte, Felipe Schwartzman, Alexander Wolman as well as seminar participants at CREATES (Aarhus), UPF, Banque de France, Kiel Institute for the World Economy, the Bundesbank, and the FRBR-UVa Research Jamboree for helpful comments. Parts of this paper were written while Matthes was visiting the Bundesbank, whose hospitality is gratefully acknowledged. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. The online appendix is available at `https://sites.google.com/site/christianmatthes/OnlineAppendix.pdf`

1 Introduction

We study over 100 years of US economic data on inflation, real output, short-term and long-term nominal interest rates as well as money growth through the lens of a time-varying parameter model to assess how the economy has changed and how the impact of monetary policy shocks has evolved over time. Doing so gives us the opportunity to study very different economic episodes using a single model and to ask if policy measures (in our case unexpected changes in monetary policy) that had a certain effect at one point can be reliably predicted to have a similar effect at a different point in time.

Our sample features two World Wars, the Great Depression, the recent financial crisis and the associated recession, technological revolutions and the founding of the Federal Reserve, so there is ample reason to believe that indeed the dynamics and co-movement of the variables we consider have changed over time.

To gauge how much the US economy has changed during our sample, we start off by calculating different measures of time variation implied by our multivariate time-varying parameter model - variation in persistence, volatility, long-run averages, and co-movement. We find that along all those lines there is substantial variation in the economy. The correlation structure between our variables of interest has changed dramatically. To give a few examples, the correlation between inflation and short-term interest rates seems to have undergone two major structural breaks (see also Cogley, Surico & Sargent (2012)). The correlation between money growth and inflation also changes dramatically, undermining efforts by economists to use growth in monetary aggregates to forecast inflation. Our model identifies the start of the Paul Volcker chairmanship at the Federal Reserve as a time of major breaks in many of our measures of time variation.

One of the most sought-after questions in macroeconomics is that of the effects of unanticipated changes in policy instruments, in particular for the case of monetary policy (Christiano, Eichenbaum & Evans (1999)). We want to tackle that question in the context of using a long sample in which the conduct of monetary policy has changed dramatically. We identify a monetary policy shock using sign restrictions, thus allowing us to remain silent on the exact choice of the monetary policy instrument. We find that effects of an 'average' (one standard deviation) shock have changed dramatically. These changes can be driven by both changes in the average size of a shock (changes in the standard deviation) and in the dynamic responses to shocks. We disentangle these possible causes and find that the size of the innovation is the major driver of the changes in the effects of a monetary policy innovation. Nonetheless, we also find meaningful changes in the dynamic response to shocks.

Our work is related to a growing literature on time-varying VARs. Papers that also use VARs with time-varying parameters and stochastic volatility include Cogley & Sargent (2005), who were the first to use this class of models; Primiceri (2005), who first identified monetary policy shocks in this class of models; Canova & Gambetti (2009), who used sign restrictions to identify monetary policy shocks just as we do; and Gali & Gambetti (2009). Ritschl & Woitek (2000) employ a time-varying parameter VAR studying the role of monetary forces during the Great Depression. A related strand of the literature studies VARs when the dynamics are governed by a discrete Markov chain. This approach has been pioneered by Sims & Zha (2006).

Recently, there has also been a growing interest in using time-varying parameter models to study longer time series. Amir-Ahmadi (2009) employs a dynamic factor model with time-varying parameters for a long U.S. sample to study the role of credit shocks over time. Sargent & Surico (2011) study the quantity theory through such a lens, for example. D'Agostino & Surico (2011) use time-varying VARs on long US time series to study changes in the properties of inflation forecasts. Benati (2010) focuses on the relationship between inflation and unemployment for the US since the late 19th century. Kliem, Kriwoluzky & Sarferaz (2013) explore the relationship of inflation and fiscal policy using US historical data. Benati & Lubik (2012) study the dynamics of inventories using long US time series, while Nason & Tallman (2013) explore the role of economic and financial factors across the business cycle. There is also a literature studying the role of the monetary transmission mechanism during the Great Depression: Sims (1999) contrasts the dynamic effects of identified monetary policy shock during the Great Depression with the post-WWII period and finds no substantive difference across those periods. Amir-Ahmadi & Ritschl (2013) study the role of monetary policy during the Great Depression in a large-dimensional factor model.
Our approach differs from these papers because we use a larger number of observables (in particular, we include yield curve information) and have a different goal: While the papers mentioned above are focused on specific time episodes of interest or specific aspects of changes in the dynamics (or, in the case of Benati & Lubik (2012), a different set of variables), we want to uncover broad stylized facts concerning changes in US economic dynamics and the impact of monetary policy.

2 The Model

We are interested in modeling the dynamics of the following vector of observables:

$$y_t = \begin{pmatrix} \Delta gdp_t \\ \pi_t \\ i_t^s \\ spread_t \\ \Delta m_t \end{pmatrix} \tag{1}$$

where Δgdp_t is the one-year difference in the log of real output, π_t is the one-year inflation rate, i_t^s is a short-term nominal interest rate, $spread_t$ is the spread between our short-term nominal interest rate and a long-term nominal interest rate, and finally, Δm_t is the one year difference in the log of a monetary aggregate. Our benchmark monetary aggregate is the monetary base[1]. Details on the data can be found in the data section.

We borrow our model from Primiceri (2005)[2] and assume that our observables follow a time-varying VAR of the following form:

$$y_t = m_t + \sum_{j=1}^{L} A_{j,t} y_{t-j} + e_t \tag{2}$$

where the intercepts μ_t, the autoregressive matrices $A_{j,t}$, and the covariance matrix Ω_t of e_t are allowed to vary over time. We set the number of lags $L = 2$. To be able to parsimoniously describe the dynamics of our model, we define $X_t' \equiv I \otimes (1, y_{t-1}' ..., y_{t-L}')$ and rewrite (2) in the following

[1]In the online appendix we study a model using M2 instead.

[2]The modeling assumptions we make are widely used in empirical macroeconomics. An overview of the methods used and assumptions made in this literature is given by Koop & Korobilis (2010).

state space form[3]:

$$y_t = X_t' \theta_t + e_t \tag{3}$$

$$\theta_t = \theta_{t-1} + u_t \tag{4}$$

The observation equation (3) is a more compact expression for (2). The state equation (4) describes the law of motion for the intercepts and autoregressive matrices. The covariance matrix of the innovations in equation (3) is modeled after Primiceri (2005):

$$e_t = \Lambda_t^{-1} \Sigma_t \varepsilon_t \tag{5}$$

The covariance state Λ_t is a lower triangular matrix with ones on the main diagonal and representative non-fixed element λ_t^i. Σ_t is a diagonal matrix with representative non-fixed element σ_t^j. The dynamics of the non-fixed elements of Λ_t and Σ_t are given by:

$$\lambda_t^i = \lambda_{t-1}^i + \zeta_t^i \tag{6}$$

$$\log \sigma_t^j = \log \sigma_{t-1}^j + \eta_t^j \tag{7}$$

To conclude the description of our model, we need to make distributional assumptions on the innovations ε_t, u_t, η_t and ζ_t, where η_t and ζ_t are vectors of the corresponding scalar innovations in the elements of Σ_t and Λ_t. We assume that all these innovations are normally distributed with covariance matrix V, which we, following Primiceri (2005), restrict

[3] I denotes the identity matrix.

as follows:

$$V = Var\left[\begin{pmatrix} \varepsilon_t \\ u_t \\ \zeta_t \\ \eta_t \end{pmatrix}\right] = \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix} \tag{8}$$

S is further restricted to be block diagonal, which simplifies inference. We estimate this model using the Gibbs sampling algorithm described in Del Negro & Primiceri (2013)[4]. A summary of this algorithm can be found in the appendix.

3 Data

In this section we describe the construction of our data (plotted in Figure 1). We use quarterly U.S. data covering the period from the first quarter 1876 to the second quarter of 2011. This time span is of specific interest as it covers the pre-Fed period as well as all chairmanships, which represent potentially different monetary policy regimes. Furthermore, the period covers 29 recession periods of different duration and depth (according to the NBER). The sample period covers two world wars and additional conflicts the United States have been engaged in, as well as several financial crises and substantial financial market deregulation during the past four decades. In the past century, the U.S. economy has also experienced four major "great events," namely the Great Depression, the Great Inflation, the Great Moderation, and the recent Great Recession. The role of monetary policy leading up to and during those events, either as a source or remedy, has been of specific interest. Therefore, this sample seems appropriate to analyze the evolution of the monetary trans-

[4]We use 500000 draws.

mission mechanism and the (in)stability of both (i) the variables affected by monetary policymaking and (ii) the variables affecting monetary policy decisionmaking. In this paper we follow Sargent & Surico (2011) and take the historical data at face value[5]. While some papers have recently started to take mismeasurement issues seriously when estimating time-varying parameter models on long historical time series, these papers only focus on one observable (see for example Cogley & Sargent (2014)). It is an open (and very interesting) question how those approaches could be efficiently generalized to multivariate models.

3.1 Annual Output Growth

Our output growth series is obtained by splicing two different real output series covering different time spans. We use real GNP series as constructed by Balke & Gordon (1986) from the first quarter of 1876 to the fourth quarter of 1947. After that, we use the real GDP series provided by the St. Louis Fed FRED database covering the first quarter of 1948 to the second quarter of 2011. The spliced series are transformed in logs and then we take year-on-year differences.

3.2 Annual Inflation Rate

The corresponding annual inflation rate is also based on the combination of two different series on the output deflator. Again the first part comes from Balke & Gordon (1986) covering the period 1876Q1-1947Q4. The second part of the series comes from the St. Louis Fed FRED database covering the time span 1948Q1-2011Q2. Again we transform the data

[5] Our data construction closely follows Sargent & Surico (2011) for the variables that appear in both papers. We deviate from them in the choice of the maturity and construction of our short-term interest rate measure.

into year-on-year growth rates.

3.3 Short-Term Interest Rate

The short term interest rate plays the role of a potential direct or indirect monetary policy instrument for at least a substantial part of the time span we analyze. There is no single series on shorter interest rates at quarterly frequency for the full sample, which requires constructing a series based on several data sources reflecting short-term borrowing conditions. From 1920Q1-2011Q2 we use data on the 90-day T-Bill rate from the secondary market.. Prior to that we backcast the series including as regressors data on call money rates and commercial paper rates. These two series and our target short term interest rate series are all available at monthly frequency. Specifically, we regress 90-day T-Bill rate on call money rates and commercial paper rates based on a sample running from February 1920 to April 1934. Combining the resulting coefficients with our regressors we can backcast our target series back to the first quarter of 1876. This way we interpolate backward the missing observations for the 90-day T-Bill rate. We thus avoid using the six-month short term interest rate, which would lead to a maturity mismatch combining the three-month and six-month rates. Furthermore, we prefer the shorter maturity rate as a potential monetary policy instrument. We use annualized interest rates throughout.

3.4 Long-Term Interest Rate

As for the term spread, we employ the difference between a constructed measure of the long-term interest rate and the short-term interest rate described in the previous section. The lack of a consistent long-term government benchmark interest rate requires the combination and backcasting of three indicators. From 1920Q1-2011Q2 we use data on the 10-year government bond yields at constant maturities. Prior to that, we backcast the series including as regressors data on railroad bond yields (high grade) and a railroad bond yields index. These two series and our target long-term interest rate series are all available at monthly frequency. Specifically, we regress 10-year government bond yields at constant maturities on railroad bond yields (high grade) and railroad bond yields index based on a sample running from February 1920 to April 1934. Combining the resulting coefficients with our regressors we can backcast our target series back to the first quarter in 1876. The long-term interest rate is expressed in annual terms.

3.5 Annual Base Money Growth

The monetary base measure to represent a direct or indirect monetary policy instrument is compiled by two series. The first part of the sample from 1876Q1-1959Q4 comes from Balke & Gordon (1986) and the second part from the FRED database covering 1919Q1-2011Q2.

4 Prior Choice

We choose priors in a way to stay as close as possible to the previous literature, while taking into account our larger sample. We use data from 1876:Q1 to 1913:Q4 to initialize the priors using a fixed coefficient VAR, similarly to Primiceri (2005). The most important prior in this class of models seems to be the prior for Q, the covariance matrix of the residuals that enter the law of motion for θ. We assume that Q, which governs the amount of time variation in the VAR coefficients, follows an inverse Wishart distribution with the following parameters:

$$Q \sim IW(\kappa_Q^2 * 152 * V(\theta_{OLS}), 152) \tag{9}$$

where the prior degrees of freedom is set to 152, which is the length of our training sample and $\kappa_Q = 0.01$ is the tuning parameter to parameterize the prior belief about the amount of time variation. Primiceri (2005) uses exactly the same approach to set his prior. Choosing the same approach allows us to keep our results comparable to his.[6]

The other priors are also set according to Primiceri (2005), adjusting for the larger size of our vector of observables. In contrast to Cogley & Sargent (2005), we do not impose the prior that the companion matrix of our VAR only has eigenvalues smaller than 1 in absolute value.

5 Results

5.1 A Useful Way to Look at Our VAR

To facilitate the discussion of the sources of volatility and persistence, we rewrite our VAR in companion form:

$$y_t = \mu_t + \sum_{j=1}^{L} A_{j,t} y_{t-j} + e_t$$

[6]The online appendix shows how our results change when we change this prior, in particular the scaling factor entering the inverse-Wishart distribution. There we find that a-priori allowing for substantially more time variation leads to forecasts that most economists would deem less sensible than the ones we will present in a subsequent section.

in the first order companion form. Note that Ω_t is the time-varying covariance matrix of e_t. We define $\mathbf{Y}_t \equiv (y_t', \ldots, y_{t-L+1}')'$, $\mu_t \equiv (m_t', 0, \ldots, 0)'$, $\mathbf{e}_t \equiv (e_t', 0, \ldots, 0)'$ and

$$\mathbf{A}_t = \begin{bmatrix} A_{1,t} & A_{2,t} & \cdots & A_{L-1,t} & A_{L,t} \\ I_K & 0 & \cdots & 0 & 0 \\ 0 & I_K & & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & I_K & 0 \end{bmatrix}$$

then we can recast the original VAR(L) into a VAR(1) using the following companion form

$$\mathbf{Y}_t = \mu_t + \mathbf{A}_t \mathbf{Y}_{t-1} + \mathbf{e}_t$$

5.2 A First Glance at Time Variation

The rest of the paper describes various ways to summarize the amount of time variation we find. As a first pass, though, we find it helpful to look at the raw results to see what patterns of time variation emerge. In Figure 2 we plot the median estimates of all elements in μ_t, $A_{1,t}$, and $A_{2,t}$. Our model is able to capture very different patterns of time variation: fixed coefficients, small (relative to the size of the coefficient) time variation, or large shifts in parameters throughout time after periods in which those parameters have been stable.

5.3 Sources of Volatility

Volatility in time series models can be traced back to two sources: the innovations (or unpredictable components) that influence the time series of interest and the systematic response to those innovations. To make this point, consider a univariate AR(1) model with Gaussian innovations:

$$z_t = \rho z_{t-1} + w_t,\, w_t \sim N(0, \sigma_w^2) \tag{10}$$

Then the j-step ahead conditional variance is given by

$$Var_t(z_{t+j}) = \sigma_w^2 \sum_{k=1}^{j} \rho^{2(j-k)} \tag{11}$$

We can see that the volatility of this process is fully characterized by the autoregressive coefficient and the variance of the innovation. The next two sections present a similar characterization for our time-varying VAR. The objects corresponding to ρ in the multivariate context are the $\mathbf{A}_t$ matrices, which are high dimensional. To study dynamics, we can focus on the eigenvalues, but even those are large in number (given that they vary over time). The section below therefore focuses on the largest eigenvalue in absolute value. This object does not fully characterize the effects of time variation in persistence on volatility, but it does give an idea about whether or not our estimated model features (locally) unstable dynamics, which in turn will have an effect on volatility.

5.4 Are There Explosive Dynamics in U.S. Time Series?

We study the probability of matrix A_t having eigenvalues larger than 1 in our sample by checking the draws of A_t that are generated by our Gibbs sampler. We can do this because, as mentioned before, we do not follow Cogley & Sargent (2005) and impose conditions on the eigenvalues of the companion matrix of our VAR. The upper left panel of Figure 3 shows this probability.

The average level of the probability until the 1940s is quite high, reaching 0.6. The probability drops first around 20 percentage points at the end of WWII. It rises again from .42 up to 50 percent until the end of the 1970s. The second big decrease in this probability following the Volcker disinflation could be interpreted in terms of a structural model in which agents have to learn about the true data-generating process (DGP): Cogley, Matthes & Sbordone (2012) show that times in which beliefs of private agents are far away from the DGP can lead to explosive dynamics, whereas the probability of explosive eigenvalues falls as beliefs move closer to the true DGP. An alternative structural model that can give temporarily explosive dynamics is given by Bianchi & Ilut (2013).

Despite the fact that high probability of explosiveness can be found in various periods in the history, the upper right panel of Figure 3 shows that the absolute value of those eigenvalues larger than 1 is only slightly larger than 1. This means that even if the economy is temporarily explosive, it takes a long time for the economy to become noticeably unstable. This also confirms conventional wisdom concerning the kind of stationarity restrictions used by Cogley & Sargent (2005). There is a substantial posterior probability of having explosive eigenvalues, making estimation algorithms with this restriction slow to converge. At the same time, the restriction itself is not far from being met for large parts of post-WWII data in the sense that the estimated eigenvalues are not far from 1[7]. To further assess how important this restriction is and what the exact pattern of locally explosive behavior implied by our estimates is, we study the distribution of time a draw of companion form matrices implies locally explosive dynamics. Do some draws always imply explosive behavior while many others never do? The lower panel of Figure 3 answers this question. We see that this distribution has its mode around 0, but the difference between the frequency at the mode and the frequency for other fractions of the sample is not large. Most of our draws feature both prolonged periods of stable behavior and unstable behavior.

5.5 Examining Stochastic Volatility

Next, we study the estimated volatilities of the innovations hitting our model. We focus on the square roots of the diagonal elements of $\Omega_t = Var(e_t)$, which incorporate both time variation in Σ_t and Λ_t.

Figure 7 plots the median as well as 18th and 86th percentile bands of these time-varying standard deviations of the reduced-form residuals. The residuals associated with real GDP growth, inflation and money growth are substantially more volatile during the first part of our sample, whereas the residuals of short-term interest rates and the spread are more volatile after World War II, in particular around the Vocker disinflation of 1980. The residual in money growth, on the other hand, does

[7] Given that we also use pre-WWII data, this approach would be harder to defend for our application.

not have a substantially larger volatility around the Volcker disinflation. The pre-WWII pattern for some variables may be partly explained by the inferior data quality prior to WWII, where measurement error seemed to be much more of a problem (for GDP, see Romer (1986)). One important take-away from this exercise is that relative to the decrease in volatility of real GDP and inflation after the Great Depression and World War II, the so-called "Great Moderation" is almost invisible in the estimated volatilities.[8]

The volatility of the forecast error in the equation for the spread shows a discrete jump in 1980. Interestingly, while average volatility in that forecast error has come down in the 1980s and 1990s, the levels remain elevated relative to pre-1980 values. Our model implies that one-quarter ahead forecasts of the slope of the yield curve have thus become less precise since 1980, an interesting hypothesis for future work.

5.6 Time t Approximations to Moments of Forecasts

To analyze the estimated time variation further, we ask what first and second moments of our observables would be if the dynamics of the observables were governed by parameter estimates that are fixed at the level estimated at one particular time t[9]. Since we do not impose stationarity on our VAR, we cannot compute the time t unconditional moments. Instead, we compute time t forecast moments for different forecast horizons, which do not require the time t estimates of the companion form matrix of the VAR having all eigenvalues (except for the eigenvalue asso-

[8] The "Great Moderation" refers to decreases of volatility in observables, not necessarily residuals, but it seems natural to expect part of this decrease to be reflected in residuals with smaller variance.

[9] Cogley & Sargent (2005) have used this approach to great effect.

ciated with the intercept) being less than 1 in absolute value.

Consider the VAR(1) representation described above. Forecast moments of y_{t+h} for a time-invariant VAR are given in Lütkepohl (2009). Denote $\Phi_{i,t} = J\mathbf{A}_t^i J'$ and $J = [I_K \; 0 \; 0 \cdots 0]$. The h-step ahead forecast mean is

$$E_t \left[y_{t+h}\right] = J \left[\left(\mathbf{I} + \mathbf{A}_t + \cdots + \mathbf{A}_t^{h-1}\right) \mu_t + \mathbf{A}_t^h \mathbf{Y}_t\right]$$

The h-step ahead forecast covariance Var_{t+h}, which is also the mean squared forecast error covariance, is given by

$$Var_t[y_{t+h}] = \sum_{i=0}^{h-1} \Phi_{i,t} \Omega_t \Phi_{i,t}'$$

The h-step ahead forecast error variance decomposition is

$$\omega_{jk,h,t} = \sum_{i=0}^{h-1} \left(\iota_j' \Theta_{i,t} \iota_k\right)^2 / Var_t[y_{t+h}]$$

where ι_k is the k-the column of I_K, $\Theta_{i,t} = \Phi_{i,t} P_t$ and P_t is a lower triangular matrix with $\Omega_t = P_t P_t'$[10]. $\omega_{jk,h,t}$ denotes the h-step ahead forecast error variance of variable j, accounted for by $e_{k,t}$ innovations at time t. The posterior statistics of the forecast moments are calculated using the corresponding smoothed posterior draws of $\mathbf{A}_{t|T}$, $\mu_{t|T}$ and $\Omega_{t|T}$ provided by our Gibbs sampler.

Figure 4 plots the medians and 18th and 86th percentile bands of the evolution of these forecast means at the 20-years-ahead horizon. A substantial part of the time variation is actually in the uncertainty surrounding the forecast means rather than in the median, which does not move too much for long periods of time for the observables we consider.

[10]We will later use P_t matrices that are consistent with our calculation of impulse responses, which we describe in the section on impulse responses.

The period from 1920 to 1940 (which encompasses the Great Depression) is represented in Figure 4 as a time of substantial uncertainty surrounding long-run values, but it is (maybe surprisingly) not associated with substantial movement in the median of the forecasts. Our model thus attributes a substantial part of the Great Depression to temporary changes in volatilities.[11]

The 1970s instead are viewed by our model as a time in which the long-run outlook was quite bleak in terms of GDP growth and inflation.

The Volcker disinflation around 1980 is seen as a major structural break in our model. Average forecasted inflation dropped dramatically, average forecasted output growth increased by 1 percent in annual terms, and the uncertainty surrounding these long-run-forecasts shrank. The recent financial crisis does not dramatically manifest itself in these long-run averages.

We use Var_{t+h} to construct time t approximations to the forecast correlations between our observables, medians and 68 percent error bands of which are depicted in Figures 5 (correlation of 20-year forecasts) and 6 (correlation of one-year forecasts). We will focus most on the long-run forecasts since they are not influenced by transitory movements and the correlations at the two horizons closely mirror each other.

There is substantial time variation in these correlations. The error bands are in general large. A substantial number of these correlations feature substantial movement in the 1970s and then a structural break at the time of the Volcker disinflation. Starting with the output growth/inflation correlation, we see that the median correlation becomes

[11]The estimates are based on all sample information. Out-of-sample forecasts using only information up to that time period would presumably look quite different.

substantially negative throughout the 1970s, implying that at high and persistent levels of inflation, the long-run levels of inflation and output growth move in opposite directions. After 1980, this strong negative relationship disappears and the median correlation becomes ever so slightly positive. A similar pattern can be observed for the output growth/interest rate relationship. The 68 percent error bands for the output growth/spread correlation contain 0 for most of the sample except for a period from the 1960s to the mid-1980s. The mid-1980s have been identified before as a point in time after which yield curve information does not carry much information for forecasting output growth.[12] Inflation and interest rates have been virtually uncorrelated for the first part of our sample (the "Gibson Paradox" studied by Cogley, Surico & Sargent (2012)). The correlation then grew throughout the 1960s and was close to 1 during the 1970s. Revisiting the by now common theme, the correlation falls dramatically with the disinflation of the early 1980s. To see why a 0 correlation between inflation and the nominal interest rate is surprising, remember the Fisher equation in its approximate linear form:

$$i_t^s = r_t + E_t \pi_{t+1} \tag{12}$$

If we think about the real interest rate r_t being roughly constant in the distant future then this equation tells us that in the long run short-term interest rates and inflation should move one-for-one.[13] A possible expla-

[12]Wheelock & Wohar (2009) state that "Several studies find that the spread has been less useful for forecasting output growth since the mid-1980s, at least for the United States."

[13]In our model we can not subtract our inflation measure from our measure of the short-term nominal interest rate to get a measure of the ex-post real rate because our short-term interest rate is a three month (annualized) interest rate, whereas we use an annual inflation measure. In terms of long-run forecasts, the difference between an annual interest rate and an annualized 3 month interest rate for a safe asset like we consider should be small. Also, we plot the correlation between inflation and the

nation for the disappearance of a significant correlation could be that during periods of low correlation between inflation and nominal interest rates inflation expectations are 'well-anchored' in that they do not move much in response to movements in variables at the time when the forecast is made. Inspecting our long-run forecast of inflation, we do indeed see little movement in forecasted inflation during times of low correlation between forecasted inflation and forecasted short-term interest rates.

Inflation and money growth are positively correlated (but not significantly so) before the mid-1960s, when the correlation becomes much stronger. The strength of this correlation disappears immediately with the beginning of the Volcker chairmanship and the associated disinflation. This again points to a positive relationship at high levels of inflation, but not at the substantially lower levels we have observed since the 1980s.

We see a substantially negative relationship between the short-term interest rate and the spread before 1980. This correlation has since become much closer to 0, meaning forecasted long-run movements in the short rate do not feed (linearly) into the slope of the yield curve. This might have implications for monetary policy - policymakers hope to influence long-term interest rates by moving their short-term target interest rates around. One might argue that our findings should not be of much concern to policymakers because of the long forecast horizon, but in Figure 6 we plot the corresponding correlation structure for a one-year forecast and find a very similar movement. These correlations, however, are not conditional on specific shocks hitting the economy. Thus, we cannot say with any certainty that our results imply less influence of policymakers

nominal interest rate 20 years in the future, whereas the Fisher equation would call for the correlation between the nominal interest rate in 20 years and the inflation rate in 20 years and one quarter. Given our long forecast horizon this seems inconsequential.

on the yield curve. Nonetheless, we do think this finding is worth noticing.

Finally, the correlation between money growth and the spread has moved into positive territory after 1980, both for one-year and 20-year forecasts. Only at the very end of the sample do these correlations move toward 0 again. Taken at face value, this implies that from 1980 to the early 2000s money growth could have been useful in predicting movements of the yield curve.

5.7 Impulse Responses to a Monetary Shock Over the Last Century

We want to analyze how the impact of unexpected movements in monetary policy has evolved over time and how much those unexpected movements have contributed to overall volatility in the economy. Defining a monetary policy shock for post-WWII data is straightforward: Economists tend to think of the Federal Reserve after WWII as choosing a path for the short-term interest rate. If we have a model (or equation) for the short-term nominal interest rate, we can then define the monetary policy shock as the residual after properly accounting for movements in all variables deemed relevant for the setting of the short-term interest rate. The same would hold true if the Federal Reserve consistently used changes in money growth as its policy instrument. In our sample we are faced with the difficulty that there has been no consistent conduct of monetary policy. We thus aim to identify monetary policy shocks not as identified shocks associated with a certain equation or variable, but rather by their impact on the economy - we use sign restrictions to identify a monetary policy shock in our model.

This section will first describe the impact of a one standard-deviation monetary policy shock on the economy. Then we will ask how important fluctuations in the size of that shock are for our results relative to changes in the dynamic response to shocks. Finally, we discuss what fraction of the overall volatility of the economy is explained by our identified monetary policy shock.

The following assumption summarizes our sign restrictions:

Assumption 1: A *monetary policy impulse vector* at time t is an impulse vector a_t, so that the impulse responses to a_t of inflation, output growth and money growth are not positive and the impulse responses for the short term interest rate are not negative, all at horizons $k = 0, \ldots, K$.

Our benchmark specification does not restrict the impulse responses of the spread. The sign restrictions are imposed for $K = 2$, hence we impose the sign restrictions at each point in time for the specified contemporaneous responses and the first and second quarter. This is in line with Uhlig (2005), who uses five months in a monthly model and Benati (2010), who imposes the restrictions on impact and for the two following quarters.[14] In contrast to the benchmark case in Uhlig (2005), we do restrict the response of output growth not to react positively following a contractionary monetary policy shock. Most theoretical macroeconomic models feature meaningful output responses to monetary policy shocks, a feature that we use to guide our identification restrictions (see Canova & Paustian (2011) for an introduction to this approach). The candidate time t im-

[14] In the online appendix we explore alternative horizons for which restrictions are imposed and find that our main conclusions hold in those cases as well.

pulse vector a_t is given by

$$a_t = \Lambda_{t|T}^{-1} \Sigma_{t|T} \alpha_t \tag{13}$$

where α_t is a column vector of conformable size drawn from the unit sphere of norm 1. To compute the impulse response of variable j to shock a_t at horizon k let $\mathbf{a}_t = [a_t', \mathbf{0}]'$[15] and calculate

$$r_{t,k}^{j,a} = \left(\mathbf{A}_{t|T}^{k} \mathbf{a}_t\right)_j. \tag{14}$$

This approach builds on Uhlig (2005), Faust (1998), Canova & Nicolo (2002), and Canova & Gambetti (2009). Additional details regarding implementation and normalization are provided in appendix B.

5.7.1 A Historical Assessment of Dynamic Consequences of Monetary Policy Shocks

We will focus our attention on a subset of available impulse responses. Figures 8 to 11 show the median responses to a one standard deviation shock[16]. We see that the response of short-term interest rates varies substantially across the sample, from small impact responses in the 1930s and 1940s to very large impact responses during the early 1980s. This might tempt some readers to assume that monetary shocks were small during the pre-WWII period, but remember that we do not associate a monetary policy shock with an innovation in interest rates alone. Figure 9 shows the impulse response of the growth rate of money. We see that

[15] $\mathbf{0}$ is a conformable vector of zeros.

[16] Following Canova & Gambetti (2009), we will focus on median responses throughout. Even though we identify impulse responses using sign restrictions only, the error bands are estimated tightly enough to make the median responses meaningful. The online appendix shows error bands for impulse responses at different horizons.

this impulse response is substantial during the exact time period when the impact response of the nominal interest rate is small.[17] The overall effects on inflation and output growth of a monetary policy shock during that time are in fact larger than at any point after that. The responses of inflation and output growth are substantially more stable after WWII. In contrast to our findings based on the estimated reduced-form VAR, we do not see a large structural break around 1980 in the responses of inflation and output growth. We will next come back to this question and ask if there might have been structural breaks after all when looking at the impulse responses through a different lens.

5.7.2 Characterizing the Evolution of the Monetary Transmission Mechanism

We have previously documented both changes in contemporaneous volatilities of forecast errors associated with our VAR as well as changes in parameters governing the dynamic responses in our model. We now try to disentangle the effects of changes in contemporaneous volatility and changes in response parameters. To do so, we want to normalize the contemporaneous effect of the monetary policy shock throughout the sample. Canova & Gambetti (2009) normalize their sign-restriction-based impulse responses by fixing the contemporaneous effect on the nominal interest rate. This strategy is useful when focusing on post-WWII data, but it seems less useful in our context. To see this, reconsider the impulse response function to the nominal interest rate and money growth

[17] The response of money growth is largest at the very end of the sample, owing to the large increase of observed money growth at the end of the sample, which our model partially rationalizes as an increase in the volatility of the forecast error associated with money growth. At that point in time, the dynamic response of the other variables in our model to changes in money growth is substantially different from the pre-WWII era.

shown before. Fixing the contemporaneous impact on the nominal interest rate to the same value throughout the sample (say an average of all contemporaneous responses throughout the sample) would imply scaling a contemporaneous response in a specific period by the ratio of the average response to the contemporaneous interest rate response in that period. For the period in the 1930s discussed above this would imply a substantial increase in all contemporaneous responses (since the contemporaneous response of the nominal interest rate is small during that time), leading to effects that most economists would deem unreasonable even for that time period. This is an artifact of implicitly choosing the nominal interest rate as the policy instrument when using that normalization. We instead choose an alternative normalization. One can think of our reduced-form model as already encoding a recursive identification scheme as a benchmark[18]: The matrix $\Lambda_t^{-1}\Sigma_t$ is lower triangular. So, if we were interested in such a decomposition, it would be natural to normalize the contemporaneous impact so that all structural shocks have a unit impact. This amounts to using Λ_t^{-1} as impact matrix. We choose the same normalization for our identification scheme. Thus any time variation in the contemporaneous impact comes from deviations of our identification scheme from the recursive ordering. Figures 12 to 14 show selected impulse responses that have been normalized this way. We see that most of the time variation in the response of inflation and output growth that we described before comes from changes in the volatility of forecast errors - changes in volatility that someone using a recursive identification scheme would assign to the structural shocks. Most of the changes throughout time of the impulse responses are driven by

[18]This has been exploited by Primiceri (2005).

changes in the volatilities of shocks. Notwithstanding, our normalized impulse reposes do uncover some meaningful changes in the transmission of monetary policy shocks. The long-run (15-20 quarters after the shock) response of output growth has increased throughout time. The impulse responses of inflation now show a clear structural break around 1980 that was masked before by the changes in volatility. We see a temporary spike in long-run responses around 1980 (those were already visible in the one standard deviation shock case) but also a substantial decrease of that response afterward - monetary policy shocks have a bigger long-run impact on inflation after 1980.

5.7.3 The Importance of Monetary Shocks Over Time

So far we have studied the impact of a monetary policy shock over time. A different question is how much monetary policy shocks have contributed to overall fluctuations in the economy over time. We tackle this question by computing the variance decomposition implied by our model and our identification scheme for the monetary policy shock. We follow the same approach we used when calculating the impulse response functions: We use the posterior distribution for parameters at each point in time and ask what the variance decomposition would be if parameters did not change in the future. This gives a sense of the local dynamics at each time t. Figures 15 to 17 display those variance decompositions (these figures show the median values across draws). For output growth, we can see a low frequency trend - at the beginning of the sample the overall importance of the identified monetary policy shock is relatively low (slightly higher than 10 percent on average) but the average value increases after WWII and then falls again after 1980. For inflation, we see a spike in the variance decompositions at short horizons around 1980. The variance decompositions also show time variation in the shape (across horizons) that is absent for output growth. The variance decompositions are relatively flat across horizons for the beginning and very end of the sample, but throughout the 1970s and 1980s we see substantially more time variation across horizons.

The variance decomposition of the nominal interest rate shows a large spike at short horizons around 1945, a decrease after WWII and an increase again after 1980 (except for the episode around 1945 when the variance decompositions are flat across horizons). How can we reconcile the spike around 1945 with the small impulse response of the nominal interest rate to our identified monetary policy shock? To do so, it is helpful to go back to Figure 7, which plots the volatility of the one-step ahead forecast errors. We see that the lowest level for the interest rate forecast volatility is reached around the same time as the spike in the variance decomposition. Thus, while the shock has little impact on the interest rate (as can be seen in the one standard deviation impulse response function), it still accounts for a large fraction of the (small) forecast error variance at short horizons of the interest rate.

6 Conclusion

In this paper we tried to uncover evidence on the amount of time variation in the U.S. economy, both when it comes to reduced-form statistics and when it comes to the response of the economy to structural (monetary policy) shocks. We found substantial evidence of time variation in volatilities of reduced form innovations and responses of the economy to those innovations. In particular, the early 1980s were a time period that our model associates with substantial shifts in the structure of the economy.

The impact of an average monetary policy shock varies substantially across time according to our model, but a large part of that time variation is driven by changes in the magnitude of this shock, not necessarily how this shock is transmitted through time.[19]

[19] This finding is in line with previous findings in the literature: Primiceri (2005) finds that there is not much time variation in impulse responses to shocks of a given size in post-WWII data using a recursive identification scheme. Sims & Zha (2006) argue that most of the time variation in post-WWII US time series is driven by changes in the volatility of innovations.

References

Amir-Ahmadi, P. (2009), Credit shocks, monetary policy, and business cycles: Evidence from a structural time varying Bayesian FAVAR, Working paper, Goethe University Frankfurt.

Amir-Ahmadi, P. & Ritschl, A. (2013), 'Depression econometrics: A favar model of monetary policy during the great depression', *mimeo* .

Balke, N. & Gordon, R. J. (1986), Appendix B Historical Data, *in* 'The American Business Cycle: Continuity and Change', NBER Chapters, National Bureau of Economic Research, Inc, pp. 781–850.

Benati, L. (2010), Evolving Phillips trade-off, Working Paper Series 1176, European Central Bank.

Benati, L. & Lubik, T. (2012), Sales, inventories, and real interest rates: A century of stylized facts, Working paper.

Bianchi, F. & Ilut, C. (2013), Monetary/fiscal policy mix and agents' beliefs, Technical report.

Canova, F. & Gambetti, L. (2009), 'Structural changes in the US economy: Is there a role for monetary policy?', *Journal of Economic Dynamics and Control* **33**(2), 477–490.

Canova, F. & Nicolo, G. D. (2002), 'Monetary disturbances matter for business fluctuations in the G-7', *Journal of Monetary Economics* **49**(6), 1131–1159.

Canova, F. & Paustian, M. (2011), 'Business cycle measurement with some theory', *Journal of Monetary Economics* **58**(4), 345–361.

Carter, C. K. & Kohn, R. (1994), 'On gibbs sampling for state space models', *Biometrika* **81**(3), 541–553.

Christiano, L. J., Eichenbaum, M. & Evans, C. L. (1999), Monetary policy shocks: What have we learned and to what end?, *in* J. B. Taylor & M. Woodford, eds, 'Handbook of Macroeconomics', Vol. 1 of *Handbook of Macroeconomics*, Elsevier, chapter 2, pp. 65–148.

Cogley, T., Matthes, C. & Sbordone, A. (2012), Optimal disinflation under learning, Working paper, NYU, UPF and NY FED.

Cogley, T. & Sargent, T. J. (2005), 'Drift and volatilities: Monetary policies and outcomes in the post WWII U.S.', *Review of Economic Dynamics* **8**(2), 262–302.

Cogley, T. & Sargent, T. J. (2014), Measuring price-level uncertainty and instability in the U.S., 1850-2012, Technical report.

Cogley, T., Surico, P. & Sargent, T. J. (2012), The return of the Gibson paradox, Working paper, NYU and LBS.

D'Agostino, A. & Surico, P. (2011), A century of inflation forecasts, CEPR discussion papers, C.E.P.R. Discussion Papers.

Del Negro, M. & Primiceri, G. (2013), Time-varying structural vector autoregressions and monetary policy: a corrigendum, Staff Reports 619, Federal Reserve Bank of New York.

Faust, J. (1998), 'The robustness of identified VAR conclusions about money', *Carnegie-Rochester Conference Series in Public Policy* **49**, 207–244.

Gali, J. & Gambetti, L. (2009), 'On the sources of the great moderation', *American Economic Journal: Macroeconomics* **1**(1), 26–57.

Kim, S., Shephard, N. & Chib, S. (1998), 'Stochastic volatility: Likelihood inference and comparison with ARCH models', *Review of Economic Studies* **65**(3), 361–93.

Kliem, M., Kriwoluzky, A. & Sarferaz, S. (2013), On the low-frequency relationship between public deficits and inflation, Technical report.

Koop, G. & Korobilis, D. (2010), Bayesian multivariate time series methods for empirical macroeconomics, Working paper, University of Strathclyde.

Lütkepohl (2009), *New Introduction to Multiple Time Series Analysis*, Springer Verlag.

Nason, J. & Tallman, E. (2013), 'Business cycles and financial crises: The roles of credit supply and demand shocks', *Macroeconomic Dynamics* .

Primiceri, G. (2005), 'Time varying structural vector autoregressions and monetary policy', *Review of Economic Studies* **72**(3), 821–852.

Ritschl, A. & Woitek, U. (2000), 'Did monetary forces cause the great depression? a bayesian var analysis for the u.s. economy', *CEPR Discussion Paper 2547, Centre for Economic Policy Research* .

Romer, C. D. (1986), 'The instability of the prewar economy reconsidered: A critical examination of historical macroeconomic data', *The Journal of Economic History* **46**(02), 494–496.

Sargent, T. J. & Surico, P. (2011), 'Two illustrations of the quantity theory of money: Breakdowns and revivals', *American Economic Review* **101**(1), 109–28.

Sims, C. (1999), 'The role of interest rate policy in the generation and propagation of business cycles: What has changed since the 30's', *Proceedings 1998 Boston FRB Annual Res. Conf., Federal Reserve Bank of Boston* pp. 121–160.

Sims, C. A. & Zha, T. (2006), 'Were there regime switches in U.S. monetary policy?', *American Economic Review* **96**(1), 54–81.

Uhlig, H. (2005), 'What are the effects of monetary policy on output? Results from an agnostic identification procedure', *Journal of Monetary Economics* **52**(2), 381–419.

Wheelock, D. C. & Wohar, M. E. (2009), 'Can the term spread predict output growth and recessions? a survey of the literature', *Federal Reserve Bank of St. Louis Review* .

A Estimation Algorithm

We use a Gibbs-Sampler to approximate the posterior distribution by generating $500,000$ draws. The exact implementation follows Primiceri (2005) including the corrigendum of Del Negro and Primiceri (2013). The algorithm proceeds as follows[20]:

1. Draw Σ^T from $p(\Sigma^T|y^T, \theta^T, \Lambda^T, V)$. This step requires us to generate draws from a nonlinear state space system. We use the approach by Kim, Shephard & Chib (1998) to approximate draws from the desired distribution. For a correct posterior sampling of the stochastic volatilities we follow the corrigendum in Del Negro and Primiceri (2013) and the modified steps therein (in particular, we first need to generate a new draw of the indicator variables used in the Kim et al. (1998) approach).

2. Draw θ^T from $p(\theta^T|y^T, \Lambda^T, \Sigma^T, V)$. Conditional on all other parameter blocks equations (3) and (4) form a linear Gaussian state space system. This step can be carried out using the simulation smoother detailed in Carter & Kohn (1994).

3. Draw Λ^T from $p(\Lambda^T|y^T, \theta^T, \Sigma^T, V)$. Again we draw these covariance states based on the simulation smoother of the previous step, exploiting our assumption that the covariance matrix of the innovations in the law of motion for the λ coefficients is block diagonal. This assumption follows Primiceri (2005), where further details on this step can be found.

4. Draw V from $p(V|\Sigma^T y^T, \theta^T, \Lambda^T)$. Given our distributional assump-

[20] A superscript T denotes a sample of the relevant variable from $t = 1$ to T.

tions, this conditional posterior of the time-invariant variances follows an inverse-Wishart distribution, which we can easily sample from.

B Algorithm to Draw Impulse Responses

Here we describe the procedure for the identification of the evolving impulse response functions to contractionary monetary policy shocks via pure sign restrictions briefly outlined in the main text. At each iteration g, for a given time period t, given a set of posterior draws of all parameters Θ_{IRF}^{g} from the stationary phase of the target distribution, we employ the following procedure c_{α} number of times to find an identified contractionary monetary policy shock:

1. Take five (which is the number of our observables) independent draws $\alpha_{M\times 1} \sim N(0,1)$ and normalize them so that the vector consisting of all those draws has unit length. Calculate the time t impulse vector at iteration g for each candidate:

$$a_t = \Lambda_{t|T}^{-1}\Sigma_{t|T}\alpha_t \tag{15}$$

2. Calculate the time t impulse responses to $\mathbf{a}_t$ (which contains a_t and a conformable number of zeros) and store the impulse responses if the sign restrictions are met.

$$r_{t,k}^{g,i} = \mathbf{A}_{t|T}^{k}\mathbf{a}_t \tag{16}$$

   where t is the time index, g the Gibbs iteration index, k the horizon of the impulse response functions and i the candidate index. Otherwise discard.

3. Redo the above procedure at each iteration g for each time period $t = 1, \ldots, T$.

4. Calculate the statistics of interest.

In our application we keep a random selection of 1000 posterior draws of all parameters and latent states from the ergodic distribution. For the calculation of the SR-IRF we evaluate a sample length 366 quarters, setting the number of candidates to $c_\alpha = 50$. At each quarter we calculate 50.000 candidates and overall we calculate $18.300.000$ candidate IRFs, which is quite time consuming.

C Figures

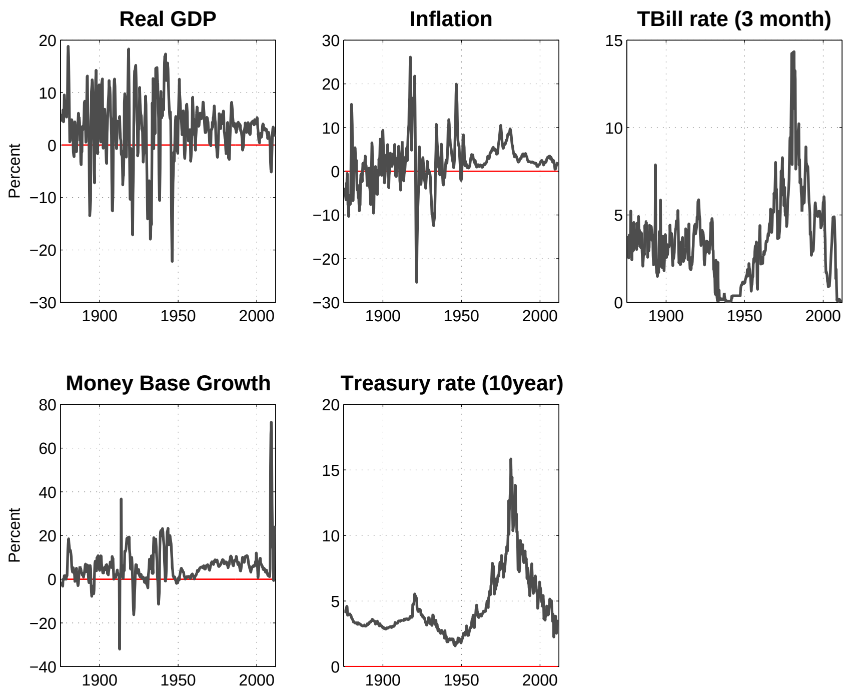

Figure 1: Data

This figure displays the transformed data as it enters our baseline VAR specification with annual real GDP growth, annual inflation rate, short-term interest rate, the term spread, and annual M2 growth. All series are in percentage units.

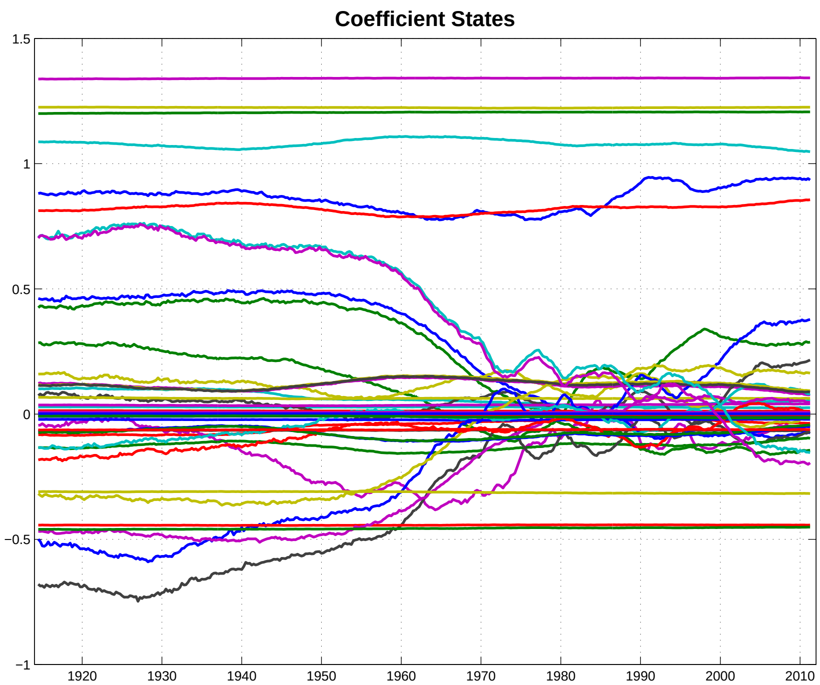

Figure 2: Time variation in coefficients

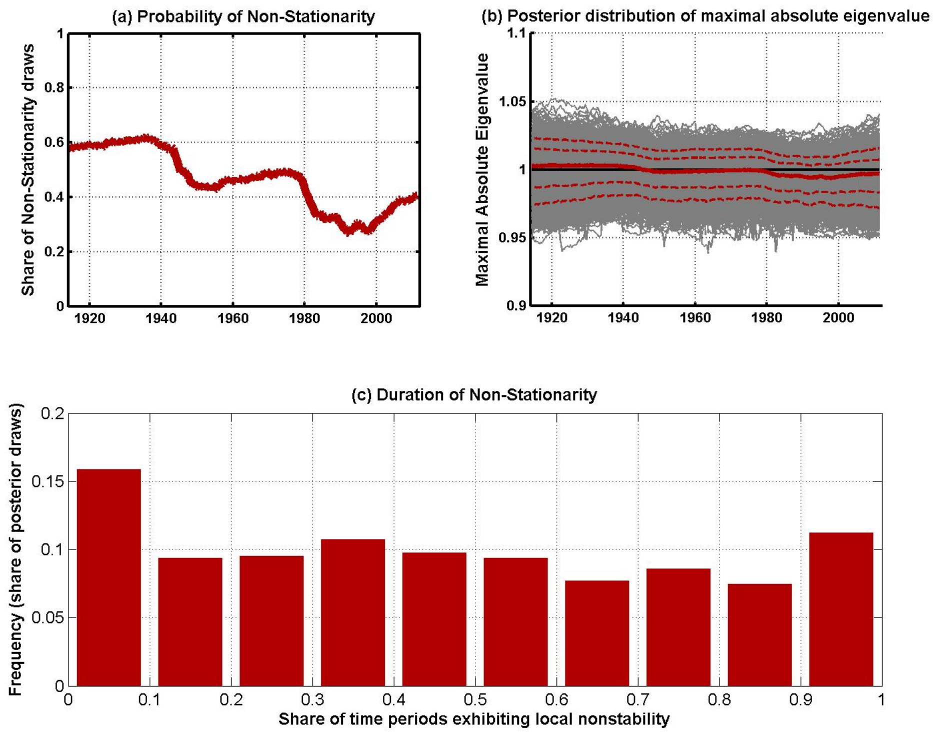

Figure 3: Explosive behavior

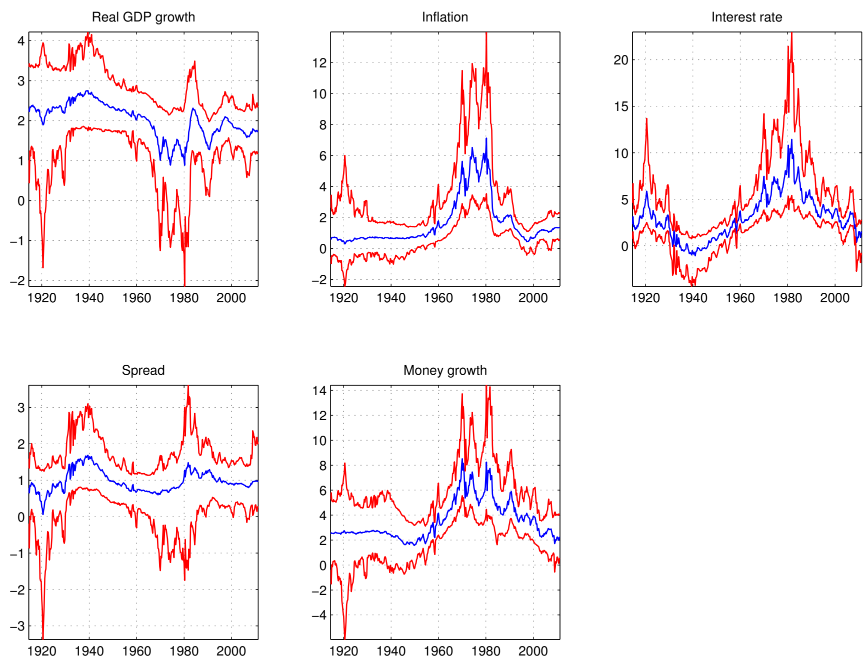

Figure 4: Evolving forecast means: 20 years ahead

This figure shows in blue the posterior median estimates of the time varying forecast means of our baseline VAR model. The red lines are the posterior 68 percent error bands.

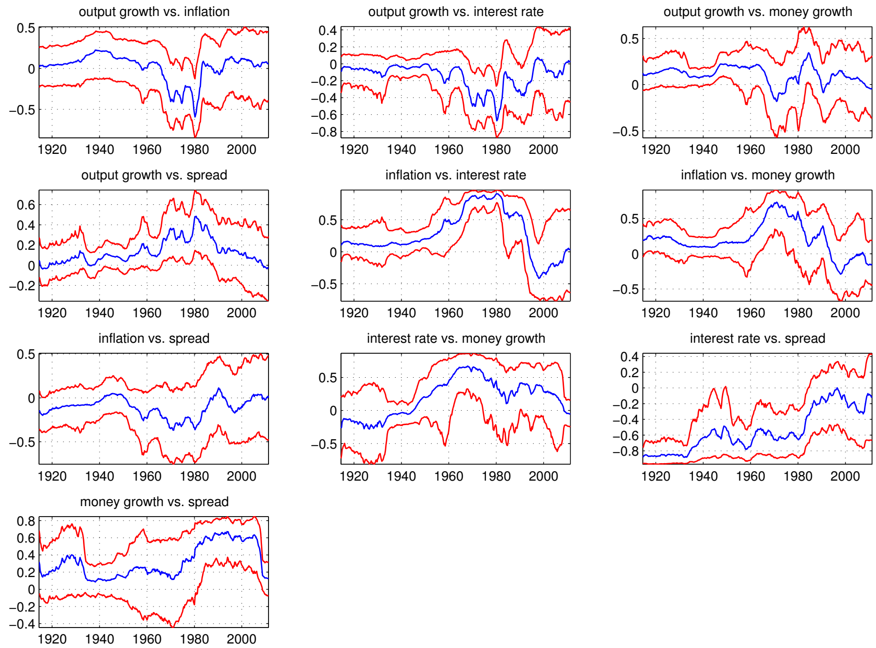

Figure 5: Forecast correlations: 20 years ahead

This figure shows in blue the posterior median estimates of the time varying forecast correlations of our baseline VAR model. The red lines are the posterior 68 percent error bands.

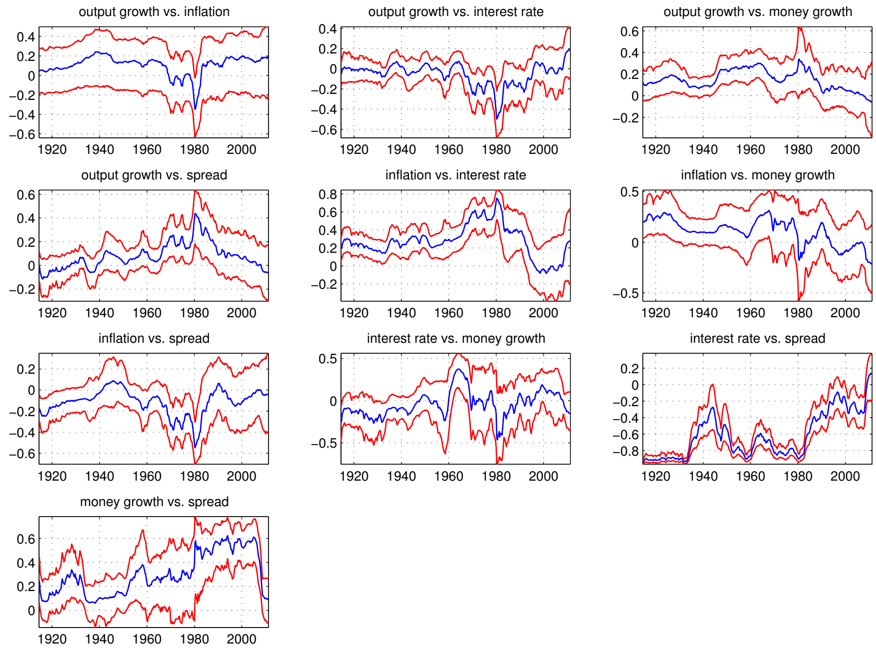

Figure 6: Forecast correlations: 1 year ahead

This figure shows in blue the posterior median estimates of the time varying forecast correlations of our baseline VAR model. The red lines are the posterior 68% error bands.

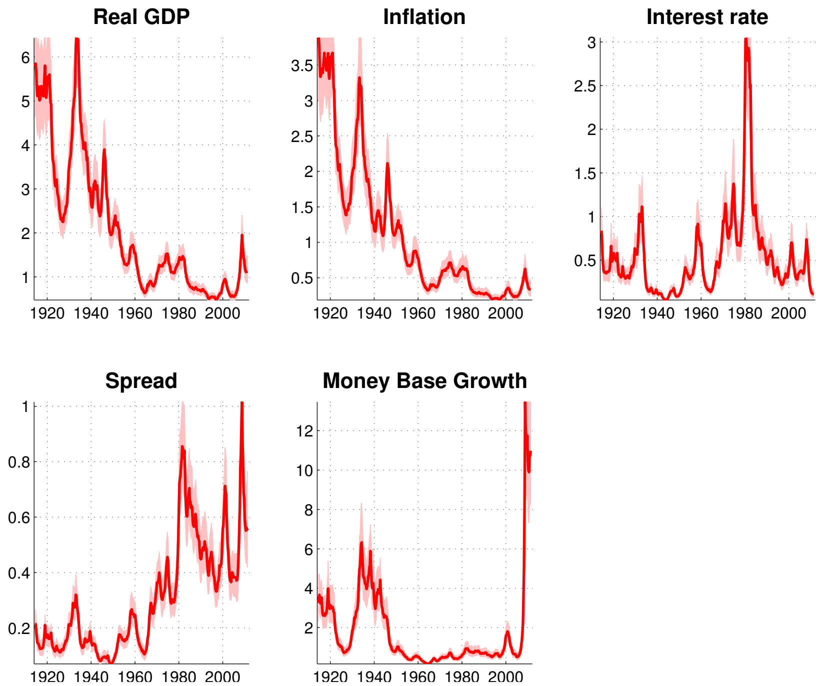

Figure 7: Volatility of reduced form residuals

C.1 Impulse Response Functions identified via sign restrictions

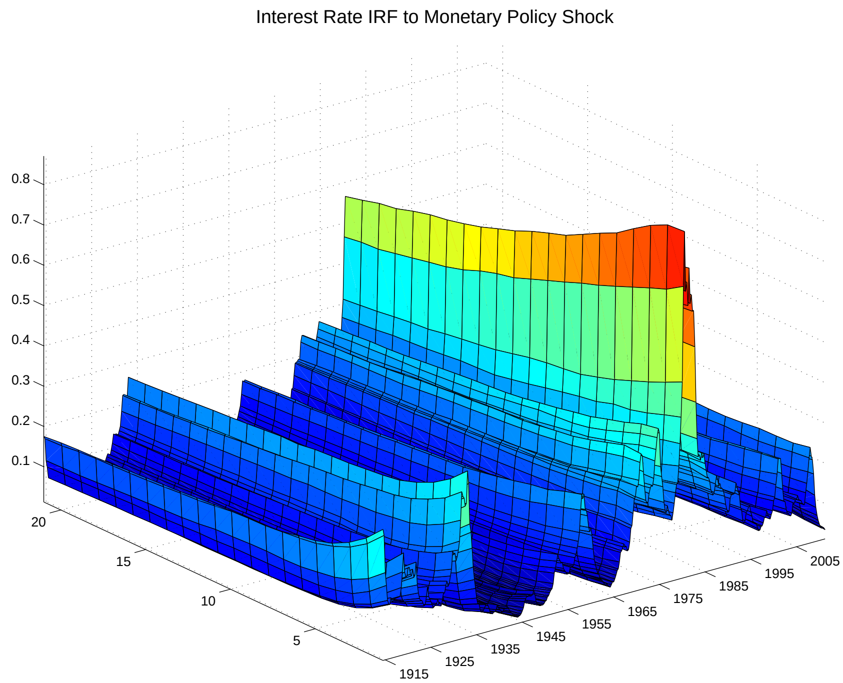

Figure 8: Evolving IRF of short-term interest rate to a monetary policy shock.

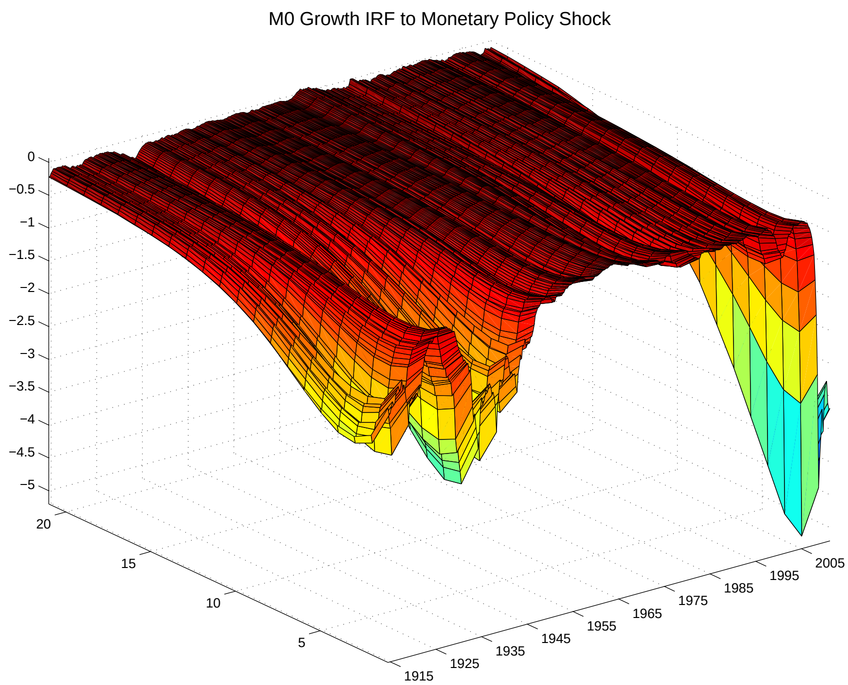

Figure 9: Evolving IRF of money growth to a monetary policy shock.

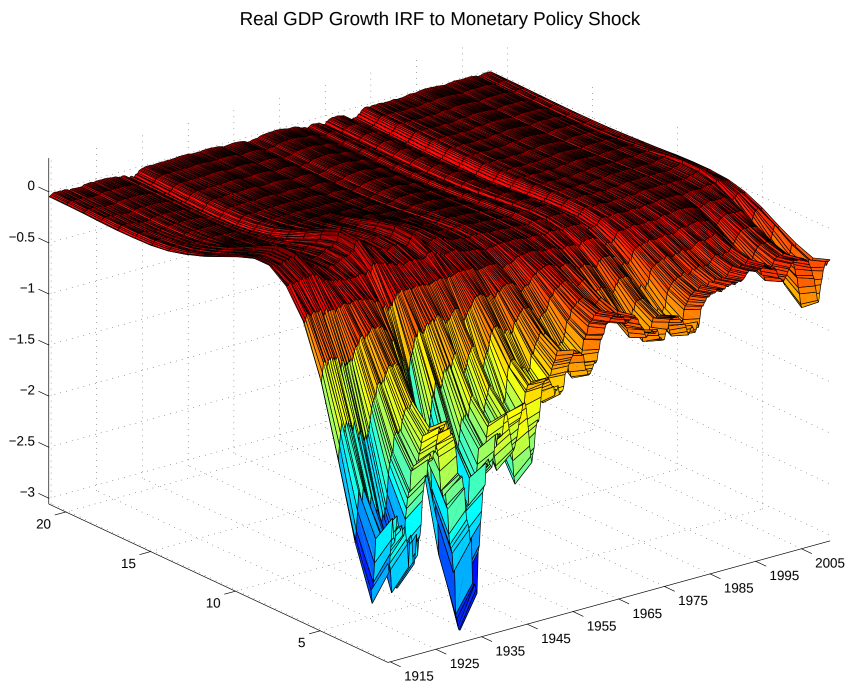

Figure 10: Evolving IRF of real GDP growth to a monetary policy shock.

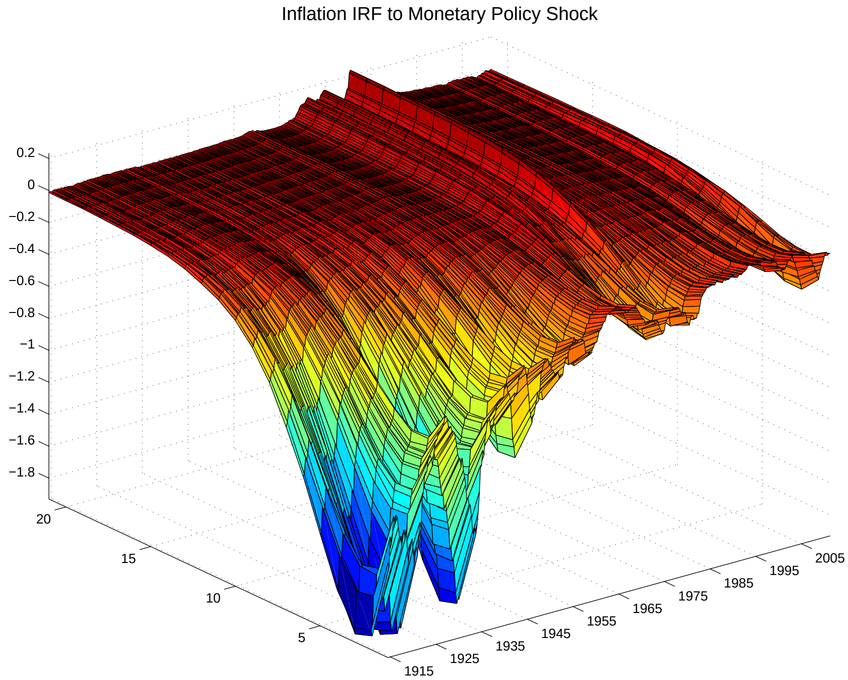

Figure 11: Evolving IRF of inflation to a monetary policy shock.

C.2 Impulse Response Functions identified via sign restrictions (normalized)

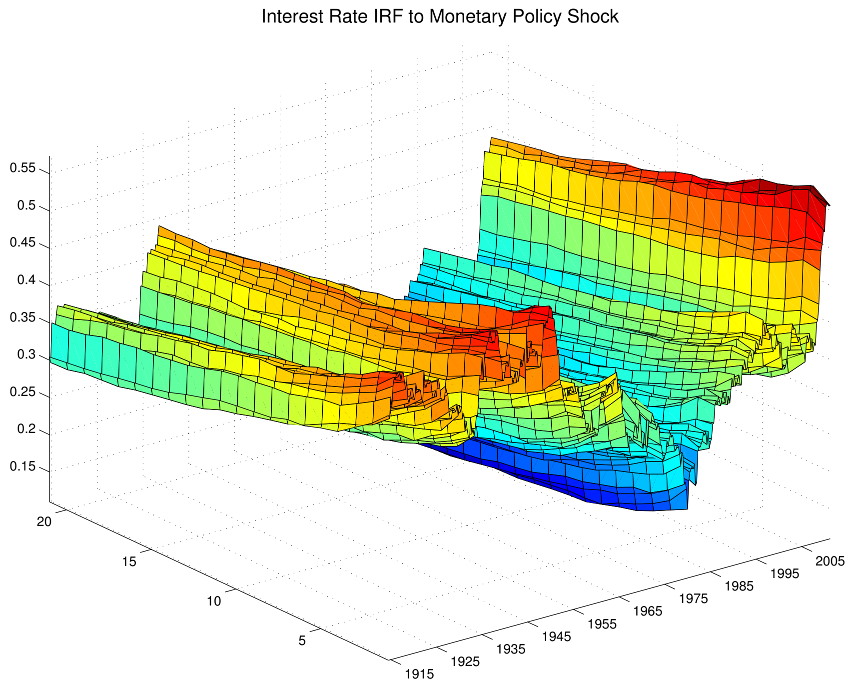

Figure 12: Evolving IRF of short term interest rate to a monetary policy shock.

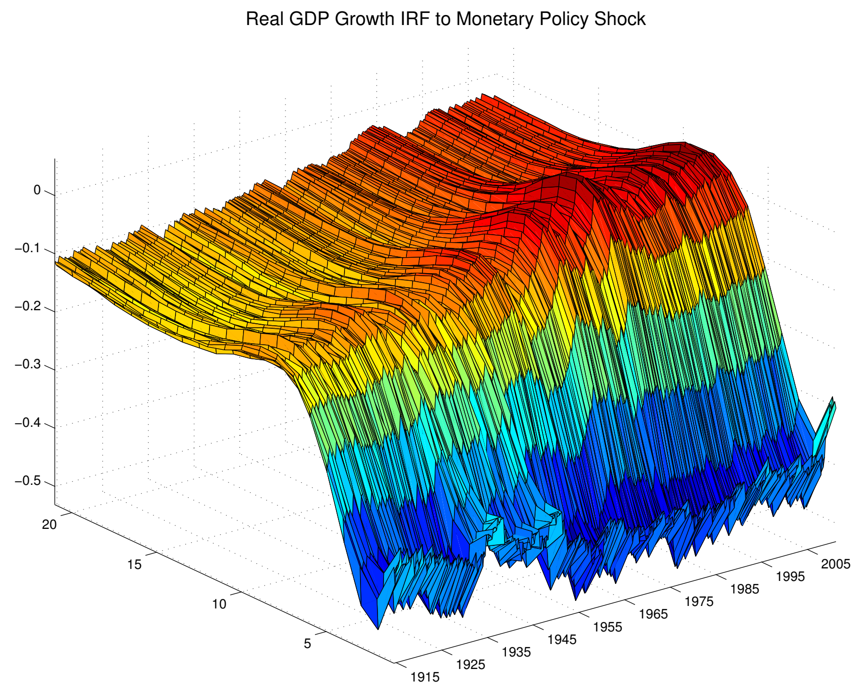

Figure 13: Evolving IRF of real GDP growth to a monetary policy shock.

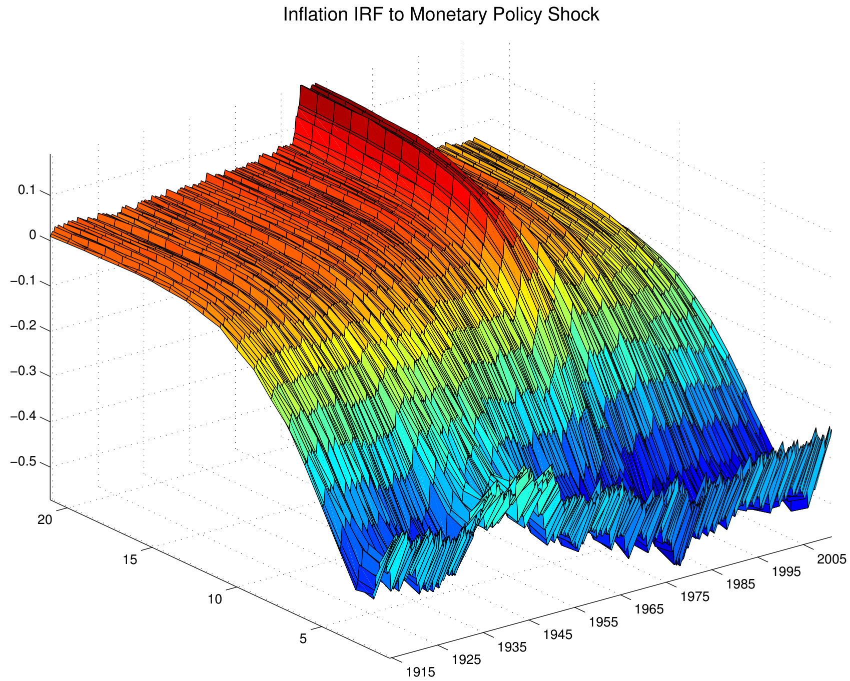

Figure 14: Evolving IRF of inflation to a monetary policy shock.

C.3 Variance Decompositions

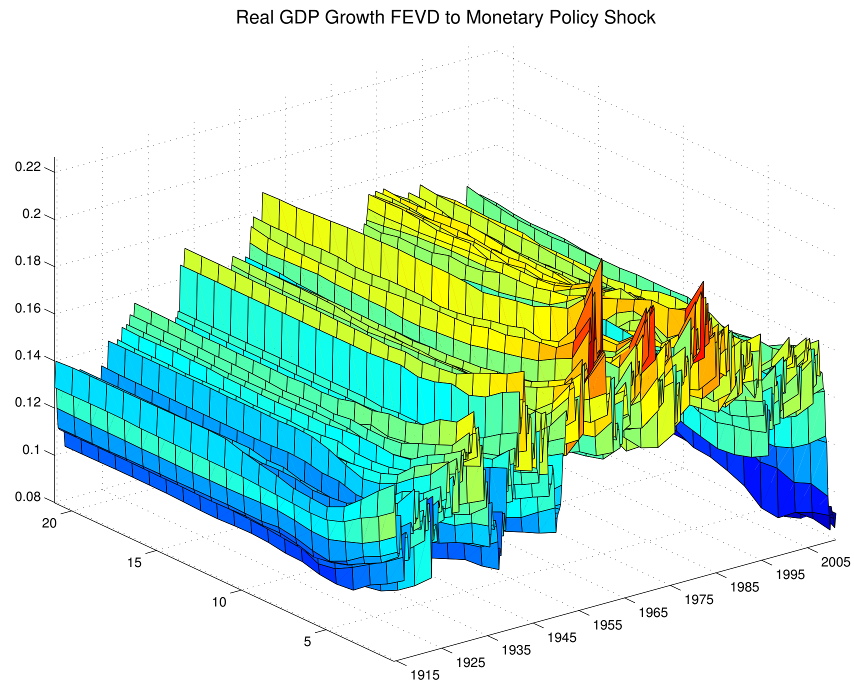

Figure 15: Evolving forecast error variance decomposition of output growth.

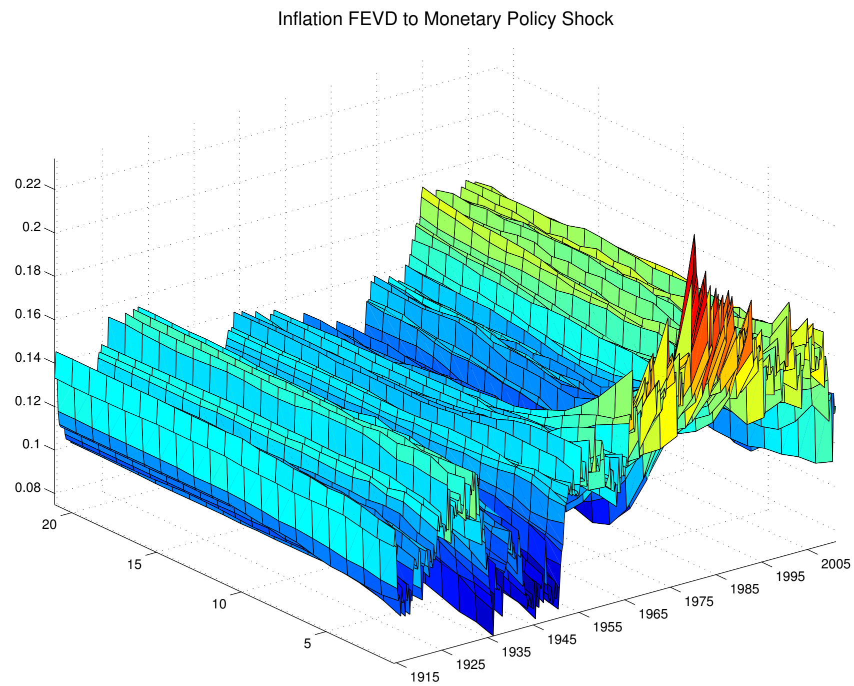

Figure 16: Evolving forecast error variance decomposition of inflation.

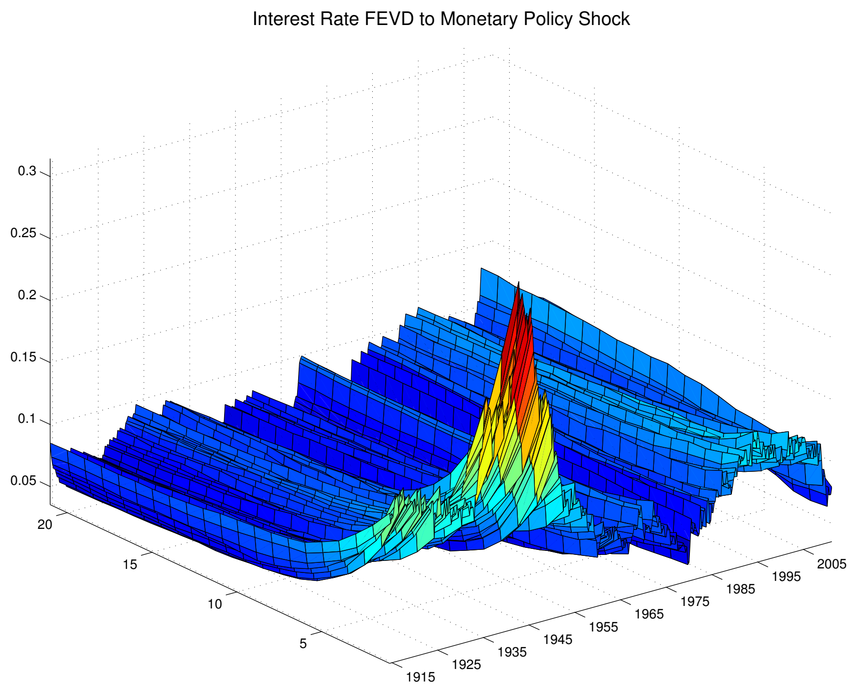

Figure 17: Evolving forecast error variance decomposition of the short-term nominal interest rate.