Measurement Errors and Monetary Policy: Then and Now

Pooyan Amir-Ahmadi*
University of Illinois at Urbana-Champaign
Christian Matthes
Federal Reserve Bank of Richmond
Mu-Chun Wang
University of Hamburg

March 2, 2017

Abstract

Should policymakers and applied macroeconomists worry about the difference between real-time and final data? We tackle this question by using a Bayesian VAR with time-varying parameters and stochastic volatility to show that the distinction between real-time data and final data matters for the impact of monetary policy shocks: The impact on final data is substantially and systematically different (in particular, larger in magnitude for different measures of real activity) from the impact on real-time data. These differences have persisted over the last 40 years and should be taken into account when conducting or studying monetary policy.

Keywords: real-time data, time-varying parameters, stochastic volatility, impulse responses

*We would like to thank Fabio Canova, Tim Cogley, and Dean Croushore as well as participants at the 2015 IAAE conference and the Cirano workshop on real-time data in Montreal for their comments. Miki Doan, Marisa Reed and Daniel Tracht provided excellent research assistance. 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.

1 Introduction

When monetary policymakers evaluate the effects of their most recent policy decisions (say, to prepare for the next round of monetary policy decisions) they only have access to preliminary real-time estimates of macroeconomic data that have been collected after the policy decision they want to evaluate. Is the difference between real-time and final macroeconomic data important enough to be considered when analyzing and conducting monetary policy? We revisit this question asked by Croushore & Evans (2006) in light of recent evidence (Aruoba (2008)) that the measurement errors in macroeconomic data are far from satisfying the properties of classical measurement errors and evidence that there is substantial time variation in the dynamics of U.S. macroeconomic time series, as emphasized by Cogley & Sargent (2005) and Primiceri (2005).

We use a Bayesian vector autoregression (VAR) with time-varying parameters and stochastic volatility estimated on data that includes real-time and final releases of macroeconomic data to uncover substantial time variation in the dynamics of measurement errors: We find that the measurement errors are significantly correlated for some variables, feature substantial changes in volatility and can be different from zero for long periods of time with magnitudes that are economically meaningful. We use a model with time-varying parameters and stochastic volatility because time variation in the dynamics and volatility of final data has been identified as important for (final) U.S. data by Cogley & Sargent (2005), Primiceri (2005), and Canova & Gambetti (2009), among others. Our paper shows that these features carry over to real-time data as well.

By using sign restrictions to identify a monetary policy shock, we establish that policymakers should indeed care about measurement errors. Differences between the impulse responses of real-time and final data on measures of real activity are significant and persist over time. As these differences are persistent over time, policymakers should take them into account.

Our work is related to the literature on time variation in macroeconomic dynamics such as Cogley & Sargent (2005), Primiceri (2005), and Gali & Gambetti (2009). The model we use to analyze time variation in our data is borrowed from those papers.[1] As pioneered by Canova & Nicolo (2002), Faust (1998) and Uhlig (2005), we use sign restrictions to identify monetary policy shocks. Canova & Gambetti (2009) use sign restrictions to identify monetary policy shocks in a VAR with time-varying parameters and stochastic volatility, but they do not consider real-time data.

Croushore & Evans (2006) tackle issues similar to ours, though in the context of a fixed coefficient VAR using either recursive or long-run restrictions. Their model of measurement error is less general than ours. For example, their models not only use fixed coefficient models, but also do not allow for biases (non-zero intercepts) in the relationship between different vintages of data. We find that these features matter, reinforcing the results by Aruoba (2008), and beyond that establish that measurement errors feature stochastic volatility and are correlated across variables.

In contrast to our work and Croushore & Evans (2006), the large majority of papers on real-time data focuses on statistical models of measurement error that do not identify effects of structural shocks. Jacobs & van Norden (2011) are motivated by the evidence in Aruoba (2008) and build a flexible model for a univariate measurement error series. In contrast to us, they model intermediate data releases, but do not consider the relationship of measurement errors across variables, time variation in the parameters, or stochastic volatility. Just as the model used in Jacobs & van Norden (2011), our model is general enough to allow for measurement errors that are correlated with either only final data ('news') or correlated only with the real-time data ('noise') as well as intermediate cases, as we show in the model section. Jacobs, Sarferaz, van Norden & Sturm (2013) build a multivariate version of Jacobs & van Norden (2011), but still abstract from stochastic volatility and time-varying parameters. Both Jacobs & van Norden (2011) and Jacobs et al. (2013) do not study the response of the economy to structural shocks, which is our main focus.

D'Agostino, Gambetti & Giannone (2013) use a VAR with time-varying parameters and stochastic volatility on real-time data to study the forecasting ability of models in this class.

[1] An overview of this literature is given in Koop & Korobilis (2010).

Fixed-coefficient VARs using various vintages of real-time data have previously been used to improve forecasting ability by Kishor & Koenig (2009) and Carriero, Clements & Galvao (2015), for example.

The issue of mismeasured data is also of utmost importance when studying long-run historical data. While scholars using historical data usually do not have access to revised data for the entire sample, they sometimes explore overlapping data sources - Cogley & Sargent (2014) do this in a model for US inflation that features stochastic volatility for true data, but in contrast to our approach their model does not feature stochastic volatility for the measurement error.

Croushore & Sill (2014) estimate a dynamic stochastic general equilibrium (DSGE) model on final data and then use the approach of Schorfheide, Sill & Kryshko (2010) to link real time data to the state variables of the estimated DSGE model. Similar to our findings, their findings show both that there are substantial differences between real-time and final data responses and that final data responses tend to be larger in absolute value.

In the next section we describe our model. We then turn to results for our benchmark specification.[2]

2 The Model

We jointly model the dynamics of the first release of any data point published - we call this real-time data - and the latest vintage available at the time of the writing of this paper - which we use as a proxy for final data. Throughout this paper, we study

[2] In an online appendix, we show that our findings are robust to alternative specifications: (i) using an alternative measure of real activity, employment growth, (ii) using an alternative identification scheme to identify monetary policy shocks, (iii) using an alternative definition of final data (iv) imposing sign restrictions on both real-time and final data, and (v) using the Wu & Xia (2016) shadow rate instead of the Federal Funds rate. The online appendix also discusses how our VAR can be motivated by a DSGE model with asymmetric information.

the dynamics of vectors of the following form:

$$y_t = \begin{pmatrix} \pi_t^{real} \\ \pi_t^{final} \\ x_t^{real} \\ x_t^{final} \\ i_t \end{pmatrix} \tag{1}$$

where π_t denotes inflation, i_t the nominal interest rate, and x_t a measure of real activity. In our benchmark, x_t will be GDP growth, but in the online appendix we also study employment growth.[3] A superscript $real$ denotes real time data, whereas the superscript $final$ denotes final data. Throughout the paper real-time data refers to the first available release of a data point. We want to recover the joint dynamics of real-time and final data and ask what those dynamics tell us about the effects of monetary policy shocks on both real-time and final data. The dynamics of y_t are given by

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

where the intercepts μ_t, the coefficients on lagged observables $A_{j,t}$, and the covariance matrix Ω_t of e_t are allowed to vary over time. Following most of the literature that has used these models on quarterly data such as Del Negro & Primiceri (2015) and Amir-Ahmadi, Matthes & Wang (2016), we set the number of lags $L = 2$.

By writing down a model for real-time and final versions of the same data series, we have also implicitly defined a model of the measurement errors η_t^π and η_t^x:

$$\begin{pmatrix} \eta_t^\pi \\ \eta_t^x \end{pmatrix} = \begin{pmatrix} \pi_t^{real} \\ x_t^{real} \end{pmatrix} - \begin{pmatrix} \pi_t^{final} \\ x_t^{final} \end{pmatrix} = S y_t \tag{3}$$

[3]Jointly modeling the dynamics of real-time and final inflation, GDP growth, employment growth, and the nominal interest rate leads to issues of numerical instabilities in the Gibbs sampler we use to estimate the model. We thus study different variants of the model including one indicator of real activity at a time. We could have reduced the lag length, but that would have made our results less comparable to others in the literature. Similar issues are documented in Benati (2014), for example. For the same reason we also refrain from including intermediate data revisions as observables.

where S is a selection matrix.[4] We thus use a flexible time series model for the measurement errors that does not impose strong restrictions on the measurement errors - they can be correlated, have non-zero means, and feature substantial time variation in conditional moments.[5] This is important since Aruoba (2008) has found that data revisions are not necessarily well behaved.

To see that our model can capture measurement errors that feature both 'news' and 'noise' components (i.e. the measurement errors can be correlated with both final and real-time data), we can use a toy version of our model for a generic scalar variable c_t without time variation, stochastic volatility, or any dynamics:

$$\begin{bmatrix} c_t^{final} \\ c_t^{real} \end{bmatrix} = e_t = \begin{bmatrix} e_{1,t} \\ e_{2,t} \end{bmatrix} \tag{4}$$

where $e_t \sim N(0, \Omega^c)$. The measurement error $\eta_t^c = c_t^{real} - c_t^{final}$ is then defined as $e_{2,t} - e_{1,t}$. Consider as an example the following model for e_t:

$$e_{1,t} = w_t$$

and

$$e_{2,t} = w_t + v_t$$

where w_t and v_t are independent Gaussian random variables. Then we have $\eta_t^c = v_t$, which is independent of the final data $c_t^{final} = w_t$. Reversing the roles of c_t^{final} and c_t^{real} shows that measurement error can be independent of real-time data in our framework. To see an intermediate case where the measurement error is correlated with both real-time and final data, assume that $e_{1,t} = w_t$, as before, but now $e_{2,t} = 2w_t + v_t$ so that $\eta_t^c = w_t + v_t$, which is correlated with both real-time and final data.

To concisely describe the model we use to study time variation in the parameters

[4]
$$S = \begin{pmatrix} 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 \end{pmatrix}$$

[5]The measurement errors inherit these features from the variables in the VAR.

of the model, we define $X_t' \equiv I \otimes (1, y_{t-1}' ..., y_{t-L}')$ and rewrite (2):[6]:

$$y_t = X_t' \theta_t + e_t \tag{5}$$

$$\theta_t = \theta_{t-1} + u_t \tag{6}$$

Following Primiceri (2005), it is convenient to break the covariance matrix of the reduced-form residuals into two parts as implied by the following equation:

$$e_t = \Lambda_t^{-1} \Sigma_t \varepsilon_t \tag{7}$$

where ε_t is a vector of independently and identically distributed (iid) Gaussian innovations with mean 0 and covariance matrix I. 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. Those elements vary over time according to:

$$\lambda_t^i = \lambda_{t-1}^i + \zeta_t^i \tag{8}$$

$$\log \sigma_t^j = \log \sigma_{t-1}^j + \nu_t^j \tag{9}$$

All innovations are normally distributed with covariance matrix V, which, following Primiceri (2005), we restrict as follows:

$$V = Var \left[\begin{pmatrix} \varepsilon_t \\ u_t \\ \zeta_t \\ \nu_t \end{pmatrix} \right] = \begin{pmatrix} I & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & T & 0 \\ 0 & 0 & 0 & W \end{pmatrix} \tag{10}$$

T is further restricted to be block diagonal, which simplifies inference. ζ_t and ν_t are vectors that collect the corresponding scalar innovations described above. We estimate this model using the Gibbs sampling algorithm described in Del Negro &

[6] I denotes the identity matrix.

Primiceri (2015)[7].

We follow Primiceri's choice of priors, adjusted for the size of our training sample. The Gibbs sampler we use is outlined in detail in Del Negro & Primiceri (2015). In contrast to Cogley & Sargent (2005), we do not impose any restrictions on the eigenvalues of the companion form matrix of the VAR. We do so both on empirical grounds (in Amir-Ahmadi et al. (2016) we show that there is a substantial probability of temporarily explosive dynamics in US data) and theoretical grounds (Cogley, Matthes & Sbordone (2015) show that temporarily explosive dynamics can emerge naturally in micro-founded dynamic equilibrium models when agents are learning).

In order to ascertain whether or not monetary policy shocks affect real-time and final data differently and if those effects have changed over time, we identify monetary policy shocks using our VAR models. As our benchmark, we use sign restrictions. An identification scheme of this sort has been used in time-varying parameter VARs with stochastic volatility by Benati & Lubik (2014), Canova & Gambetti (2009), and Amir-Ahmadi et al. (2016), among others.

Structural models used by macroeconomists give us a good sense of the signs of the effects of monetary policy shocks on final data. The corresponding effects on real-time data are less clear, and depend on the specifics of any particular DSGE model with both real-time and final data (we present one such model in the next section). This consideration leads us to only use sign restrictions on final data, not on real-time data. We are thus not imposing any restrictions on the impulse response functions of real-time data. We restrict the nominal interest rate to not decrease after a positive monetary policy shock and both final inflation and final GDP growth to not increase after a positive monetary policy shock. We impose those restrictions on impact and for the first two periods after impact - this is the same number of periods as chosen by Benati (2010), for example. While Uhlig (2005) did not impose restrictions on output in his application of sign restrictions, it is by now commonplace in the literature to impose restrictions on output or output growth as well (see, for example, Canova &

[7] We use 250,000 posterior draws, out of 200,000 are used as burn-in. We then keep every 10th draw of the remaining 50,000, resulting in 5000 stored draws. We have assessed and ensured convergence of the Markov Chain using the standard diagnostics.

Gambetti (2009)).

The equation for the nominal interest rate that is recovered using our identification scheme gives, by construction, the nominal interest rate as a function of lagged real-time and final data. The lagged final data is not directly observable by the central bank when it makes its decisions every period. As such, we do not directly interpret the nominal interest rate equation as a monetary policy rule (in contrast to Canova & Gambetti (2009)), but instead interpret it as the central bank responding to observables such as survey and forecast data, which in turn depend on both real-time and final data. This assumption can be justified by referring to micro-founded structural models where the private sector has an informational advantage and thus knows the final data before the central bank does. The online appendix describes one DSGE model with these features. That model shares features with work by Aoki (2003), Nimark (2008), Lubik & Matthes (2014), and Svensson & Woodford (2004), for example.[8]

Given that, for computational reasons, we can not include additional observables such as intermediate data releases, the only viable alternative would have been to restrict the central bank to only react to lagged real-time (i.e. first release) observables. This approach would have substantially underestimated the information available to the central bank. We think our approach better approximates the actual (large) information sets considered by central banks when making their decisions.[9]

[8] To keep their models tractable, those papers either assume relatively simple stochastic processes for the measurement error that can not match our findings on the properties of measurement errors, or they assume that the true realization of the data is observed by the central bank, but only with a lag.

[9] In future work, we plan to relax this assumption and instead of using 'final' data use the latest available vintage of data each quarter as new data becomes available. We view this as a separate project since in such a VAR we could no longer study the joint dynamics of real time and final data, which is our main goal in the current paper.

3 Data

As our benchmark, we use the Philadelphia Fed's real-time database (Croushore & Stark (2001)) to construct a sample of annualized quarterly real-time and final inflation (based on the GNP/GDP deflator) and annualized quarterly real-time and final real GNP/GDP growth[10]. Real-time growth rates are calculated using all available data when an estimate of the latest level of the corresponding series is first available - the growth rate of GDP over the last year at any point in time is defined as the ratio of the latest real GDP release to the most current available vintage of real GDP one quarter earlier, for example.

As a proxy for final data, we use the most recent vintage available to us. Other approaches are certainly possible - Aruoba (2008) defines the final data as the vintages available after a fixed lag (for most variables 3 years). In the online appendix we present additional results that use this alternative definition of final data.

The real-time data is available starting in the fourth quarter of 1965. The last vintage we use is from the second quarter of 2014 (incorporating data up to and including the first quarter of 2014). We use 40 observations to initialize the prior for our time-varying-VAR model. For the nominal interest rate (which is measured without error), we use the average effective Federal Funds rate over each quarter.[11]

Figure 1 plots the real-time data and the measurement error as defined in the previous section. To get the final data, we have to subtract the measurement error from the real-time data - positive measurement error implies that the real-time measurement is higher than the final data. To convince yourself that the difference between real-time and final data can be meaningful, it suffices to look at the mid-1970s: Real-time GDP growth actually was lower than in the most recent recession, but there were substantial revisions to that data later on - the measurement errors associated with those errors is negative, meaning that data was revised upward substantially.[12]

[10]From now on we will refer to this variable as GDP growth.

[11]In the appendix, we also present a version of the model that uses the shadow rate from Wu & Xia (2016).

[12]Lubik & Matthes (2014) use a learning model to model the choices of a central bank that only has access to real-time data as it makes its decisions. Just as Orphanides (2002), they highlight that mis-measured data had a big impact on U.S. monetary policy in the 1970s. In the current paper, we

In the following section we will analyze the time series properties of the measurement errors and check how their behavior has changed over time. We will see that it is indeed important to allow for time variation in the dynamics of these series.

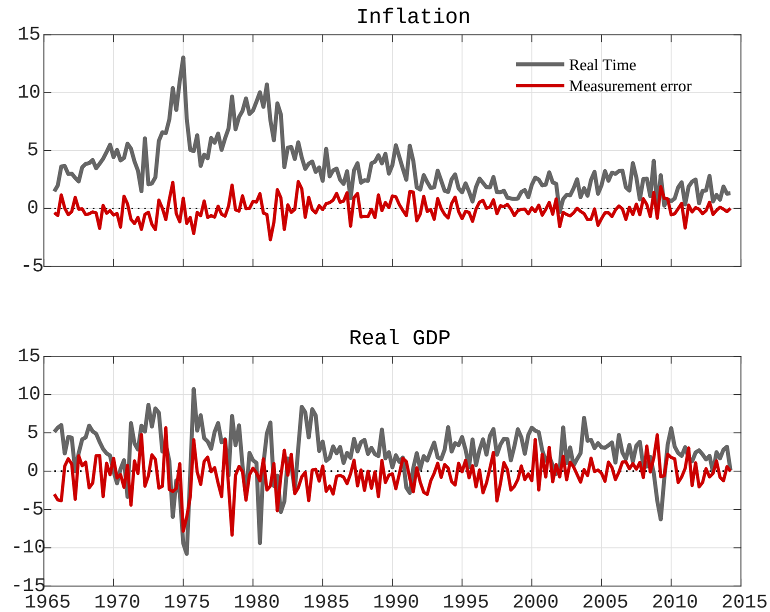

Figure 1: Real-time data and measurement error

4 Results

4.1 The Time-Varying Properties of Measurement Errors

First, we want to describe how the properties of the measurement errors in inflation and GDP growth have changed over time. Cogley & Sargent (2005) have pioneered the use of 'local to time t' moments to study changes in the dynamics of VARs with time-varying parameters and stochastic volatility. In short, they calculate (unconditional) moments of the data governed by equation 2 at each point in time assuming that the coefficients will remain fixed over time. That way they recover a sequence of moments over time. This is feasible in their setup because they impose restric-

instead focus on the impact of monetary policy shocks on both real-time and final data.

tions on the eigenvalues of the companion form matrix of the VAR. We, on the other hand, do not impose any such restrictions for the reasons mentioned before. Instead, we study forecasts from our model based on smoothed or full-sample parameter estimates (assuming, similar to Cogley & Sargent (2005), that the coefficients will not change in the future) and calculate the moments of forecasts of the measurements errors. It is important to emphasize that we use these moments of forecasts as low-dimensional summary statistics that capture the dynamics of our model. We can think of these moments as finite horizon versions of the summary statistics used by Cogley & Sargent (2005). We focus here on one-year ahead forecasts. Increasing the forecast horizon substantially would increase the uncertainty surrounding the estimated forecast moments exactly because we do not impose any restrictions on the dynamics of the VAR.

Figure 2 plots the median and 68 % posterior bands for the one-year ahead forecasts

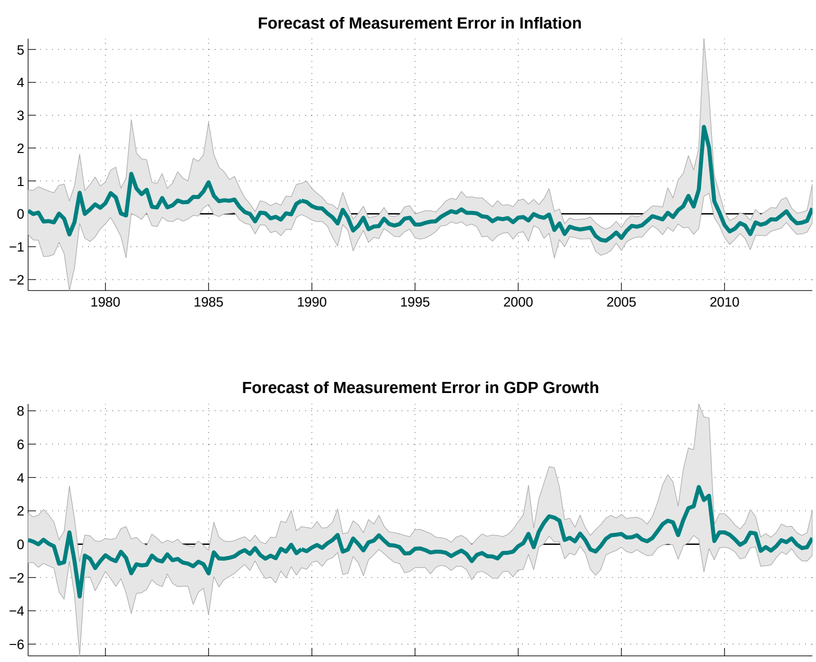

Figure 2: One-year ahead forecasts of measurement errors

of the measurement error based on the model with GDP growth.[13] Our model pre-

[13]The date on the x-axis represents the date of the conditioning information.

dicts substantial measurement errors one year in advance. We can think of these one-year ahead forecasts as a proxy for trends or more generally the persistent parts of the measurement errors.[14] Our results confirm those in Aruoba (2008), who finds that measurement errors in many variables do not have a mean of zero. Throughout most the 1980s the one-year ahead forecast in the measurement error of inflation is positive, of an economically meaningful size, and borderline statistically significant - inflation was initially overestimated during that period. During the 1990s and up to the financial crisis, inflation instead tended to be initially underestimated. During the financial crisis, inflation was substantially overestimated initially.

The measurement error in real GDP growth is negative during the 1980s (meaning that GDP growth was initially estimated to be lower than the final data suggests), before turning statistically insignificant during the 1990s. From 2000 to the financial crisis we see an initial overestimation of GDP growth (with a substantial overestimation during the financial crisis).

We now turn to higher moments of the forecasts. Figure 3 plots the volatilities of the one-year ahead forecasts of measurement errors and the associated correlation between the forecasted measurement errors. Both volatilities share a similar pattern[15] - high volatility in the 1970s and early 1980s, a decline afterward and a noticeable uptick in volatility during the recent financial crisis. Interestingly, the correlation between the measurement errors is significantly negative throughout our sample, but has an upward trend for most of our sample that is only broken during the early 2000s. A negative correlation implies that an increase in the measurement error of GDP growth (real-time GDP growth becomes larger relative to final GDP growth) is associated with a decrease in the measurement error in inflation (final inflation becomes larger relative to the real-time measurement), so that an initial overestimation of GDP growth tends to be associated with an underestimation of inflation. Since the magnitude of the correlation decreased substantially over time, this pat-

[14] Interpreting forecasts as trends has a long tradition in empirical macroeconomics going back to Beveridge & Nelson (1981).

[15] It is common in models of the type we use here, used in conjunction with the type of data that we analyze, that the volatilities of the variables in the VAR share a similar pattern - see for example Del Negro & Primiceri (2015).

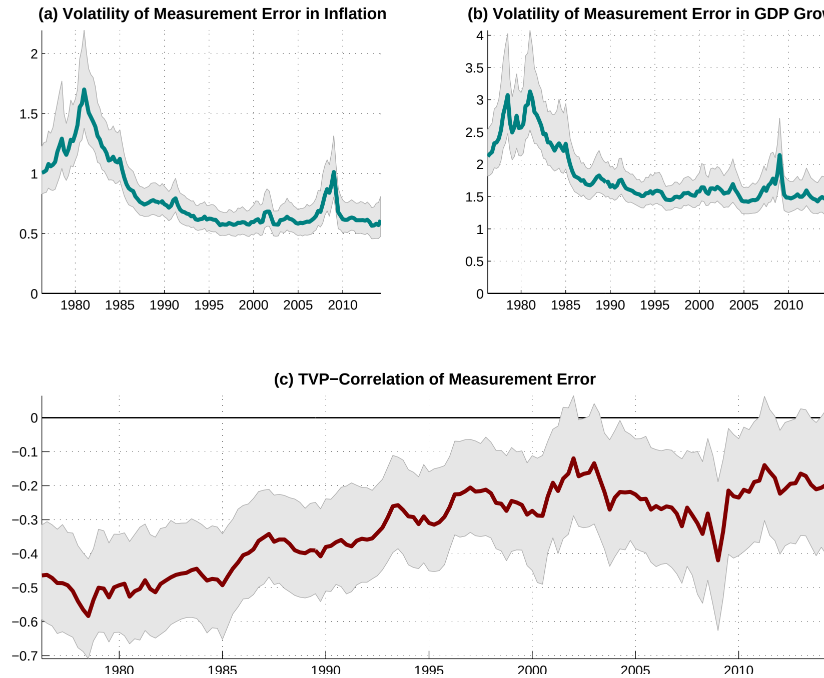

Figure 3: Volatility and correlation of forecasted measurement errors

tern has become weaker over time. To summarize, a simple model of measurement errors that models them as being independent across variables and having constant innovation variance can miss important features of observed measurement errors.

4.2 The Effects of Monetary Policy Shocks Over Time

We first show impulse responses for different periods. We follow the standard approach in the literature to construct these impulse responses: For each time period, we draw parameters from the posterior distribution for that period and then keep these coefficients fixed as we trace out the effects of a monetary policy shock. We focus on impulse responses at short horizons because that is where we find the largest difference between real-time and final data. Since we are interested in the differences between the effects on real-time and final data (rather than changes in the impulse responses functions over time *per se*), we use one standard deviation shocks, where the standard deviation changes over time. This will give us a sense of how the impact of a usual shock has changed over time. Figure 4 plots the evolution of the nominal interest rate to such a shock. The black line gives the pointwise median response and the gray bands cover the area from the 15th to the 85th percentile of the response with each of the 5 shades of gray covering the same probability. We can see that there are differences on impact over time (in particular, the standard deviation of monetary policy shocks decreases), but the overall median pattern remains stable over time. In contrast, there is substantial time variation in the uncertainty surrounding the median response. At some points in time there are some draws that imply explosive behavior of the nominal interest rate.

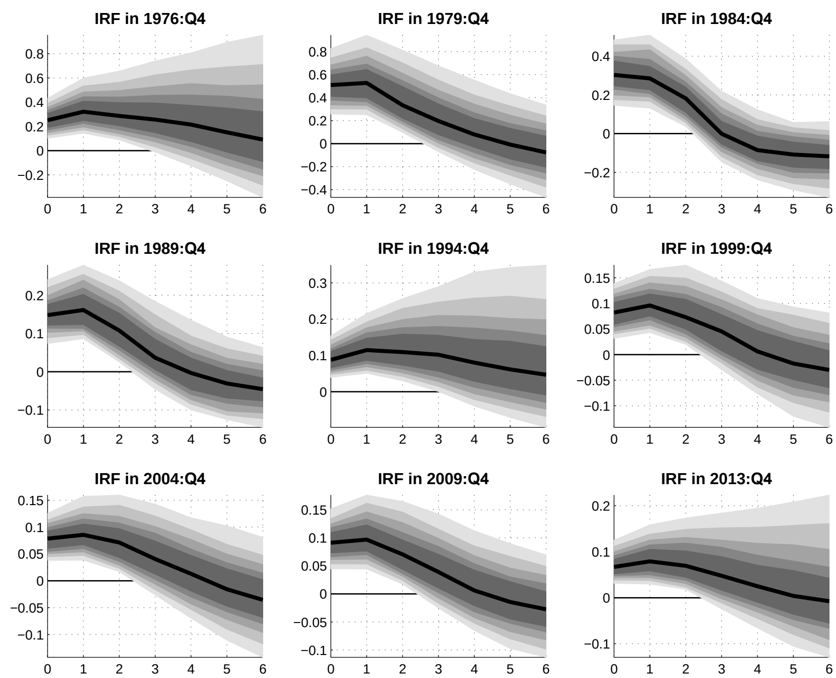

Figure 4: Impulse response functions for the nominal interest rate to a one standard deviation monetary policy shock.

Figure 5 plots the responses of real-time and final inflation to the same monetary policy shock. The black line and gray areas correspond to the median and the 15th to 85th percentiles of real-time data responses, whereas the red lines represent the responses of final data. The bold red line is the median and the outer dashed red bands correspond to the same percentiles as the outermost error bands for real-time data (the 15th and 85th percentiles). For the most part the responses of real-time and final inflation are very similar, especially after 4 to 5 periods. The sign restrictions are mostly satisfied by responses of real-time inflation even though we do not impose those restrictions.[16] Nonetheless, we do find significant differences. For example, in 1979 the median impact response of real-time inflation is twice as large as that of final data. Broadly speaking, we see a larger difference (on impact) for the first part of our sample (through the 1980s). Substantial differences in the responses between real-time and final inflation are present in the late 1970s to the late 1980s.

The impulse responses for GDP growth in figure 6 show a different pattern with more pronounced differences. On impact and for the first few periods after the shock hits, final GDP growth is lower than real-time GDP growth. This pattern is most pronounced in 1984 and 1989, but persists throughout our sample. The magnitude of those differences is economically significant - it matters if the response to a contractionary monetary policy shock on impact is a reduction of 0.25 percentage points in annualized GDP growth or 0.75 percentage points (these are roughly the magnitudes in 1984:Q4). We can also see that the sign restrictions we impose on final data are also met for most draws of the real-time data response.

So far we have studied the marginal distributions of the impulse responses to real-time and final data and compared them to each other. We are also interested in the evolution of the joint distribution of impulse responses across real-time and final data. Our estimation algorithm allows us to study the joint posterior of impulse responses for a given horizon at each point in time. For each of those time/horizon pairs, we calculate an estimate of the joint posterior of real-time and final impulse responses (for each horizon and date this can be thought of as a scatterplot). We call $r_t^{real,i}(j)$ the impulse response at horizon j of real-time variable i ($i \in \{\pi, GDP, emp\}$)

[16] In the online appendix we show that our results also hold of we impose sign restrictions on both real-time and final data.

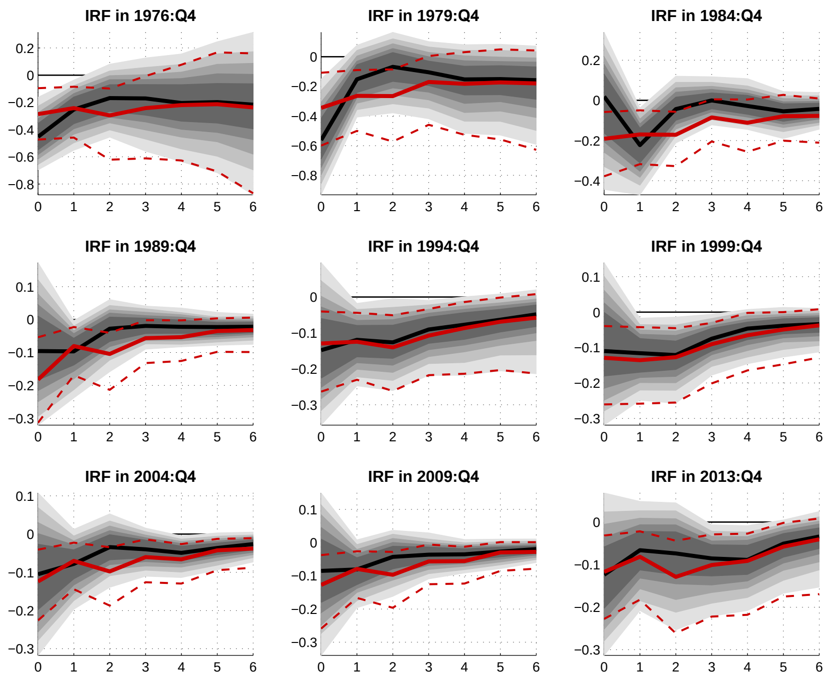

Figure 5: Impulse response functions for real-time (gray/black) and final (red) inflation to a one standard deviation monetary policy shock.

calculated using a draw of VAR coefficients at time t and $r_t^{final,i}(j)$ the response of the corresponding final variable (both calculated using the same parameter draw).

We first plot the median and the 15th and 85th percentile bands for the difference between final data and real-time impulse responses of GDP growth[17] on impact (i.e. at horizon 0): $r_t^{final,GDP}(0) - r_t^{real,GDP}(0)$. A negative number means that the final data response is smaller than the corresponding real-time response. Figure 7 reveals that the median difference has been negative throughout our sample with a maximum of -0.1 and a minimum of -0.5 percentage points,[18] meaning that the response of final data is larger in magnitude than the response of real-time data as the final response

[17]The median difference for inflation is centered at 0 for most of the sample, so we omit it here. This finding is also evident from figure 8.

[18]Remember that we use annualized values throughout this paper.

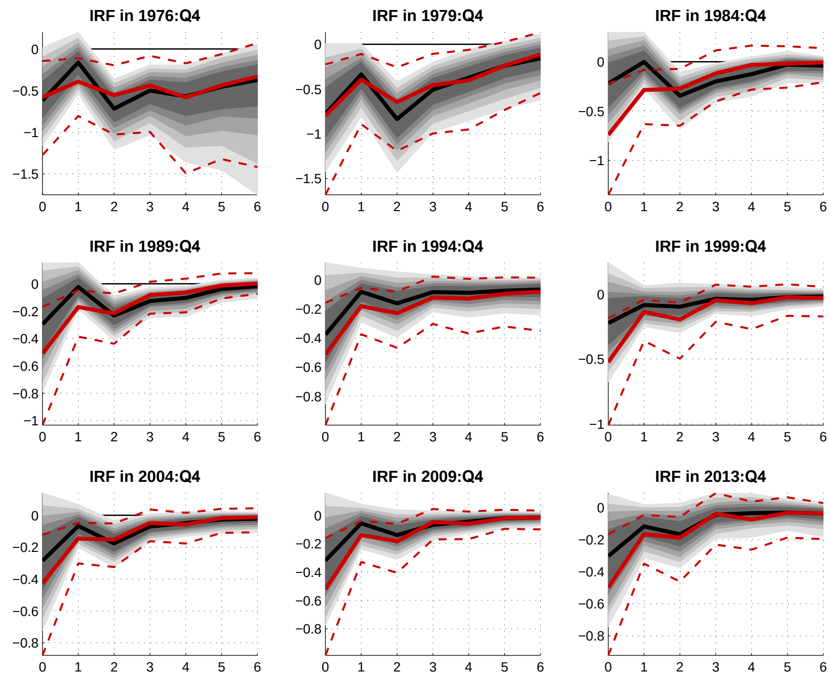

Figure 6: Impulse response functions for real-time (gray/black) and final (red) real GDP growth to a one standard deviation monetary policy shock.

is restricted to be negative on impact and the real-time response is negative for most draws. This again emphasizes that the differences are economically meaningful - central banks would care about these magnitudes. The 85th percentile of the difference hovers around 0. Thus, there is a positive probability that the difference is close to 0 at any point in our sample as can be seen by the point-wise error bands[19]. However, the fact that the median difference is negative throughout and of a eco-

[19]Note that the error bands are calculated based on the marginal distribution of the differences each period. They do not directly take into account information about the difference in the proceeding and following periods (i.e. the joint distribution of the difference across periods). The bands based on the marginal distribution only take into account information about other periods in an indirect fashion since they are based on smoothed (full sample) parameter estimates. For fixed coefficient VARs with sign restrictions, issues with pointwise error bands have been highlighted by Inoue & Kilian (2013), for example. It is not clear how to extend their methods to VARs with time-varying coefficients and stochastic volatility in general and to our question at hand in particular.

nomically significant magnitude leads us to believe that it is indeed important to take the difference between real-time and final data seriously. Furthermore, policymakers regularly worry about worst case outcomes. We can see that the difference between the impact response for final and real-time data could be substantially larger in magnitude than what is suggested by the median numbers.

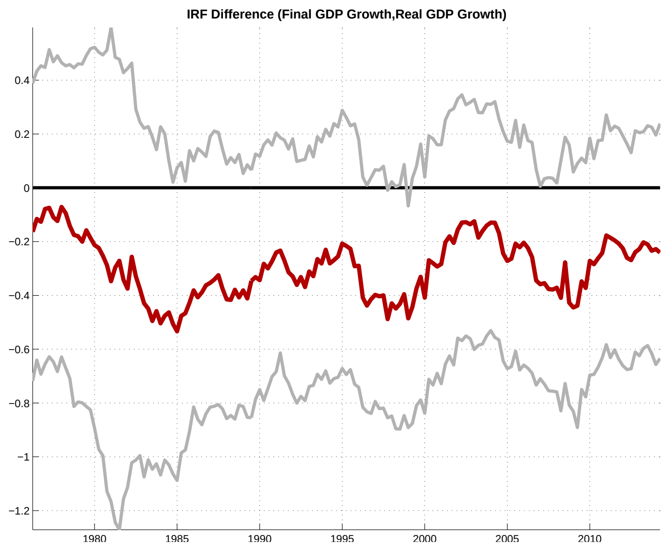

Figure 7: Differences between GDP growth impulse responses on impact: The distribution of $r_t^{final,GDP}(0) - r_t^{real,GDP}(0)$ over time.

For each period in our sample, we then regress the real-time responses at that point in time on the final responses at the same point in time and a constant:

$$r_t^{real,i}(0) = \alpha_t^i + \beta_t^i r_t^{final,i}(0) + u_t \tag{11}$$

Thus, each hypothetical scatterplot is summarized by two numbers, the constant α_t^i and the slope β_t^i. We focus on the contemporaneous response since the differences are largest for small horizons. Since the sample size for each regression is given by the number of draws we use to calculate the impulse responses, we do not report standard errors for the coefficients - these standard errors would be tiny. If responses based on real-time data are just a noisy version of the responses based on final data, we would expect the intercept α_t^i to be zero and the coefficient on the responses for final data β_t^i to be 1. Otherwise there is a bias in the real-time data responses relative to the responses based on final data that economists are actually interested in. Figure 8 shows how the intercept and the coefficient on the final-data response vary over time for the case of the contemporaneous response to a monetary policy shock. The gray line represents the slope of the regression β_t^i (right axis) and the red line represents the intercept α_t^i (left axis). Both paths show a similar pattern: Until 1980 there is a clear bias. After 1980 the coefficients quickly move toward values that imply no bias.

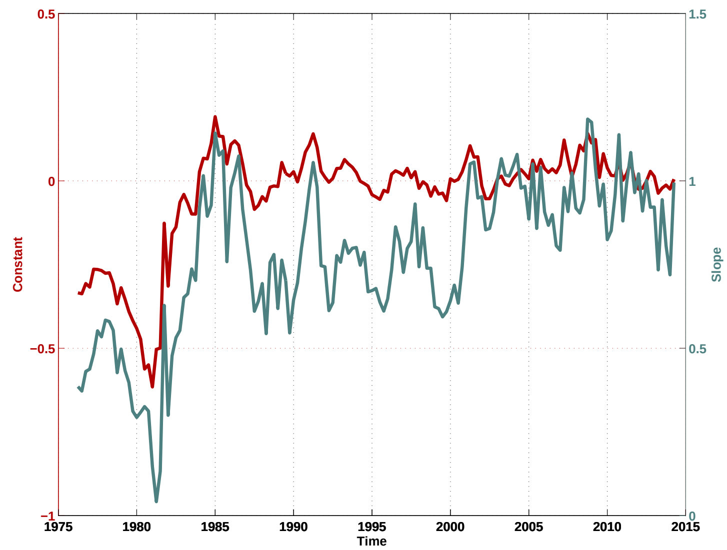

Figure 8: Relationship between inflation real-time and final data based impact impulse responses over time. Intercept α_t^π in red and slope β_t^π in gray.

Figure 9 shows the results for the same regressions in the case of the contemporaneous response of real-time and final GDP growth. We see a broadly similar pattern for the intercept that moves toward zero after 1980. There is no substantial shift in the behavior of the slope, though. The slope is never as small as the minimum slope for inflation, but it also does not substantially move toward 1 after 1980. Real-time GDP growth responds differently than final GDP growth to a monetary policy shock on impact in systematic fashion throughout our sample. We think of these results as a cautionary tale about the information content of real-time data releases of GDP growth. It is important to remember here that we try to recover the true response of real-time data to a monetary policy shock, not the response to a monetary policy shock that can be recovered in real-time.

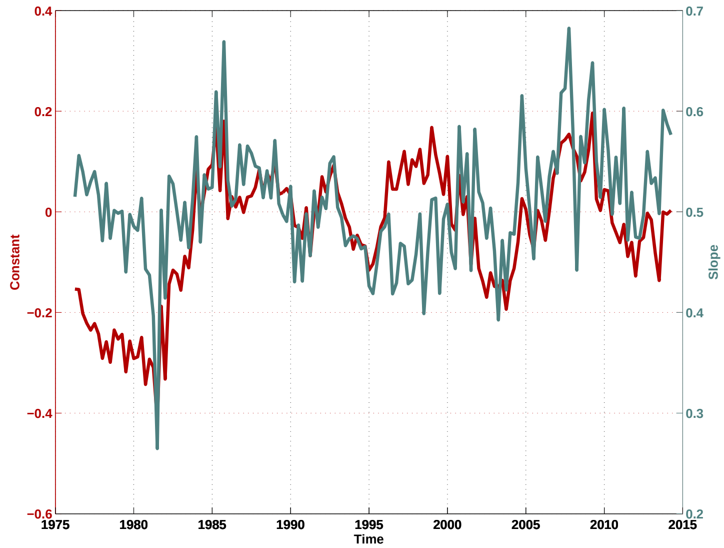

Figure 9: Relationship between GDP growth real-time and final data based impact impulse responses over time. Intercept α_t^{GDP} in red and slope β_t^{GDP} in gray.

5 Conclusion

Measurement errors are pervasive in real-time macroeconomic data. We extend the insights of Aruoba (2008) to incorporate time varying dynamics and document that these measurement errors feature substantial time-varying volatility, can be correlated with a time-varying correlation, and are not centered around zero. Thus, modeling real-time data as the sum of the final data and a simple independent noise process can miss important features of the data.

We show that these facts are not a curiosity, but have policy implications: (i) These differences between real-time data and final data manifest themselves in the substantially different ways that real-time and final data respond to monetary policy shocks, and (ii) the real-time responses can be substantially biased. Furthermore, the responses of various measures of real activity are larger in magnitude for final data. How does this directly affect policymakers? Policymakers must base their decision on real-time data for the last few periods and can not run a VAR specification with *both* real-time and final data that we use throughout this paper. Nonetheless, the finding that final GDP growth reacts more strongly to a monetary policy shock than real-time GDP growth can be used by policymakers: If real-time data points to a response of GDP growth of a given magnitude to a monetary policy surprise[20], policymakers should be aware that the actual response of final GDP growth will tend to be larger in magnitude and plan their next policy decision accordingly. A policymaker could even go further: Under the assumption that recent changes in real-time GDP growth are mainly driven by a monetary policy surprise (so that other shocks do not play a significant role), one could use recent estimates of the coefficients in equation (11) (which links the response of real-time and final data) to get an estimate of the associated response of final data.[21]

[20]Whether or not a given policy action is a policy *surprise* can be deduced from measures of monetary policy surprises that are available in real-time, see for example Kuttner (2001).

[21]There would be an approximation error because the estimates of the coefficients are not available for the current period. It might be advisable then to use an average of the coefficients over the last few available periods to smooth out any high frequency noise in the changes of those coefficients.

References

Amir-Ahmadi, P., Matthes, C. & Wang, M.-C. (2016), 'Drifts and Volatilities under Measurement Error: Assessing Monetary Policy Shocks over the Last Century', *Quantitative Economics* **7**(2), 591–611.

Aoki, K. (2003), 'On the optimal monetary policy response to noisy indicators', *Journal of Monetary Economics* **50**(3), 501–523.

Aruoba, S. B. (2008), 'Data Revisions Are Not Well Behaved', *Journal of Money, Credit and Banking* **40**(2-3), 319–340.

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

Benati, L. (2014), Economic Policy Uncertainty and the Great Recession, Working paper, University of Bern.

Benati, L. & Lubik, T. (2014), 'Sales, inventories, and real interest rates: A century of stylized facts', *Journal of Applied Econometrics* p. forthcoming.

Beveridge, S. & Nelson, C. R. (1981), 'A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'', *Journal of Monetary Economics* **7**(2), 151–174.

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.

Carriero, A., Clements, M. & Galvao, A. B. (2015), 'Forecasting with Bayesian Multivariate vintage-based VARs', *International Journal of Forecasting* **31**, 757768.

Cogley, T., Matthes, C. & Sbordone, A. M. (2015), 'Optimized Taylor rules for disinflation when agents are learning', *Journal of Monetary Economics* **72**(C), 131–147.

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.

Croushore, D. & Evans, C. L. (2006), 'Data revisions and the identification of monetary policy shocks', *Journal of Monetary Economics* **53**(6), 1135–1160.

Croushore, D. & Sill, K. (2014), Analyzing Data Revisions With a Dynamic Stochastic General Equilibrium Model, Working Paper 14-29, Federal Reserve Bank of Philadelphia.

Croushore, D. & Stark, T. (2001), 'A real-time data set for macroeconomists', *Journal of Econometrics* **105**(1), 111–130.

D'Agostino, A., Gambetti, L. & Giannone, D. (2013), 'Macroeconomic forecasting and structural change', *Journal of Applied Econometrics* **28**(1), 82–101.

Del Negro, M. & Primiceri, G. (2015), 'Time-varying structural vector autoregressions and monetary policy: a corrigendum', *Review of Economic Studies* .

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.

Inoue, A. & Kilian, L. (2013), 'Inference on impulse response functions in structural VAR models', *Journal of Econometrics* **177**(1), 1–13.

Jacobs, J. P. A. M., Sarferaz, S., van Norden, S. & Sturm, J.-E. (2013), Modeling Multivariate Data Revisions, CIRANO Working Papers 2013s-44, CIRANO.

Jacobs, J. P. & van Norden, S. (2011), ‘Modeling data revisions: Measurement error and dynamics of true values’, *Journal of Econometrics* **161**(2), 101–109.

Kishor, N. K. & Koenig, E. F. (2009), ‘VAR Estimation and Forecasting When Data Are Subject to Revision’, *Journal of Business & Economic Statistics* **30**(2), 181–190.

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

Kuttner, K. N. (2001), ‘Monetary policy surprises and interest rates: Evidence from the Fed funds futures market’, *Journal of Monetary Economics* **47**(3), 523–544.

Lubik, T. A. & Matthes, C. (2014), Indeterminacy and Learning: An Analysis of Monetary Policy in the Great Inflation, Working Paper 14-2, Federal Reserve Bank of Richmond.

Nimark, K. (2008), ‘Monetary policy with signal extraction from the bond market’, *Journal of Monetary Economics* **55**(8), 1389–1400.

Orphanides, A. (2002), ‘Monetary-Policy Rules and the Great Inflation’, *American Economic Review* **92**(2), 115–120.

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

Schorfheide, F., Sill, K. & Kryshko, M. (2010), ‘DSGE model-based forecasting of non-modelled variables’, *International Journal of Forecasting* **26**(2), 348–373.

Svensson, L. E. O. & Woodford, M. (2004), ‘Indicator variables for optimal policy under asymmetric information’, *Journal of Economic Dynamics and Control* **28**(4), 661–690.

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.

Wu, J. C. & Xia, F. D. (2016), 'Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound', *Journal of Money, Credit and Banking* **48**(2-3), 253–291.