

Measuring the Effects of Expectations Shocks

Michael P. Clements	Ana Beatriz Galvão*
ICMA Centre	Warwick Business School
Henley Business School	University of Warwick
University of Reading	Ana.Galvao@wbs.ac.uk
M.P.Clements@reading.ac.uk	

December 6, 2019

Abstract

We show that expectation shocks – revisions in expectations unrelated to changes in current economic fundamentals – have positive significant effects on US economic activity. To measure the expectation shocks, we estimate a mixed-frequency VAR model that allows economic conditions in the current quarter to affect current-quarter GDP expectations. The expectations shock is estimated with real-time data so such shocks do not suffer a ‘look-forward’ bias by incorporating future data revisions. Dynamic responses are estimated with the aid of a quarterly VAR and using older vintages as instruments to account for measurement errors in the observed values. Expectations shocks explain 10% of the two-year variation of output, investment, consumption and hours. We find that expectations shocks are correlated with alternative belief-based shocks, but nevertheless have significant additional short-run effects.

Key words: mixed-frequency Vector Autoregressive Models; real-time data; measurement errors; expectational shocks.

JEL code: E37.

*Michael Clements is also an Associate member of the Institute for New Economic Thinking at the Oxford Martin School, University of Oxford. Corresponding author: Ana Beatriz Galvao; email: ana.galvao@wbs.ac.uk. We are grateful to Tony Garratt, Ivan Petrella, Giovanni Ricco, Tomas Götz and participants of the Escoe conference and the Bundesbank Macro Workshop for helpful comments.

1 Introduction

There is a venerable literature on the effects of changes in agents' beliefs about the future on economic fluctuations dating back at least to Pigou (1927) and Keynes (1936). Changes in expectations drive business cycles because they reflect either news about future fundamentals (Beaudry and Portier (2014)) or fluctuations in sentiment (Milani (2017)) and beliefs (Angeletos, Collard and Dellas (2018)). To disentangle expectations-driven shocks from other fundamental shocks, authors have proposed identification strategies for technology news shocks (Barsky and Sims (2011)), sentiment shocks (Levchenko and Pandalai-Nayar (2017), Fève and Guay (2016) and Lagerborg, Pappa and Ravn (2019)), and confidence shocks (Barsky and Sims (2012)). Angeletos *et al.* (2018) evaluate the effects of confidence shocks using an empirical DSGE model. They claim that an empirical strategy based on vector autoregressive models and a measure of confidence would have difficulties in identifying confidence shocks because one needs to remove the effects of all anticipated future fundamental shocks. Leduc and Sill (2013) provide evidence that expectations shocks have positive effects on economic activity using a vector autoregressive model.

In this paper, we propose a method to improve the measurement of expectations-driven business cycles by addressing two major issues. The first one is that agents do not observe the true values of current and recent macroeconomic variables when they form their expectations. To be able to identify expectations shocks, agents' knowledge of the macroeconomic environment has to be correctly modelled by avoiding 'look-ahead' bias.¹ The second is that at the time that the expectation for the current and future quarters of macroeconomic variables is formed, high-frequency economic and financial indicators on the current quarter may be available. These indicators reflect the realization of fundamental shocks within the quarter have to be taken into account to identify the autonomous expectations shocks.

We propose a two-step method. In the first step, expectations shocks are estimated using US SPF (Survey of Professional Forecasters) forecasts of output growth in a real-time mixed-frequency VAR model. In the second step, we compute impulse responses using a VAR estimated with the assumed 'true' values of the underlying time series. We show how to obtain impulse response estimates under two assumptions. The first one is that after many rounds of revisions, we observe

¹The literature has generally either ignored the issue of data uncertainty, or made a judicious choice of variable so that it does not arise. For example, Leduc and Sill (2013) choose unemployment expectations, precisely because data revisions to the unemployment rate are small and can be ignored. An exception is the analysis of monetary policy shocks by Croushore and Evans (2006), and their analysis guides our approach.

the ‘true’ values (as in, e.g., Cunningham, Eklund, Jeffery, Kapetanios and Labhard (2009), Kishor and Koenig (2012) and Garratt, Lee, Mise and Shields (2008)) and the second that ‘true’ values are never observed (as in, e.g., Jacobs and van Norden (2011)).

Our empirical results suggest that expectations shocks lead to business cycle comovement across macroeconomic variables, and that the information on our measure of expectations shocks is not encompassed by confidence and technology news shocks. Expectations shocks explain 8% of the two-year variation in output and investment. In contrast to Leduc and Sill (2013), we find that expectations shocks have small effects on monetary policy.

We show that the use of latest-vintage data (e.g., the vintage available at the time the study is undertaken) will not correctly estimate the expectations shocks or their impact on the macroeconomy. The expectations shock can, however, be consistently estimated using real-time data.

As well as using only the available vintages to estimate expectations shocks in real time, we need to use all the key information available to the forecasters, at the time. Some studies based only on quarterly data neglect the higher-frequency data which in practice may be a valuable source of information. For forecasting GDP, for example, key monthly indicators such as industrial production and non-farm payroll employment for the first month of the quarter in question will generally provide relevant information, and reflecting this receive much media attention. Ignoring relevant information, monthly or otherwise, would lead to incorrect estimates of the expectations shocks. To correctly estimate expectations shocks and their impacts, we consider mixed-frequency VAR (MF-VAR) models (see, e.g., Ghysels (2016)) enabling the monthly data to be included in a convenient way.

Although we use a real-time mixed-frequency VAR to estimate the expectations shocks, the transmission of those shocks through the macroeconomy is studied in a seven-variable VAR model in log-levels, consisting of variables typically of interest to macroeconomists - we refer to this as the Macroeconomist’s vector autoregressive model. The appropriate method of estimating the response of the macroeconomy to these shocks depends upon a number of factors: whether data revisions are news or noise (as defined by Mankiw and Shapiro (1986)), whether after a number of rounds of revisions we can assume the true values are revealed, and whether the responses of activity variables (e.g., output) to expectations shocks ought to be specified in terms of true values or early-release values. Suppose the researcher wishes to measure the macroeconomic responses in a model using the ‘true’ values of the variables. Next, suppose that the true values of the variables subject to revision are eventually revealed. Then we can simply discard the more recent observations from

the latest data release, and so exclude values still subject to news revisions, when we estimate the dynamic responses to real-time expectations shocks. Alternatively, if we are not able to assume that the true values are revealed, an instrumental variable approach is required to accomplish the same end (as proposed by Croushore and Evans (2006)).

We show that the empirical findings may be altered both quantitatively and qualitatively if real-time, monthly data, is not used to measure expectations shocks. For example, we show that the dynamic effects on output, consumption, investment and hours will be incorrectly confined to the first year after the shock. The responses of prices and monetary policy are significant and positive. Estimating expectations shocks allowing agents to have observed the shocks that have taken place during the first month of the quarter is one reason why are findings are qualitatively different.² In addition, data revisions to key macroeconomic data are far from negligible (see, e.g., Aruoba (2008)), and we show that whether ‘true’ values are assumed to be eventually observed may alter our assessment of the long-run effects of expectations shocks.

We provide evidence that expectation shocks are a largely complementary source of business cycle fluctuation since they lead to cyclical comovements across macroeconomic aggregates even if orthogonalized to technological news and confidence shocks. However, our measured expectations shocks may be proxying other expectations-based shocks, such as news shocks about alternative economic fundamentals, and sentiment shocks. Indeed, SPF forecasts have been employed previously in the estimation of DSGE models to pin down expectations and to help the measurement of news shocks (Miyamoto and Nguyen (2019)) and sentiment shocks (Milani (2017)). Empirically, we find that our responses resemble those for confidence shocks found by Angeletos *et al.* (2018) in a DSGE model, where the confidence shocks are caused by frictional coordination due to higher-order beliefs.

The plan of the rest of the paper is as follows. Section 2 discusses the calculation of the expectations shocks using a real-time VAR. Section 3 explains how we determine the responses of the macroeconomy to the shocks, and presents the results. In section 4 we explore the relationship between the expectations shocks and alternative belief-based shocks. Finally, section 5 offers some concluding remarks.

²This is important, because in our expectations data, agents form their beliefs around the middle of the quarter, and will be aware of developments up to that point.

2 Measuring Expectation Shocks

In this section we explain how we estimate expectations shocks. Expectations shocks are the revisions in economic agents' forecasts of future GDP values that are not due to changes in current economic conditions. They may represent 'news' shocks, such as the future technology innovations (Beaudry and Portier (2006) and Barsky and Sims (2011)); 'animal spirits' (autonomous shocks to expectations when all the fundamental factors have been included) and confidence, as in Barsky and Sims (2012) and Angeletos *et al.* (2018); exogenous waves of optimism or pessimism, as in Milani (2011); or sentiment, as in Fève and Guay (2016).

2.1 Forecasts Updates and Available Information

We estimate expectations shocks using the U.S. Survey of Professional Forecasts (SPF) as the source of expectations data.³ SPF forecasts are made around the middle of the middle month of the quarter (survey questionnaires are required to be returned by around the end of the third week of the middle month). At quarter t , the value of the target variable Y_t is not available because of publication delays. Hence, in response to a survey at time t , there is a nowcast $Y_{t|t}$, and forecasts for the next three quarters $Y_{t+1|t}$, $Y_{t+2|t}$, $Y_{t+3|t}$ where the conditioning is on the survey t information.⁴ Forecasts updates are then defined as:

$$Y_{t+n|t} - Y_{t+n|t-1},$$

where $Y_{t+n|t}$ is the cross-sectional median of respondents' forecasts at time t of Y at $t+n$. Forecasts are recorded for real GDP in dollars. Because of changes in the base index to compute the GDP deflator, the level of real GDP will exhibit periodic level shifts. These rebasing effects would need to be removed if real GDP were to be used in levels (as in Garratt *et al.* (2008), for example). An alternative solution used in this section is to work in terms of growth rates, and growth rate forecasts, so that rebasing effects are largely absent. In section 3 the Macroeconomist's VAR is expressed in log levels.

Figure 1 presents forecast updates for $n = 0, 1, 2$ for quarterly GDP growth at annualized values for $t = 1968Q4 - 2016Q3$. These were computed as $Y_{t+n|t} = 400(\log(Z_{t+n|t}) - \log(Z_{t+n-1|t}))$, using

³Freely available at <https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters>

⁴From 1982, the SPF also includes predictions for $Y_{t+4|t}$. Our baseline results in this paper use forecasts from 1968.

SPF median forecasts for real GDP values $Z_{t+n|t}$. Use of consensus expectations, calculated as an average of the individual respondents' expectations, is standard practice in the literature.⁵

Table 1 presents summary statistics for the each forecast horizon. It is clear from the Figure and Table that the variability of forecast updates to nowcasts is twice as large as the variability of two-quarter ahead forecasts, and the correlation between nowcast updates and two-quarter ahead updates is small (8%). One would naturally expect smaller updates at longer horizons, due to the stationarity of the output growth series. Nevertheless, Bordalo, Gennaioli, Ma and Shleifer (2018) find that SPF consensus forecasts at longer horizons tend to under-react to economic news. We use the shortest horizon forecasts to measure expectation shocks. This choice is compatible with the argument in Angeletos *et al.* (2018) that variations in confidence are related to changes in the short-term outlook. Updates to longer horizon forecasts may be more related to news shocks, as suggested by Miyamoto and Nguyen (2019). We check the robustness of our main results to this choice in the next section.

2.2 Expectations Shocks and GDP Revisions

A forecast update at time t may be due to: *i*) new information that has arrived between $t-1$ and t , *ii*) sluggish adjustment to previous news due to inattentiveness or sticky information (see, e.g., Sims (2003), Mankiw and Reis (2002)), or *iii*) changes in confidence (Barsky and Sims (2012)) and news about fundamentals (Beaudry and Portier (2006)). We aim to remove the effects of the first two possibilities by filtering forecast updates using a vector autoregressive model. An advantage of using a VAR model to calculate the shocks, as opposed to a nowcasting model (as surveyed by Bańbura, Giannone, Modugno and Reichlin (2013)), is that we can obtain a shock series which is orthogonal to a set of other shocks. We use recursive identification based on a Cholesky decomposition, where the ordering of the variables - the nature of the recursive structure - is suggested by the time-line of the various data releases and when forecasts are produced.

The literature on the sources of business cycles and the propagation of shocks typically employs the latest-available vintage on Y_t (and other variables subject to revision) to undertake the analysis.⁶ It may be tempting to discount the importance of accounting for data revisions, but data revisions can be large relative to the variability in the series (see, e.g., Aruoba (2008)). Croushore (2011b,

⁵That said, concerns have been raised about the use of the consensus, especially when when the composition of the panel is changing over time, most notably by Engelberg, Manski and Williams (2011).

⁶This literature is surveyed by Beaudry and Portier (2014).

2011a) provide review articles, and Jacobs and van Norden (2011), Cunningham *et al.* (2009), Kishor and Koenig (2012) and Garratt *et al.* (2008) are key papers considering various ways of modelling data subject to revision, and Clements and Galvão (2019) provide a recent review.

Were we to download a recent data vintage, say the 2016Q4 vintage, Y_t^{16Q4} for $t = \dots, 16Q2, 16Q3$, we would be implicitly assuming forecasters observe Y_{t-1}^{16Q4} at the time t that the forecast is made. In real-time, however, the forecaster at time t only has access to the first release of GDP in the previous period, that is, Y_{t-1}^t . The data revision ($Y_{t-1}^{16Q4} - Y_{t-1}^t$) is in part unpredictable at t (see, for example, the literature cited in the previous paragraph). This is important because we wish to measure the effect of expectations shocks, and care is required not to contaminate these with data revisions. Section 2.3 illustrates with a simple model.

2.3 A Simple Model Illustrating the Effects of Data Revisions

A simplified model will help clarify the impact of data revisions on the measurement of expectations shocks and their transmission. Suppose the bivariate structural VAR (SVAR) for expectations updates about the fundamental variable Y_t , $Y_{t|t} - Y_{t|t-1}$, and the first-release values of the fundamental variable, Y_t^{t+1} , takes the form:

$$Y_{t|t} - Y_{t|t-1} = a_{11}(Y_{t-1|t-1} - Y_{t-1|t-2}) + a_{12}Y_{t-1}^t + u_t^{\text{exp}} \quad (1)$$

$$Y_t^{t+1} = a_{21}(Y_{t-1|t-1} - Y_{t-1|t-2}) + a_{22}Y_{t-1}^t + a_{0,12}u_t^{\text{exp}} + u_t^{\text{fund}}, \quad (2)$$

for $t = 2, \dots, T$. For simplicity we ignore intercepts and consider a first-order VAR. Our aim is to estimate u_t^{exp} , which under the assumption that $\text{cov}(u_t^{\text{exp}}, u_t^{\text{fund}}) = 0$, is the expectations shocks of interest. Here u_t^{fund} denotes the structural shock to output. Note that eqs. (1) and (2) describe a structural VAR model, and that $a_{0,12}$ could be obtained by applying a Cholesky decomposition to the variance-covariance matrix of the bivariate reduced-form VAR model (that is, by imposing recursive identification). The recursive identification is justified by the fact that the first release of GDP, Y_t^{t+1} , is published at least two months after professional forecasters are surveyed for their nowcasts $Y_{t|t}$. The expectations shock has a contemporaneous effect on output, but forecasters cannot respond to unobserved output innovations in the same period. This is the identification restriction used by Leduc and Sill (2013). As in the monetary policy analysis of Croushore and Evans (2006), we assume that agents' expectations $Y_{t|t}$ only respond to the first-release available at time t , that is to Y_{t-1}^t (as opposed to, say Y_{t-1}^{16Q4}). We define the output equation (2) in terms

of first-release output. The main advantage of this assumption is that we can identify both u_t^{exp} and u_t^{fund} without the confounding effects of data revisions, as will be clear below. If as Croushore and Evans (2006), the aim is to compute responses for the true value of output growth Y_t , we will need to employ a two-step approach, with the second step described in the next section.

What would happen if we were to estimate (1)-(2) using the latest-vintage data? That is, using Y_t^{16Q4} for $t = 2, \dots, T$, or generally, Y_t^{T+1} , to estimate the equations instead of the first-release output, Y_t^{t+1} (for $t = 2, \dots, T$). To make progress, we need a model of the relationships between the different data vintages. Assume that the period $t + s$ vintage (observed) estimate of the value of Y in period t , denoted Y_t^{t+s} , where $s = 1, \dots, l$, consists of the true value Y_t , as well as news and noise data revisions components, v_t^{t+s} and ω_t^{t+s} , so that

$$Y_t^{t+s} = Y_t + v_t^{t+s} + \omega_t^{t+s}.$$

A revision is news if the initially-released data is an optimal forecast of the revised data, so news revisions are not correlated with the earlier-release, i.e., $Cov(v_t^{t+s}, Y_t^{t+s}) = 0$. Data revisions are noise when each new release of the data is equal to the true value of Y_t , denoted Y_t , plus noise, so that noise revisions are not correlated with the truth, $Cov(\omega_t^{t+s}, y_t) = 0$.⁷ Suppose that the total number of observations T is large enough relative to the number of rounds of revisions considered, $l - 1$, such that the latest vintage available Y_t^{T+1} is a good approximation of true values Y_t . This implies that the vintage- $T + 1$ value incorporates all the l -news revisions terms, and has no measurement error. Then the first-release values are related to the fully-revised ‘true’ values by:

$$Y_t^{t+1} = Y_t^{T+1} + v_t^{t+1} + \omega_t^{t+1}$$

and:

$$Y_{t-1}^t = Y_{t-1}^{T+1} + v_{t-1}^t + \omega_{t-1}^t.$$

The implication of these assumptions is that by estimating eq. (1) with the latest vintage

⁷Jacobs and van Norden (2011) provide a model which generates news and noise revisions with the required properties.

available, we obtain:

$$Y_{t|t} - Y_{t|t-1} = a_{11}(Y_{t-1|t-1} - Y_{t-1|t-2}) + a_{12}Y_{t-1}^{T+1} + \zeta_{1t} \quad (3)$$

$$Y_t^{T+1} = a_{21}(Y_{t-1|t-1} - Y_{t-1|t-2}) + a_{22}Y_{t-1}^{T+1} + a_{0,12}u_t^{\text{exp}} + \zeta_{2t} \quad (4)$$

by substituting for Y_{t-1}^t , where $\zeta_{1t} = u_t^{\text{exp}} + a_{12}(v_{t-1}^t + \omega_{t-1}^t)$ and $\zeta_{2t} = u_t^{\text{fund}} - (v_t^{t+1} + \omega_t^{t+1} - a_{22}(v_{t-1}^t + \omega_{t-1}^t))$. If current-vintage data is used to estimate the model, then the recursively-identified structural shocks will be contaminated by the data revisions. Data revisions have another effect: the OLS estimator may deliver inconsistent estimates of the parameters. Whether or not we can obtain consistent estimates of the unknown parameters in (3) and (4) will depend on whether the data revisions are news or noise. Suppose that data revisions are news. Then the errors ζ_{1t} and ζ_{2t} are correlated with the explanatory variable Y_{t-1}^{T+1} , because $\text{Cov}(Y_{t-1}^{T+1}, v_{t-1}^t) \neq 0$, and the parameter estimates will be inconsistent. However, under the assumption that revisions are noise, the errors and Y_{t-1}^{T+1} are not correlated, because $\text{Cov}(Y_{t-1}^{T+1}, \omega_{t-1}^t) = 0$. News revisions - i.e., data revisions which add new information - will induce a correlation even if most of the observations have been heavily revised. The evidence suggests that early GDP revisions add new information: Croushore and Stark (2003) suggest revisions up to one year can be characterized as news, but thereafter appear to be a mixture. As a consequence, expectations shocks ought to be estimated using eqs. (1) and (2), and not eqs. (3) and (4).

To summarize, the data revisions included in latest-available-vintage data will contaminate the econometrician's estimates of the expectations shocks experienced in real-time by the forecasters. Using a recursive identification approach, we advise use of the first-release of output, instead of the latest-vintage of data, to estimate the reduced-form VAR. As in Leduc and Sill (2013), the VAR can be then used to compute the dynamic transmission of expectations shocks with the caveat that we are disregarding data revisions that might improve the measurement of the variables in the model. In Section 3 we discuss an alternative method to compute responses that considers true values instead of first-releases in eq. (2), following Croushore and Evans (2006).

In the next section we report the results for a real-time VAR employing first releases of real GDP published with about a 30-day delay (so at the end of the first month of the following quarter). This corresponds to estimating equations (1) and (2). Real-time data on GDP is from the Philadelphia Fed Real-Time dataset for Macroeconomists (RTDSM).⁸ In section 3.2 we show how a VAR using

⁸Freely available at <https://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data>

the best available estimates of output can be consistently estimated using an instrumental variables estimator.

2.4 The Real-Time Mixed Frequency VAR

In the real-time VAR we allow forecasters to make use of a wider information set than in the studies which only permit quarterly data. Not using information available to the forecaster to calculate expectations shocks will result in erroneous estimates. We assume forecasters consider monthly variables when updating their forecasts, and use a mixed-frequency VAR, as in Ghysels (2016).⁹

At the time the forecast is surveyed, the forecaster will have access to first-month of the quarter data on stock (SP500 index) returns, SP , inflation π and the short-rate R , as well as first-month of the quarter data on variables closely correlated to Y_t , such as industrial production (IP) and non-farm payroll (NF). Industrial production and payroll employment are key indicators watched by economic commentators, and might be relevant for nowcasting $Y_{t|t}$ (see the literature surveyed by Bańbura *et al.* (2013)). Inflation and the short-rate are key macroeconomic variables observed by professional forecasters because they provide information for current and future changes in monetary policy that may affect GDP growth. We include stock returns as Clements and Galvão (2017) suggest that equity market prices during the month in which the first release of GDP is published carries information that can be used to predict subsequent GDP data revisions.

Two of the monthly series in the information set, industrial production and non-farm payroll, are also subject to data revisions and publication delays. Their first releases are published up to 30 days after their reference month. We include only first releases of these time series, that is, we include $X_{t,m}^{t,m+1}$, where, as before, the subscript value is the reference quarter, month and the superscript indicates the publication date that in this case is a month later. We use monthly growth rates, computed as $X_{t,m}^{t,m+1} = 100(\log(Z_{t,m}^{t,m+1}) - \log(Z_{t,m-1}^{t,m+1}))$, where Z are the values in levels.

We apply the strategy of using GDP first releases, Y_t^{t+1} , and the recursive identification scheme discussed in section 2.2 to the mixed-frequency VAR. The ordering is given by the time-line of the release of the different variables, their vintages, and the timing of the SPF questionnaire. As explained before, five monthly series are included. We set the vector of monthly variables at quarter t and month m as:

$$\mathbf{x}_{t,m}^{t,m+1} = [SP_{t,m}, R_{t,m}, IP_{t,m}^{t,m+1}, NP_{t,m}^{t,m+1}, \pi_{t,m}]'$$

⁹In principle one could use higher frequency data, as in Ghysels and Wright (2009), for example.

Stock returns, $SP_{t,m}$, and the short-term rate $R_{t,m}$ are ordered before the other monthly series because they are observed with no publication delay. Only $IP_{t,m}$ and $NP_{t,m}$ are subject to data revisions so the notation incorporates the fact we use first releases for these variables. We evaluate the impact on our empirical results from including additional financial variables in section 3.

We include both the time series of nowcasts $Y_{t|t}$ and of one-quarter-ahead forecasts $Y_{t+1|t}$ in the VAR. This allow us to identify shocks to forecast revisions. Empirically, the specification with $Y_{t|t}$ and $Y_{t+1|t}$ included separately, as here, is slightly better than including the variables as $Y_{t|t} - Y_{t|t-1}$, i.e., restricted to appear as an update as in eqs. (1) and (2), although qualitatively similar results are obtained.

The first month of the current quarter information on the monthly variables $\mathbf{x}_{t,1}^{t,2}$ is available before the forecasts $Y_{t|t}$ and $Y_{t+1|t}$ are made. But the following months $\mathbf{x}_{t,2}^{t,3}$ and $\mathbf{x}_{t,3}^{t+1,1}$ are available later. As consequence, we order $\mathbf{x}_{t,1}^{t,2}$ before the forecasts, but $\mathbf{x}_{t,2}^{t,3}$ and $\mathbf{x}_{t,3}^{t+1,1}$ are included after.

The first release of GDP, Y_t^{t+1} is published at the end of the first month of the next quarter, so it is available after $\mathbf{x}_{t,3}^{t+1,1}$.

The 18×1 vector of endogenous variables of the stacked mixed-frequency VAR is then set as:

$$\mathbf{y}_t = [\mathbf{x}_{t,1}^{t,2'}, Y_{t|t}, Y_{t+1|t}, \mathbf{x}_{t,2}^{t,3'}, \mathbf{x}_{t,3}^{t+1,1'}, Y_t^{t+1}]'. \quad (5)$$

The vector of reduced form innovations is ε_t , and if A_0 is the lower triangular matrix from the Cholesky decomposition of $\text{var}(\varepsilon_t) = \Sigma$, then the structural shocks are $u_t = A_0^{-1} \varepsilon_t$. The expectations shock is the sixth structural shock:

$$u_{t|t}^{\text{exp}} = Y_{t|t} - E[Y_{t|t} | Y_{t|t-1}, \mathbf{x}_{t,1}^{t,2'}, \mathbf{x}_{t-1,2}^{t-1,3'}, \mathbf{x}_{t-1,3}^{t,1'}, Y_{t-1}^t, \dots].$$

Our standard recursive identification scheme means that the expectations shock may only contemporaneously affect variables lower down in the ordering given in (5), such as the advance estimate of Y_t , Y_t^{t+1} , but not variables above it. By using only first releases of all variables subject to revisions, we are avoiding the the ‘look-forward’ bias affecting both shock identification and transmission of the shocks described in section 2.2.

As described above our expectations shocks are based on forecast revisions observed in the nowcast $Y_{t|t}$, but our identification strategy could be also applied for $n = 1, 2, 3$, and we show results for $n = 3$ in the next section. Our emphasis on nowcasts is due to the fact that they are updated more as new data is released or as monetary policy (for example) is changed, relative to

longer-horizon forecasts. This is natural, because the value of forecast origin information diminishes as the horizon increases, for stationary growth rates. This means that we might be able to more accurately determine the short-horizon expectations shocks, such as $u_{t|t}^{\text{exp}}$.

2.5 Empirical Estimates of Expectations Shocks

To deal with parameter uncertainty arising from the large number of parameters in the 18-variable real-time VAR, we use the Minnesota prior with hyper-parameters estimated by the Bayesian VAR MCMC estimation approach proposed by Giannone, Lenza and Primiceri (2015). We set $p = 5$, while prior hyper-parameters are estimated in a MCMC algorithm calibrated to accept around 40% of the candidate draws.¹⁰ The large number of parameters arise because we want to ensure the expectations shocks are ‘shocks’, and do not include components known to the forecasters but wrongly missing from our VAR.

Figures 2 give our estimates of the expectations shock, $u_{t|t}^{\text{exp}}$, based on the real-time mixed-frequency VAR estimates at the posterior mean for the sample period from 1968Q4 to 2016Q3. We also draw in the forecast updates (as shown in Figure 1) for comparison purposes. The correlation between the forecasts updates and expectation shocks is 53%, suggesting that the variables included in the mixed-frequency VAR explain a half of the variation in the nowcast updates.

2.5.1 The Role of Monthly Data

How important is it to use a mixed-frequency VAR that includes monthly data rather than just quarterly indicators? We estimate a quarterly latest vintage VAR as $\mathbf{y}_t^{\text{final},q} = [Y_{t|t}, Y_{t+1|t}, \mathbf{x}_t^{16M12'}, Y_t^{16Q4}]$ and use the recursive identification to obtain $u_{t|t}^{\text{exp,final,quart}}$, graphed in Figure 2. It is clear the information in the model is then only able to explain a smaller share of the variation of the forecast updates. Indeed only 1/4 of the forecast update variation is then explained, as the shocks have a correlation of 76% with forecast updates.

2.5.2 The Role of Real-time Data

What happens if we use latest-vintage data? We re-estimate the mixed frequency VAR in (5) using Y_t^{16Q4} instead of Y_t^{t+1} , and latest-vintage values for $IP_{t,m}^{16M12}$ and $NP_{t,m}^{16M12}$. Then using parameters at the posterior mean, we compute $u_{t|t}^{\text{exp,final}}$ using the recursive identification scheme

¹⁰We thank Giorgio Primiceri for making the code for Giannone *et al.* (2015) available from his website. We also thank Danilo Cascaldi-Garcia for sharing his code to identify news shocks.

as above. The correlation between $u_{t|t}^{\text{exp},\text{final}}$ and $u_{t|t}^{\text{exp}}$ is 84% and the correlation of $u_{t|t}^{\text{exp},\text{final}}$ with forecast updates is 54%. So the use of real-time data does affect the estimation of the expectation shocks.

A final exercise is to compare the responses we obtain using the real-time data approach, with those from estimating the model on the latest-available vintage. We present responses for the three key endogenous variables: output, prices and the short-rate. Because the models are estimated with monthly inflation and the short-rate, we present responses for the last month in the quarter (note that the shocks have zero impact only on the first month in the quarter). The responses presented in Figures 3 and 4 are accumulated responses for output and prices. For the real-time mixed-frequency VAR in Figure 3, we show responses for first-release GDP values. In Figure 4 the responses are for the latest vintage GDP values.

The results in Figure 3 clearly indicate that expectation shocks have positive significant effects on output, prices and the short-rate at short and medium horizons. In contrast, if we use the mixed-frequency VAR with values that were not available at the time the forecasts are made (because they include future revisions), we only find a positive, significant effect at the one-quarter horizon, as indicated in Figure 4. There is a clear attenuation effect in the dynamic transmission of the shocks by including future data revisions when we estimate the dynamic transmission of the shock.

A problem with measuring the impact of expectations shocks with the real-time mixed frequency VAR as in Figure 3 is that we are not using the ‘best’ estimates of GDP and other key macroeconomic variables, as is standard practice in empirical business cycle analysis. A comparison between Figures 3 and 4 is also problematic because we are estimating responses for different values of output. We solve these issues in section 3 by using a Macroeconomist’s VAR, and proposing two methods to estimate the dynamic responses depending on whether true values are assumed to be observed or not.

2.6 Understanding Expectations Shocks

We define expectations shocks as updates in the consensus forecasts for GDP growth that are not related to new information on macroeconomic fundamentals. In this section we map our empirical estimates of these shocks to the business cycle, and to economic developments more generally. In table 2 we list the dates of the ‘large’ expectation shocks, where large means shocks of two standard deviations. The time series of shocks is plotted in Figure 2 from 1968.

We first investigate how the shocks are related to business cycle phases (the second column

of Table 2) and record their signs. Three out of the eight shocks occur during recessions. Given the preponderance of expansionary periods during this period, large expectations shocks are more likely to occur during recessions. Two occur at times of turning points: at the onset of recession in 1980Q1, and the end of a recession in 2001Q4. There is also a large expectations shock in 1984Q2 as a high growth recovery phase comes to an end.

We consider the possibility that the expectations shocks are linked to other shocks in the literature, and in particular the monetary policy shocks identified by Romer and Romer (2004). The large negative shock in 1979Q3 occurs at a time of change in monetary policy regime. More recently, the large negative shock in 2009Q1 coincides with the major financial crisis. Nevertheless, there is generally only limited overlap, as might be expected, since our shocks ought to be unrelated to inflation and interest rates shocks observed in the first month of the quarter.

We find large positive expectations shocks in 1978Q2 and 2006Q1, corresponding to large increases in the growth rate of GDP (as described in Table 2). These large unexplained forecast updates may reflect changes in the consensus forecasts in anticipation of the publication of strong growth rate figures that have not been yet reflected in the monthly economic data included in the VAR.

Our expectations shocks are calculated from consensus expectations, but might be correlated with cross-sectional forecaster disagreement, which is correlated with uncertainty, which is a counter-cyclical variable.¹¹ If we measure disagreement using the interquartile range (IQR, published by the Philadelphia Fed) for the current-quarter forecasts (the nowcasts), the correlation between disagreement and our expectations shocks is negligible. But the correlation between the IQR and the squared the shocks is 40%. Large expectations shocks are more likely in periods of high disagreement. Figure 5 presents the squared expectations shocks. The Figure also includes the predictions for the expectations shocks based on projecting the series of shocks on the IQR series. The IQR series explains about 20% of the variation in the squared shock series. An inspection of Figure 5 suggests that expectations shocks are more likely when forecasters' beliefs are more heterogenous, which usually happens during recessions. Heterogeneity of beliefs might be caused by higher-order beliefs, the main driver of the transmission of confidence shocks in Angeletos *et al.* (2018).

¹¹See, e.g., Zarnowitz and Lambros (1987) and Giordani and Söderlind (2003) on disagreement and uncertainty, and Bloom (2009), Bachmann, Elstner and Sims (2013), Baker, Bloom and Davis (2016) and Girardi and Reuter (2017), *inter alia*, on the impact of uncertainty shocks on economic activity.

3 The Dynamic Effects of Expectations Shocks

We calculated expectations shocks from a ‘real-time VAR’ estimated using only data which would have been available in real-time to professional forecasters and market participants, and using an identification scheme based on the timing of events: data releases and the filling in of survey questionnaires. However, macroeconomists are often interested in measuring the effect of shocks on the ‘true values’ of a set of key macroeconomic variables of interest. These true values are given by the latest data available at the time of the study, instead of the estimates that would have been available in real-time each time a survey return was made.

We follow the tradition of calculating responses to previously estimated structural shocks using the latest-available data. The recent literature suggests we may calculate responses to the shock from a series of suitably defined local projections, following Jorda (2005), or we could include the shock in a second VAR, as an observed variable, and obtain the impulse responses by iteration in the usual way: see Plagborg-Møller and Wolf. (2018) and Stock and Watson (2018).¹²

We proceed by considering the expectations shocks in Figures 2 as observed time series, and add them to the Macroeconomist’s VAR to measure the effects of these shocks on the variables of interest.

3.1 The Macroeconomist’s VAR

Our Macroeconomist’s VAR model uses a set of variables that commonly feature in analyses of the responses to belief-based shocks, as in Leduc and Sill (2013), Barsky and Sims (2012), Levchenko and Pandalai-Nayar (2017) and Fève and Guay (2016). These include utilization-adjusted Total Factor Productivity computed by Fernald (2014), real GDP, aggregate consumption and investment, total weekly hours, the CPI and the 3-month TBill rate. Data on investment, consumption, GDP and hours are the 2016Q4 vintage values from the Philadelphia Fed real-time dataset. The vector of variables is given by $\mathbf{x}_t = [TFP_t, Inv_t, Const_t, GDP_t, H_t, CPI_t, R_t]'$.¹³

An advantage of employing the Macroeconomist’s VAR to measure the effects of the expectations

¹²The former establishes the equivalence between local projections and VARs, setting aside small-sample considerations. The latter discusses the two methods when there are ‘external’ or ‘proxy’ variables that aid the identification of the structural shocks. A key requirement to calculate the structural shocks is that the VAR is invertible, that is, the VAR reduced form shocks and the structural shocks must span the same space. Then it is possible in principle to calculate the structural shocks from the variables which comprise the VAR (see, e.g., the discussion in Stock and Watson (2018), Ramey (2016) or Kilian and Lütkepohl (2017)).

¹³This matches that used by Forni, Gambetti and Sala (2019), in their evaluation of the fundamentalness or invertibility of the Structural VAR.

shocks as described above is that we are able to compute forecast error variance decompositions, that is, we are able to measure how much of the business cycle variation of this set of key macroeconomic variables is explained by expectation shocks.

Following the literature on measuring the effects of macroeconomic shocks surveyed in Ramey (2016), we employ a VAR in log-levels to allow the possibility of common trend components across the macroeconomic variables. An exception is the short-rate that is taken in levels.

3.2 Data Revisions and the Macroeconomist’s VAR

Our aim is to show how to compute responses to expectations shocks assuming that the macroeconomist’s VAR is estimated with true values for all macroeconomic variables subject to data revisions. Let \mathbf{y}_t denote a vector of the true values of the macroeconomic variables included in the VAR. Assume also, for simplicity, that all the m variables in the vector \mathbf{y}_t are subject to revision, the VAR order is 1, and we ignore intercepts. The VAR is then:

$$\mathbf{y}_t = \Phi \mathbf{y}_{t-1} + \mathbf{C}_0 \hat{u}_{t|t}^{\text{exp}} + \boldsymbol{\varepsilon}_t, \quad (6)$$

where $\hat{u}_{t|t}^{\text{exp}}$ is estimated consistently using the real-time VAR. Note here we recognize that the impact of $\hat{u}_{t|t}^{\text{exp}}$ on the true value of the macroeconomic variables may differ from the impact on the first release.

If we are able to assume that true values are eventually observed and are equal to ‘heavily revised’ values, we are able to estimate the eq. (6) as follows. Because national accounts data are typically subject to three rounds of annual revisions, we can shorten the sample by removing the last three years or so of observations (as Croushore and Evans (2006)), such that all the values included in the analysis have been ‘heavily revised’. Let \mathbf{y}_t^{T+1} denote the vintage T -vector of variables relating to observation t , then the VAR is estimated for observations $t = p + 1, \dots, T - l + 1$, assuming that $l = 16$. Empirically, we use the 2016Q4 vintage of data for the variables Inv_t , $Cons_t$, GDP_t , H_t (which includes data through 2016Q3), and Fernald’s 2017-M3 vintage for TFP_t with data up to 2016Q4.¹⁴ This implies that by assuming that true values are well approximated by ‘heavily revised’ values, we can use observations up to 2012Q4 from the 2016Q4 vintage to estimate the VAR.

¹⁴The reference is Fernald (2014), with updates to the series available from <https://www.frbsf.org/economic-research/economists/john-fernalld/>

Assume now that true values of the variables subject to revision are never observed, but that revisions after the first 3 annual revisions are noise, that is, they just remove earlier measurement error, such that, $y_{i,t}^{T+1} = y_{i,t} + \omega_{i,t}^{T+1}$ for $t = 1, \dots, T-l+1$ and $i = 1, \dots, m$, and $l = 16$. Each variable in the vector \mathbf{y}_t is observed with an independent measurement error, compatible with pure noise revisions.

If we use observations up to $T-l+1$ from the $T+1$ vintage to estimate eq. (6), then we have:

$$\mathbf{y}_t^{T+1} = \Phi \mathbf{y}_{t-1}^{T+1} + \mathbf{C}_0 \hat{u}_{t|t}^{\text{exp}} + \boldsymbol{\varepsilon}_t + \boldsymbol{\omega}_t^{T+1} - \Phi \boldsymbol{\omega}_{t-1}^{T+1}, \quad (7)$$

where $\boldsymbol{\omega}_t^{T+1} = (\omega_{1,t}^{T+1}, \dots, \omega_{m,t}^{T+1})$. Because the regressors are correlated with the disturbances, OLS will deliver inconsistent estimates. Notice that $\boldsymbol{\omega}_t^{T+1} = \mathbf{0}$ and $\boldsymbol{\omega}_{t-1}^{T+1} = \mathbf{0}$ correspond to the true values being observed, $y_{i,t}^{T+1} = y_{i,t}$ for $t = 1, \dots, T-l+1$, and similarly $y_{i,t-1}^{T+1} = y_{i,t-1}$. As already noted, we then revert to (6) which is consistently estimated on the data \mathbf{y}_t^{T+1} , $t = p+1, \dots, T-l+1$.

If we prefer not to make this assumption, we can use an instrumental variable estimator instead, following Croushore and Evans (2006). Specifically, we use a vintage for observations $t = 1, \dots, T-l+1$ published $l-1$ quarters earlier. The first-stage regression is:

$$y_{t,i}^{T+1} = \beta_{0,i} + \beta_{1,i}t + \beta_{2,i}t^2 + \beta_{3,i}y_{t,i}^{T-l+2} + \tilde{\omega}_{i,t}^{T+1} \text{ for } t = 1, \dots, T-l+1 \text{ and } i = 1, \dots, m \quad (8)$$

where $y_{t,i}^{T-l+2}$ is the instrument for $y_{t,i}^{T+1}$, and variables are in log-levels. Instrument relevance is (virtually) assured, as both the $T-l+2$ and $T+1$ -vintages are estimates of the same true values.

Moreover, when the series are in log-levels and potentially integrated of order one, as in our application, we would expect the different vintage estimates to be cointegrated. The inclusion of the quadratic trend in the first stage regression in (8) accommodates the fact that the effects of changes in the base year between $T-l+2$ and $T+1$ may vary across observations.¹⁵ Instrument validity requires the exclusion restriction:

$$E \left[\left(\varepsilon_{i,t} + \omega_{i,t}^{T+1} - \phi_i \omega_{i,t-1}^{T+1} \right) \left(y_{i,t-1} + \omega_{i,t-1}^{T-l+2} + v_{i,t-1}^{T-l+2} \right) \right] = 0$$

where the first term is the error in (7) and the second the instrument. The instrument includes a news revision term $v_{i,t-1}^{T-l+2}$, because we allow the l -most recent observations to still be subject

¹⁵Empirically, we set $l = 16$, and there is one change of base year for the chain linking deflator between the 2013Q1 and the 2016Q4 vintages.

to news and noise revisions. The exclusion restriction holds because the time series $y_{t,i}^{T+1}$ for $t = 1, \dots, T - l + 1$ is assumed not to include observations subject to news revisions, since $y_{i,t}^{T+1} = y_{i,t} + \omega_{i,t}^{T+1}$. In addition, the restriction holds because the expectations of all the products are zero: $E\left(\omega_t^{T+1}\omega_{t-1}^{T-l}\right) = E\left(\omega_{t-1}^{T+1}\omega_{t-1}^{T-l}\right) = 0$, ω and y are uncorrelated, and $\varepsilon_{i,t}$ is assumed uncorrelated with ω and y .

Using $\hat{y}_{i,t}$ as the fitted value of the regression in (8) computed for each i , and then added to vector $\hat{\mathbf{y}}_t$, we measure the transmission of expectations shocks on the true unobserved values using the following second-stage system of equations:

$$\mathbf{y}_t = \Phi\hat{\mathbf{y}}_{t-1} + \mathbf{C}_0\hat{u}_{t|t}^{\text{exp}} + \varepsilon_{IV,t}. \quad (9)$$

3.3 Empirical Implementation

In practice, not all variables in \mathbf{x}_t are subject to revisions. To implement the instrumental variable approach described earlier, we instrument the first five variables such that:

$$\hat{\mathbf{x}}_t = [\widehat{TFP}_t, \widehat{Inv}_t, \widehat{Cons}_t, \widehat{GDP}_t, \widehat{H}_t, CPI_t, R_t]'$$

For GDP, consumption, investment and hours, we replace the 2016Q4 vintage data with the predictions from a regression as in (8), where we use data from the 2013Q1 vintage to obtain predictions for values in the 2016Q4 vintage for observations up to 2012Q4. For TFP, we use as an instrument the 2013M3 vintage, and as before we use observations up to 2012Q4.

As our focus is the dynamics effects of $\hat{u}_{t|t}^{\text{exp}}$, we include $\hat{u}_{t|t}^{\text{exp}}$ as the first variable in the VAR, that is, $\mathbf{z}_t = [\hat{u}_{t|t}^{\text{exp}}, \hat{\mathbf{x}}_t']$. Define also the vector using instead the first-stage regression fitted values, that is, $\hat{\mathbf{z}}_t = [\hat{u}_{t|t}^{\text{exp}}, \hat{\mathbf{x}}_t']$. The system to be estimated is:

$$\mathbf{z}_t = \mathbf{c} + \sum_{\tau=1}^p \mathbf{A}_\tau \hat{\mathbf{z}}_{t-\tau} + \mathbf{v}_t, \quad (10)$$

with $\mathbf{v}_t \sim N(0, \Sigma_v)$.

The identification assumption required for recursive identification is that innovations to the expectations shocks are not contemporaneously correlated with the shocks to the macroeconomic variables. As $\hat{u}_{t|t}^{\text{exp}}$ is by definition orthogonal to the shocks affecting the first release of output, we can assume it is orthogonal to the shocks to the subsequently revised values of the macrovariables

since these refer to information not available at the time expectations were formed. As a consequence, we can obtain estimates of \mathbf{C}_0 in eq. (9) using the Cholesky decomposition applied to Σ_v .

We set $p = 5$, and adopt both the Minnesota prior and the ‘dummy-initial-observation’ prior (using values in \mathbf{z}_t to obtain initial values), while prior hyper-parameters are estimated in a MCMC algorithm calibrated to accept around 40% of the candidate draws as in Giannone *et al.* (2015). We compute impulse responses and forecast error variance decompositions using 20,000 draws from the posterior distributions.

3.4 Empirical Results

Figures 6 and 7 show impulse responses using the instrumented Macroeconomist’s VAR, and Table 3 shows the forecast error variance decomposition of output. In these tables and figures we present the mean values and also the 68% bands.

Figure 6 shows the responses of output to expectations shocks, $\hat{u}_{t|t}^{\text{exp}}$, measured in three different ways. The first one reproduces the results for the real-time VAR in Figure 3, so these are responses of the first-release of GDP. The other two are obtained using the Macroeconomist’s VAR. The lines in black are based on the estimation of the VAR with data from the 2016Q4 vintage up to 2012Q4 (and TFP data from the 2017M3 vintage), that is, a VAR for $\mathbf{z}_t = [\hat{u}_{t|t}^{\text{exp}}, \mathbf{x}'_t]$. These estimates assume we are interested in computing responses to true values, and we are willing to suppose true values of the data are given in the 2016Q4 vintage for observations that have undergone at least 3 rounds of annual revisions. The lines in grey show the responses computed using eq. (10), that is, it considers the case that the true values are not observed.

An inspection of Figure 6 suggests that at the impact of the shock, responses of the true value of output are significantly higher than the ones obtained for first-release output. Short-run effects are in general very similar across specifications, but there are differences in how fast the effect of the expectations shock dies out. If we assume that true values are not observed because the national account estimates are always contaminated by non-transitory measurement errors due to sampling uncertainty and other conceptual assumptions, then expectations shocks have significant effects on aggregate output only up to 3 years.

Table 3 shows the share of the variation of output explained by expectations shocks, computed with the same set of candidate methods to measure the transmission of expectations shocks. The effect of expectations shocks at impact varies from 4% to 6%, with larger values obtained using

first-releases. The proportions at medium horizons ($h = 8$) vary more widely across specifications from 5% to 12%. These results suggest that as the statistical office incorporates new information in the estimates of the national accounts data, the relative importance of the expectations shocks in explaining business cycle variation in output increases.

Figure 7 shows the responses of TFP, Investment, Consumption and Hours to expectations shocks computed using the instrumented Macroeconomist VAR. For comparison we also include estimates using the Macroeconomist VAR. The arguments in section 3.2 suggest using the instrumented macroeconomist VAR estimates unless we are willing to assume the latest-vintage values can be treated as true values, i.e., that they are free of noise. Recall that whether we use OLS or IV, we need the maintained assumption that the estimation sample is truncated at $T - l + 1$ in order that observations subject to news revisions are excluded.

Measurement errors are relatively larger for investment and TFP as suggested by the size of the deviations between black and grey lines in Figure 7 and supported by the fact that the estimated variance of the measurement error computed using eq.(8) is almost 10 time larger for these variables in comparison with consumption and output. Interestingly, the effect of using the instrumented Macroeconomist’s VAR is observed even for short horizons when evaluating responses for TFP. The effects of TFP measurement errors on estimates of the transmission of shocks using VAR models has been studied by Cascaldi-Garcia (2017) and Kurmann and Sims (2017), amongst others.

By removing the impact of measurement errors, we find that the effects of expectations shocks at horizons longer than two years are attenuated for all variables in Figure 7, delivering a clearer evidence of short-term effects from expectations shocks. If we compute the responses without instrumenting TFP, but including instruments for the four activity measures, we find a smaller long horizon attenuation since the use of predicted values from the first-stage regression visibly deteriorates the fit of the TFP equation in the VAR model (twofold increase in mean squared errors), but this deterioration is very limited in the VAR equations of the other macroeconomic variables.

The next section compares expectations shocks with alternatives in the literature, to determine whether expectations shocks constitute a new source of business cycle variation.

3.5 Comparison with baseline identification strategy

At this point, we consider the consequences of two key elements of our approach for the understanding of the effects of expectations shocks on the macroeconomy. We bring out the effects of allowing

for data uncertainty, and the real-time nature of the forecasters information set, by comparing the responses obtained using our two-step approach to a benchmark specification. The baseline strategy simply includes the forecasts updates $Y_{t|t} - Y_{t|t-1}$ as the first variable in the macroeconomist VAR, and identifies expectations shocks using a recursive approach. Red lines in Figure 8 indicate the posterior mean responses of applying this strategy with data up to 2012Q4 from the 2016Q4 vintage. Black lines describe the responses using our two-step strategy, including 68% bands.

The empirical results in Figure 8 suggest that the responses of macroeconomic variables are attenuated in the baseline approach. That is, the cost of disregarding data uncertainty and the real-time nature of the information set is to clearly under-estimate the responses of variables such as consumption and investment. Instead of significant positive effects up to two years after the shock, we find effects only over the first year. In contrast, responses of prices and interest rates are significantly positive, matching the findings of Leduc and Sill for expectations shocks calculated for unemployment. The baseline strategy suggests characterizing expectations shocks as short-term demand shocks, inducing monetary policy reactions that are stronger than the ones obtained with the two-step strategy. We discuss further the economic implications of our empirical results in Section 4.

3.6 Robustness Checks: Additional Financial Variables and Longer-Horizon Expectations

In this section, we evaluate the effects of two choices we made when we specified the real-time mixed-frequency VAR in Section 2.4. We consider whether our findings (described in Figures 6 and 7) are robust to those choices by considering reasonable alternatives.

The first choice regards the monthly variables included in the real-time mixed-frequency VAR. Because expectations shocks are changes in expectations not related to fundamentals updates, assumptions about the information set employed in the estimation of the shock may matter. Our results so far use five monthly variables: two key monthly measures of economic activity, one financial expectational variable (stock returns) and two important macroeconomic variables (inflation and the short-rate) that agents observe (including their shocks) in real-time. This would appear to cover some of the main sources of information a forecaster would likely draw on. Nevertheless, we check whether the results are robust to the inclusion of six additional monthly variables.

The second choice regards the use of nowcast updates ($n = 0$) instead of updates for longer horizon expectations ($n = 2$). We re-estimate the real-time mixed frequency VAR in eq. (5) using

$Y_{t+2|t}$ and $Y_{t+3|t}$ instead of $Y_{t|t}$ and $Y_{t+1|t}$.

The black lines in Figure 9 describe responses with the same specification as in Figures 6 and 7. The red lines are the posterior mean responses when expectations shocks are computed using 11 monthly variables in the mixed frequency VAR. None of the additional variables are subject to revision. They include one additional economic activity variable (unemployment), oil prices and four additional financial variables (long-term (10-year) rate, corporate bond spread, US/CAN exchange rate and stock market volatility). The estimated expectations shocks (which is then also orthogonal to shocks observed in the first month of the quarter for all these variables) has a correlation of 45% with the original expectations updates. Although we have added more variables, this compares with a correlation of 53% between the original shocks (in Figure 2) and the expectations updates, suggesting a relatively small change. As shown in Figure 9 this is indeed the case, because the responses using this alternative estimate of expectations shocks are all within the 68% bands of our benchmark estimate. As a consequence, we conclude that the set of five monthly variables included in our original model appear to provide a reasonable approximation to the forecasters information set, because expanding this information set leaves the results largely unchanged.

Figure 10 shows a comparison between our benchmark responses (black lines) with the posterior mean responses obtained for an alternative estimate of expectations shocks that uses $Y_{t+2|t}$ and $Y_{t+3|t}$ in eq.(5). The responses obtained with expectations shocks estimated on longer-horizon forecasts are similar to the ones obtained in the central case using updates to nowcasts. All red lines are within the 68% bands. Qualitatively, responses of prices and inflation show more significant negative responses of these variables, suggesting that long-term expectations shocks effects resemble the effects of supply instead of demand shocks.

4 Comparison of Expectations shocks and Alternative Belief-Based Shocks

As described in the introduction, the expectations shocks may be related to news about future fundamentals, such as anticipated technology shocks, or to changes in consumer confidence. Of interest is the extent to which the shock is an alternative measurement of extant shocks identified in the literature, as opposed to complementing the literature with an additional source of business cycle variation. To address this question, we compare our estimated shock with the technology news shocks of Barsky and Sims (2011) and the confidence shocks of Barsky and Sims (2012).

4.1 Measuring Technological News and Consumption Confidence Shocks

News shocks in Beaudry and Portier (2006) and Barsky and Sims (2011) are future technological changes that are anticipated today. Following the news shocks literature surveyed in Beaudry and Portier (2014), we use two forward-looking variables to identify news shocks. The first one is an equity market index, namely, the S&P500 index. The second is the confidence variable, ‘E5Y’, from the Michigan survey, used as a measure of consumer confidence by Barsky and Sims (2012), that we label $Conf_t$. The measure of total factor productivity is the utilization-adjusted series of Fernald (2014), from the latest available vintage (2017M3), TFP_t^{17M3} .

The vector of endogenous variables in the VAR is $\mathbf{z}_t = [TFP_t^{17M3}, SP500_t, Conf_t, Y_t^{16Q4}, P_t, R_t]'$. We estimate a VAR in log-levels to allow that the variables may exhibit cointegration, as in Beaudry and Portier (2006) and Barsky and Sims (2011). The news shocks are identified by the twin requirements that i) they maximize the forecast error variance decomposition of TFP after 40 quarters, and ii) they have a zero effect on TFP at impact (i.e., only affect future values). This is the identification scheme proposed by Barsky and Sims (2011), which also allow us to compute an unexpected technology shock (or ‘surprise’ technology shocks) which are allowed to have an impact effect on technology. Because the identification scheme does not identify the sign of the shock, the restriction that the impact effect of news shocks on $SP500$ is non-negative is also imposed. We compute the time series of news shocks at the posterior mean using a VAR with five lags estimated by Bayesian methods with quarterly data from 1968 to 2016.

Barsky and Sims (2012) use a three variable VAR: $\mathbf{z}_t = [C_t^{16Q4}, Y_t^{16Q4}, Conf_t]'$ (where C_t^{16Q4} refers to aggregate consumption in the latest available vintage, 2016Q4 in our case) to calculate a confidence shock. Their preferred recursive identification places confidence last in the VAR, implying that confidence shocks have no impact effect on the two macroeconomic aggregates. As before, we estimate a VAR in log-levels with five lags using the MCMC algorithm, described in Giannone *et al.* (2015). We computed the time series of confidence shocks at the posterior mean, using a Cholesky decomposition to compute the historical shocks.

4.2 Correlations of Shocks

Table 4 presents correlation estimates between technological news and confidence shocks and expectations shocks. Expectations shocks are significantly correlated with the consumer confidence shock. However, the correlations are no larger than 20%, suggesting that expectations and confi-

dence shocks have distinct elements. Neither can the expectation shocks be regarded as essentially a technology news shock.

These results support the claim that expectations shocks may be an alternative measure of the confidence shocks considered by Barsky and Sims (2011) and Angeletos *et al.* (2018), since they are significantly (although mildly) correlated with consumer confidence.

4.3 The Macroeconomic Effects of Expectations Shocks

To better understand the nature of the expectations shocks relative to the technological news and confidence shocks, we calculate a time series of expectations shocks $u_{t|t}^{\text{exp}*}$ that is orthogonal to the news and confidence shocks. We do this by regressing the expectations shock $u_{t|t}^{\text{exp}}$ on the news and confidence shocks to generate the orthogonalized shock $u_{t|t}^{\text{exp}*}$. The analysis of the correlations between the shocks in section 4.2 indicates the expectations shock is distinct from either the news shock or consumer sentiment, but it is informative to consider the responses of the macro-variables to the component of the expectations shock not encompassed by the other two shocks.¹⁶

Table 5 shows estimates for the variance decomposition of five macroeconomic variables using either the expectation shocks, $u_{t|t}^{\text{exp}}$, or the expectation shocks that is orthogonalized to confidence and news shocks $u_{t|t}^{\text{exp}*}$ using the instrumented Macroeconomist’s VAR, which uses data up to 2012Q4. The results in Table 5 suggest that the orthogonalization has only a small effect on the explanatory power of the expectations shocks for business cycle variation. Interesting, we find evidence that $u_{t|t}^{\text{exp}*}$ can explain a larger variation of hours at short horizons (2% to 4%) than $u_{t|t}^{\text{exp}}$. The contribution of the orthogonalized expectations shocks to business cycle variation at the two-year horizon is of 8% for output, 6% for consumption, but 8% for investment and 10% for hours. Expectations shocks also contribute to explain 4% of the variation in TFP at long horizons ($h = 40$).

In summary, the expectations shock is a largely complementary source of business cycle variation. Changes in agents’ expectations that are unrelated to a representative monthly set of economic fundamentals, to consumer confidence shocks, and to technological news shocks still account for business cycle variation in key economic variables.

Figure 11 shows the responses to the $u_{t|t}^{\text{exp}*}$ shocks. These responses are for horizons up to 10

¹⁶Often in structural VAR analysis the shocks are calculated within one over-arching model so that the shocks are mutually orthogonal by construction. This might be possible here, but it seems simpler to estimate the news and sentiment shocks as originally proposed, and then perform the orthogonalization in an additional step.

years (40 quarters) to observe long-run effects. It is clear that all long-run effects are statistically zero. In the short-run, however, we find that expectations shocks lead to a significant positive comovement in aggregate output, consumption, investment and hours. There is a temporary decline in TFP, and a temporary increase in CPI and the short-rate.

If we compare the empirical responses in Figure 11 with the ones computed for confidence shocks by Angeletos *et al.* (2018), we find qualitatively similar effects, since we have evidence of positive significant comovement in the key activity variables, and the effects on TFP, inflation and the short-rate, are very small. However, consistent with the evidence in Fève and Guay (2016) for sentiment shocks, we find the expectations shocks explain only a small part of the business cycle variation (around 10% based on the estimates in Table 4) instead of the 40-60% reported by Angeletos *et al.* (2018). Because we consider only one type of news shocks (news about future technological changes), the dynamic transmission of expectations shocks described in Figure 11 may also capture the effects of news shocks about other fundamentals, such as those of Miyamoto and Nguyen (2019).

5 Conclusions

We propose a two-step approach to measure the effects of expectations shocks on the macroeconomy. In the first step, we approximate the information set available to the agents forming expectations with a mixed-frequency VAR, estimated using first-release values of the variables subject to data revisions. We show that our approach can consistently estimate expectations shocks identified as being unrelated to fundamental changes. If instead we employ the latest-available data vintage to estimate the mixed-frequency VAR, data revisions will contaminate the expectations shock estimates. The second step of our approach estimates the responses of ‘true values’ of the macroeconomic variables. We show how to estimate responses under either of two assumptions about the true values: (i) they are eventually observed after many rounds of revisions or (ii) they are never observed because statistical office estimates of macroeconomic variables always include measurement error. If (i) holds, one can estimate the VAR using ‘heavily revised’ data taken from the latest vintage, and a shortened sample. Under (ii) an instrumental variable approach as in Croushore and Evans (2006) is required.

The dynamic responses characterize the business cycle comovement (in output, consumption, investment and hours) that is triggered by an expectation shocks. The responses of these key

aggregates are positive and significant at horizons up to two years, while the responses of TFP, prices and the short-rate are small and often insignificant. This pattern is similar to the responses to confidence shocks of Angeletos *et al.* (2018), but differs from the expectations shocks responses in Leduc and Sill (2013). We show that expectations shocks contain information that goes beyond confidence and technology news shocks.

In summary, we provide empirical evidence that changes in expectations that are unrelated to current and past fundamentals, but may be related to sentiment, confidence and non-technological news, have a role in explaining business cycle variation. The identification and estimation strategy behind these empirical results rely on a new methodology that takes into account the effects of data uncertainty and the real-time nature and scope of the agents' information set.

References

- Angeletos, G., Collard, F., and Dellas, H. (2018). Quantifying Confidence. *Econometrica*, 86(5), 1689–1726.
- Aruoba, S. B. (2008). Data revisions are not well-behaved. *Journal of Money, Credit and Banking*, 40, 319–340.
- Bachmann, R., Elstner, S., and Sims, E. R. (2013). Uncertainty and economic activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics*, 5(2), 217–249.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593–1636.
- Bañbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2013). Now-casting and the real-time data flow, chapter 4. In Graham Elliott, A. T. (ed.), *Handbook of Economic Forecasting, volume 2A*, pp. 195–327: Elsevier.
- Barsky, R. B., and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3), 273–289.
- Barsky, R. B., and Sims, E. R. (2012). Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review*, 102(4), 1343–77.
- Beaudry, P., and Portier, F. (2006). Stock Prices, News, and Economic Fluctuations. *American Economic Review*, 96(4), 1293–1307.
- Beaudry, P., and Portier, F. (2014). News-Driven Business Cycles: Insights and Challenges. *Journal*

of Economic Literature, 52(4), 993–1074.

- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, **77**, 623–685.
- Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2018). Over-reaction in Macroeconomic Expectations. NBER Working Papers 24932, National Bureau of Economic Research, Inc.
- Cascaldi-Garcia, D. (2017). News Shocks and the Slope of the Term Structure of Interest Rates: Comment. *American Economic Review*, 107(10), 3243–3249.
- Clements, M. P., and Galvão, A. B. (2017). Predicting Early Data Revisions to US GDP and the Effects of Releases on Equity Markets. *Journal of Business and Economic Statistics*, **35:3**, 389 – 406.
- Clements, M. P., and Galvão, A. B. (2019). Data revisions and real-time forecasting. *The Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press. doi:10.1093/acrefore/9780190625979.013.248.
- Croushore, D. (2011a). Forecasting with real-time data vintages, chapter 9. In Clements, M. P., and Hendry, D. F. (eds.), *The Oxford Handbook of Economic Forecasting*, pp. 247–267: Oxford University Press.
- Croushore, D. (2011b). Frontiers of real-time data analysis. *Journal of Economic Literature*, **49**, 72–100.
- Croushore, D., and Evans, C. L. (2006). Data revisions and the identification of monetary policy shocks. *Journal of Monetary Economics*, 53(6), 1135–1160.
- Croushore, D., and Stark, T. (2003). A real-time data set for macroeconomists: Does the data vintage matter?. *The Review of Economics and Statistics*, **85**, 605–617.
- Cunningham, A., Eklund, J., Jeffery, C., Kapetanios, G., and Labhard, V. (2009). A state space approach to extracting the signal from uncertain data. *Journal of Business & Economic Statistics*, **30**, 173–180. doi:10.1198/jbes.2009.08171.
- Engelberg, J., Manski, C. F., and Williams, J. (2011). Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition. *Journal of Applied Econometrics*, 26(7), 1059–1078.
- Fernald, J. G. (2014). A quarterly, utilization-adjusted series on total factor productivity. 2012-19, FRBSF Working Paper (updated March 2014).
- Fève, P., and Guay, A. (2016). Sentiments in SVARs. Tse working papers TSE-656, Toulouse

School of Economics (TSE).

- Forni, M., Gambetti, L., and Sala, L. (2019). Structural VARs and Non-invertible Macroeconomic Models. *Journal of Applied Econometrics*.
- Garratt, A., Lee, K., Mise, E., and Shields, K. (2008). Real time representations of the output gap. *Review of Economics and Statistics*, **90**, 792–804.
- Ghysels, E., and Wright, J. (2009). Forecasting Professional Forecasters. *Journal of Business and Economic Statistics*, **27**, 504–516.
- Ghysels, E. (2016). Macroeconomics and the reality of mixed frequency data. *Journal of Econometrics*, **193**, 294–314.
- Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior Selection for Vector Autoregressions. *The Review of Economics and Statistics*, *97*(2), 436–451.
- Giordani, P., and Söderlind, P. (2003). Inflation forecast uncertainty. *European Economic Review*, **47**(6), 1037–1059.
- Girardi, A., and Reuter, A. (2017). New uncertainty measures for the EURO area using survey data. *Oxford Economic Papers*, **69**, 278–300.
- Jacobs, J. P. A. M., and van Norden, S. (2011). Modeling data revisions: Measurement error and dynamics of ‘true’ values. *Journal of Econometrics*, **161**, 101–109.
- Jorda, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, *95*(1), 161–182.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. London: Mcmillan.
- Kilian, L., and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*. Cambridge: Cambridge University Press.
- Kishor, N. K., and Koenig, E. F. (2012). VAR estimation and forecasting when data are subject to revision. *Journal of Business and Economic Statistics*, **30**(2), 181–190.
- Kurmann, A., and Sims, E. (2017). Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks. Nber working papers 23142, National Bureau of Economic Research, Inc.
- Lagerborg, A., Pappa, E., and Ravn, M. O. (2019). Sentimental Business Cycles. Discussion paper. Mimeo, EUI Florence.
- Leduc, S., and Sill, K. (2013). Expectations and Economic Fluctuations: An Analysis Using Survey Data. *The Review of Economics and Statistics*, *95*(4), 1352–1367.

- Levchenko, A. A., and Pandalai-Nayar, N. (2017). TFP, News, and ‘Sentiments’: The International Transmission of Business Cycles. Working papers, University of Michigan.
- Mankiw, N. G., and Shapiro, M. D. (1986). News or noise: An analysis of GNP revisions. *Survey of Current Business* (May 1986), US Department of Commerce, Bureau of Economic Analysis, 20–25.
- Mankiw, N. G., and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, **117**, 1295–1328.
- Milani, F. (2011). Expectation Shocks and Learning as Drivers of the Business Cycle. *Economic Journal*, *121*(552), 379–401.
- Milani, F. (2017). Sentiment and the U.S. business cycle. *Journal of Economic Dynamics and Control*, *82*(C), 289–311.
- Miyamoto, W., and Nguyen, T. L. (2019). The Expectational Effects of News in Business Cycles: Evidence from Forecast Data. *Journal of Monetary Economics*. Forthcoming.
- Pigou, A. C. (1927). *Industrial Fluctuations*. London: Mcmillan.
- Plagborg-Møller, M., and Wolf, C. K. (2018). Instrumental variable identification of dynamic variance decompositions. Working paper, Princeton University.
- Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. In Taylor, J. B., and Uhlig, H. (eds.), *Handbook of Macroeconomics*, Vol. 2A, Ch. 2, pp. 71–162: Elsevier.
- Romer, C. D., and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, *94*(4), 1055–1084.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, **50**, 665–690.
- Stock, J. H., and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, *128*(610), 917–948.
- Zarnowitz, V., and Lambros, L. A. (1987). Consensus and uncertainty in economic prediction. *Journal of Political Economy*, **95**(3), 591–621.

Figure 1: Expectations Updates measured by the SPF forecasts of quarterly real GDP growth (annualised).

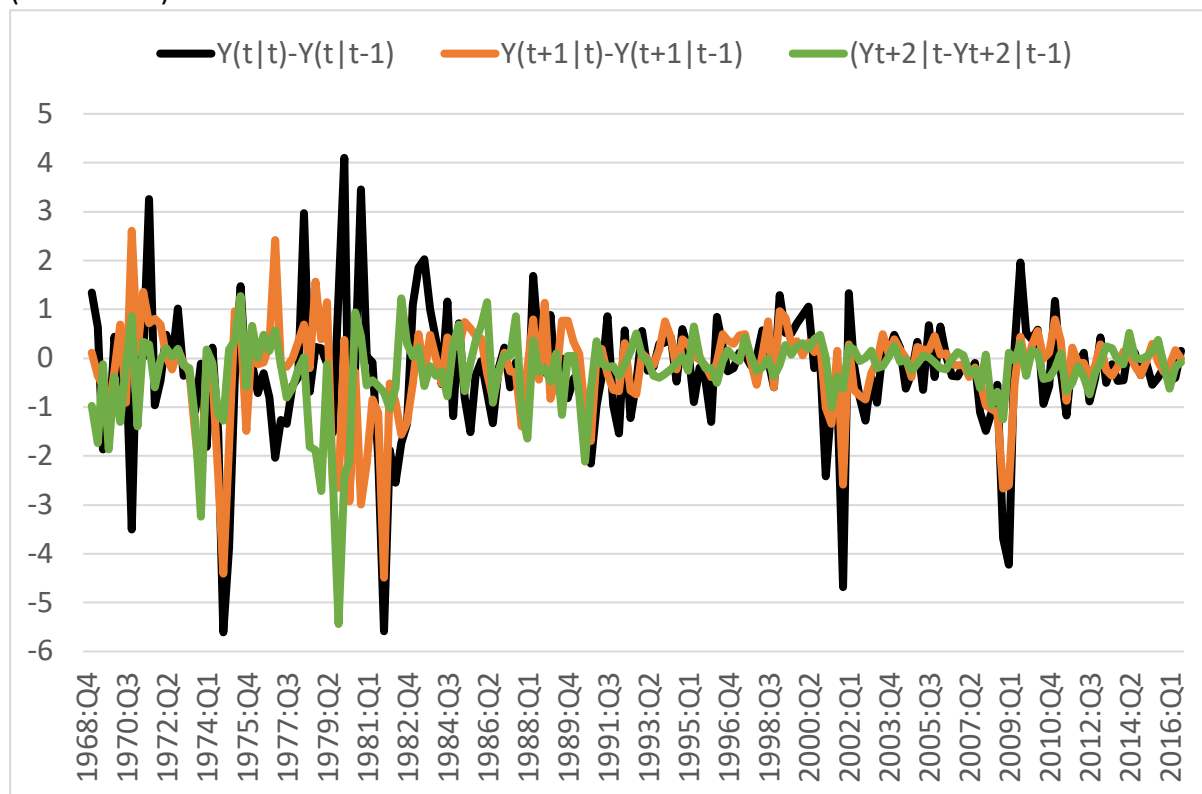
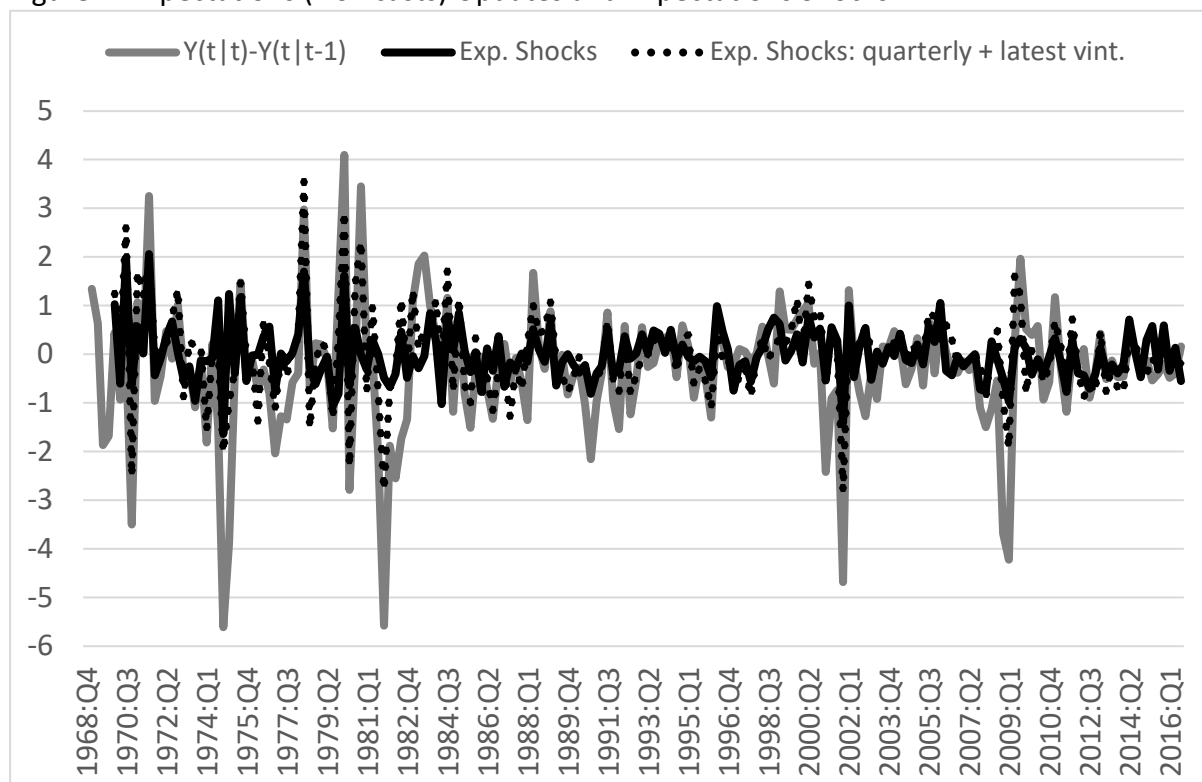



Figure 2: Expectations (Nowcasts) Updates and Expectations Shocks



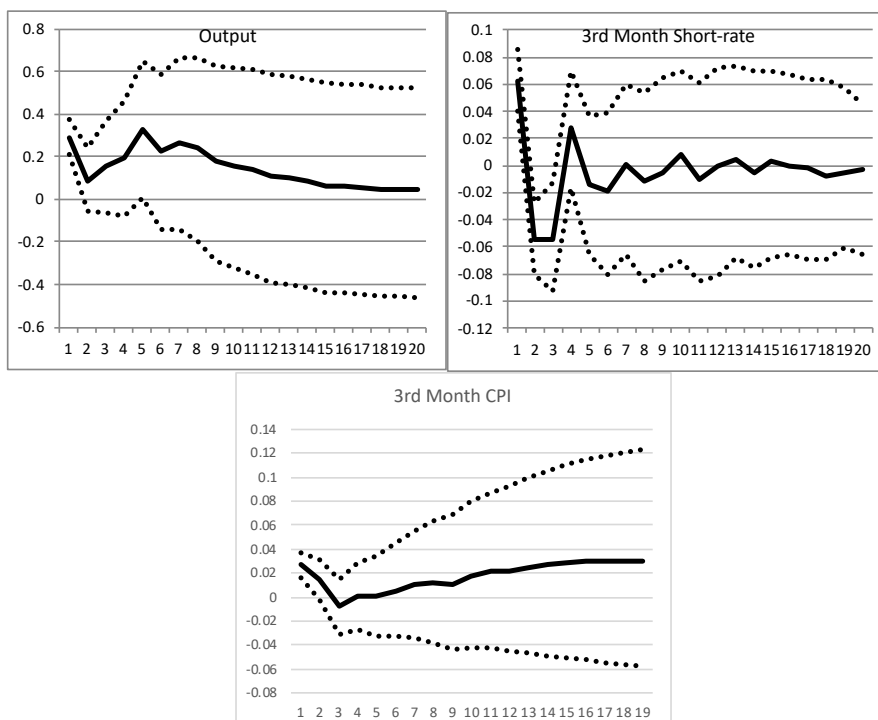
Note: Expectations shocks are computed with the real-time mixed-frequency VAR estimated for the 1969-2016 sample with parameters at the posterior mean. The “quarterly + latest vint.” denotes the use of a quarterly VAR model estimated with data from the latest-available vintage to compute the expectations shocks.

Figure 3: Responses to Expectations Shocks with the Real-Time Mixed-Frequency VAR




Notes: Values are the mean response computed for 20,000 draws from the posterior distribution of the parameters. Dotted lines are 68% confidence bands. Sample period: 1968Q4-2016Q3.

Figure 4: Responses to Expectations Shocks with the “latest-vintage” Mixed-Frequency VAR




Notes: Values are the mean response computed for 20,000 draws from the posterior distribution of the parameters. Dotted lines are 68% confidence bands. Sample period: 1968Q4-2016Q3.

Figure 5: Expectations Shocks and Forecaster Disagreement



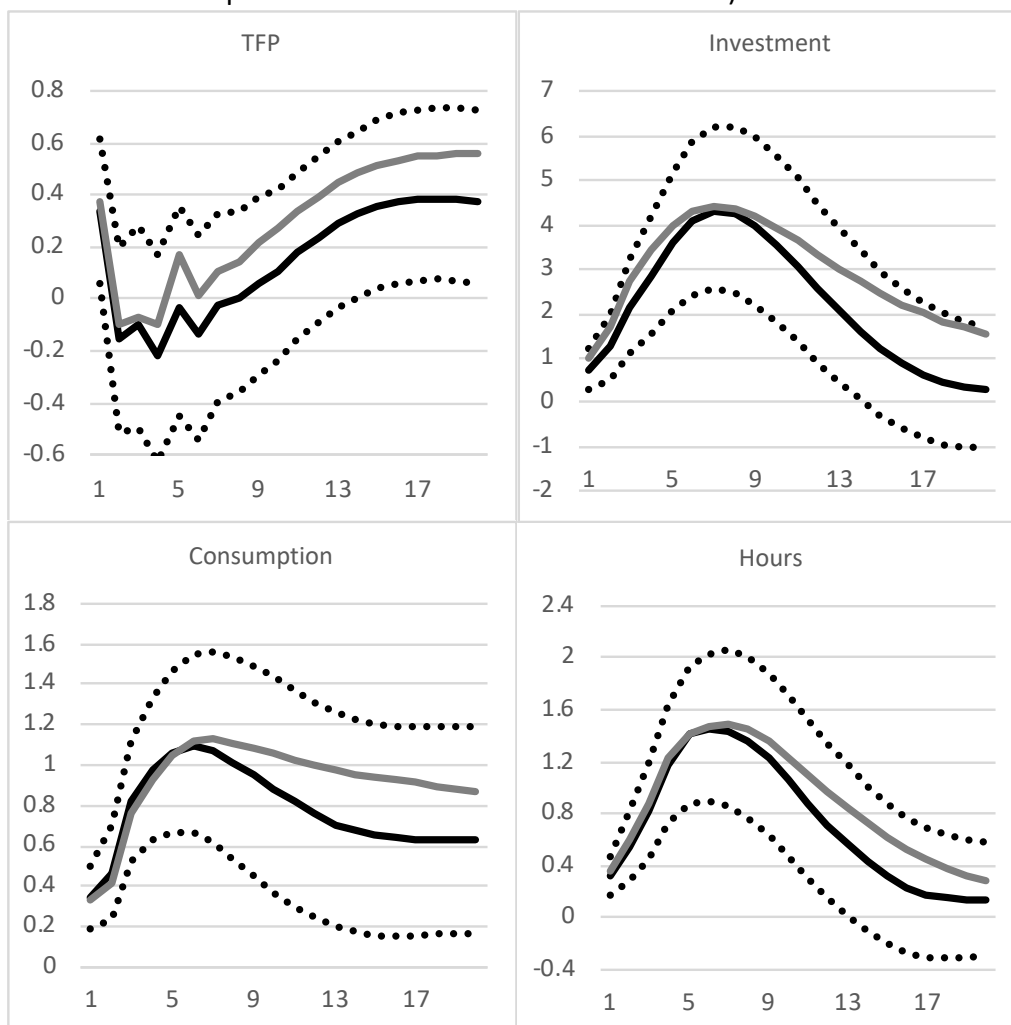
Note: The regression of the squared shocks on the disagreement series yielded a t -statistic of 4.32 and an R^2 of 19%.

Figure 6: Responses of Output to expectations shocks: first-release (black_X) in the RT-VAR, in the macroeconomist VAR (black) and in the instrumented macroeconomist VAR (grey).



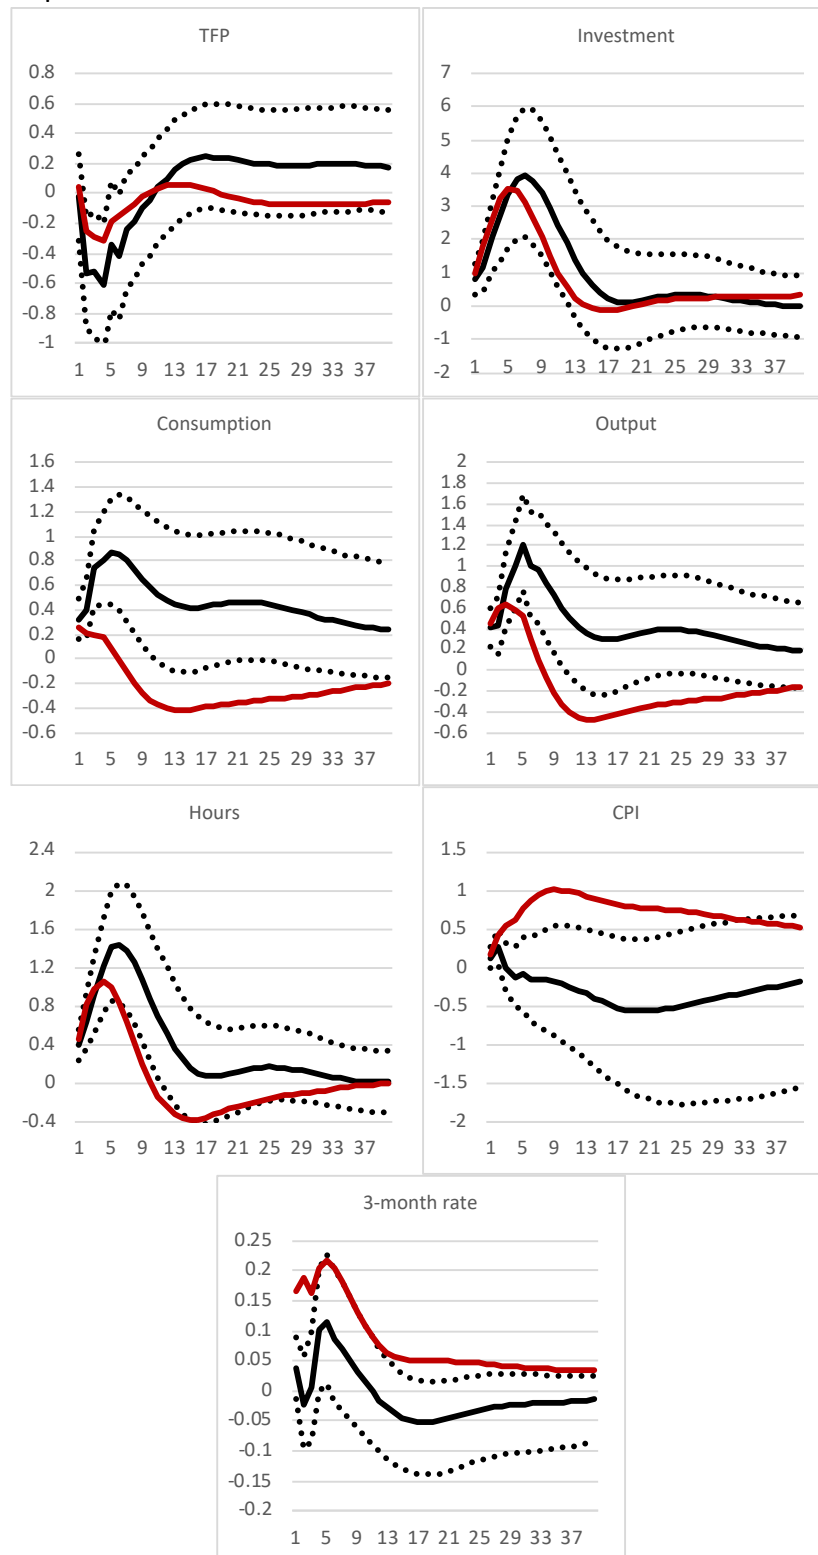
Notes: Values are the mean response computed for 20,000 draws from the posterior distribution of the parameters. Dotted lines are 68% confidence bands. Sample period: 1970Q1-2012Q4.

Figure 7: Responses to Expectations Shocks with the Instrumented Macroeconomist's VAR (grey lines are mean responses without the use of instruments)



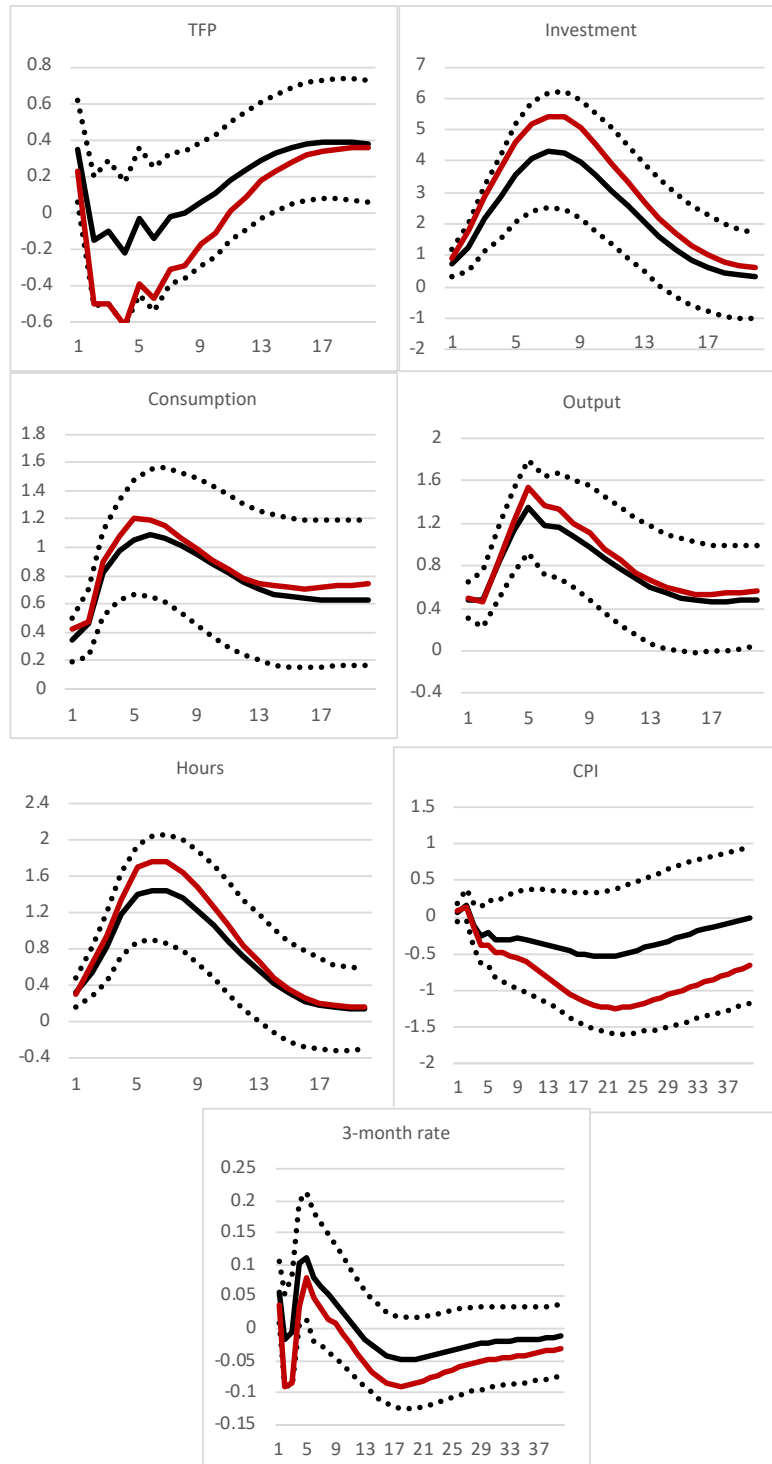
Notes: Values are the mean response computed for 20,000 draws from the posterior distribution of the parameters. Dotted lines are 68% confidence bands. Sample period: 1970Q1-2012Q4.

Figure 8: Responses to Expectations Shocks: $\hat{u}_{t|t}^{exp}$ vs recursive identification applied directly to expectations updates.



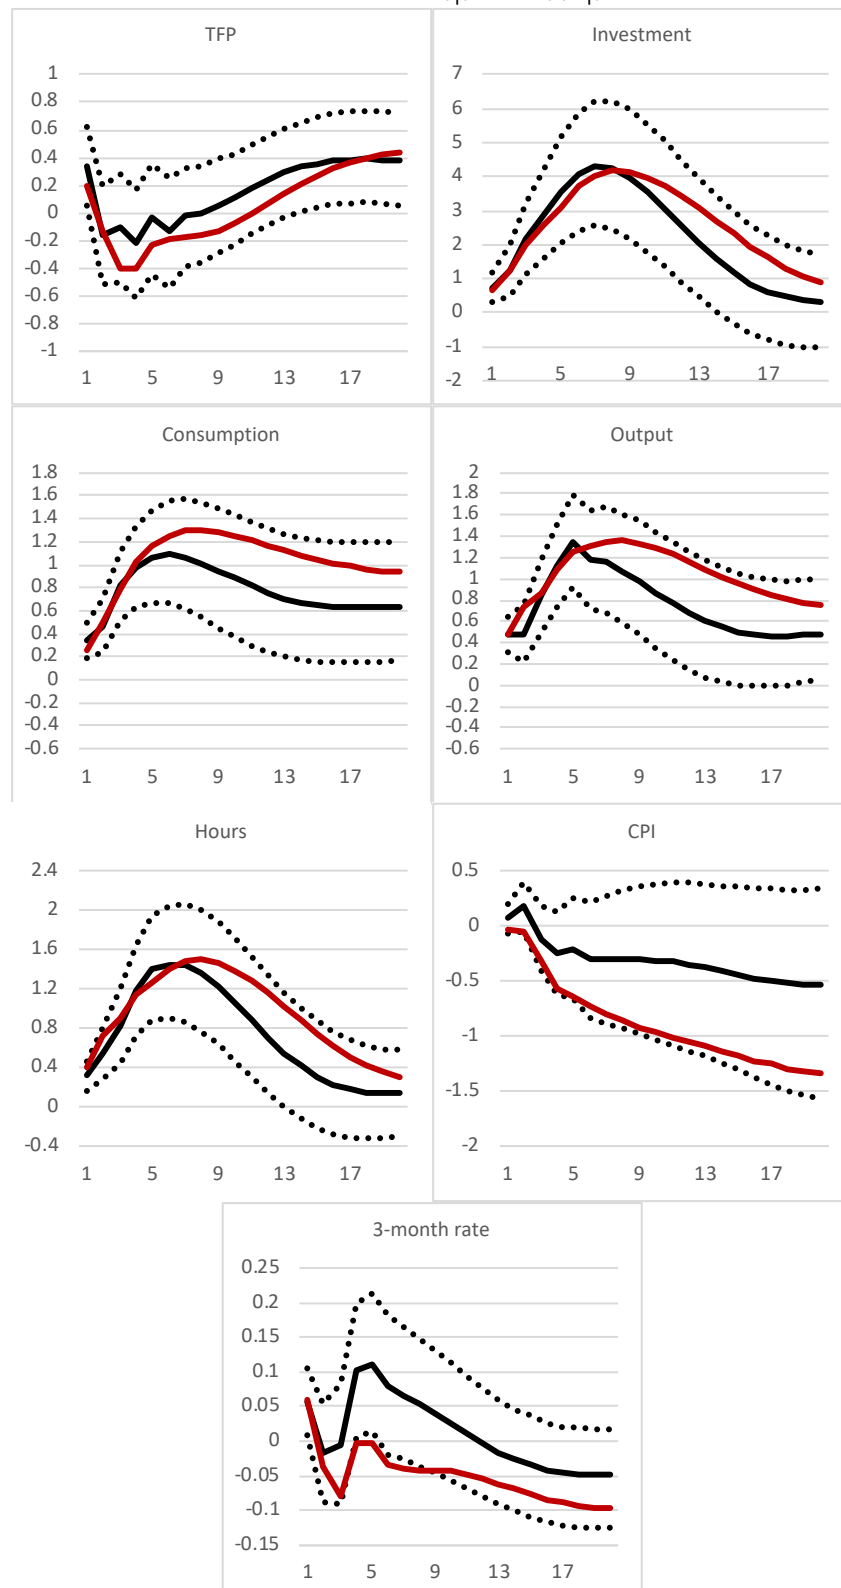
Note: The black lines are responses computed using the instrumented Macroeconomist's VAR using expectation shocks estimated with the real-time mixed-frequency VAR ($\hat{u}_{t|t}^{exp}$). Dotted line are 68% bands. The red lines are posterior mean responses computed using a recursive identification with expectations updates ($Y_{t|t} - Y_{t|t-1}$) as the first variable (instead of the shocks) in the Macroeconomist's VAR (estimated with data up to 2012Q4, not instrumented).

Figure 9: Responses to Expectations Shocks: $\hat{u}_{t|t}^{exp}$ vs expectations shocks estimated using a mixed-frequency VAR with 11 monthly variables.



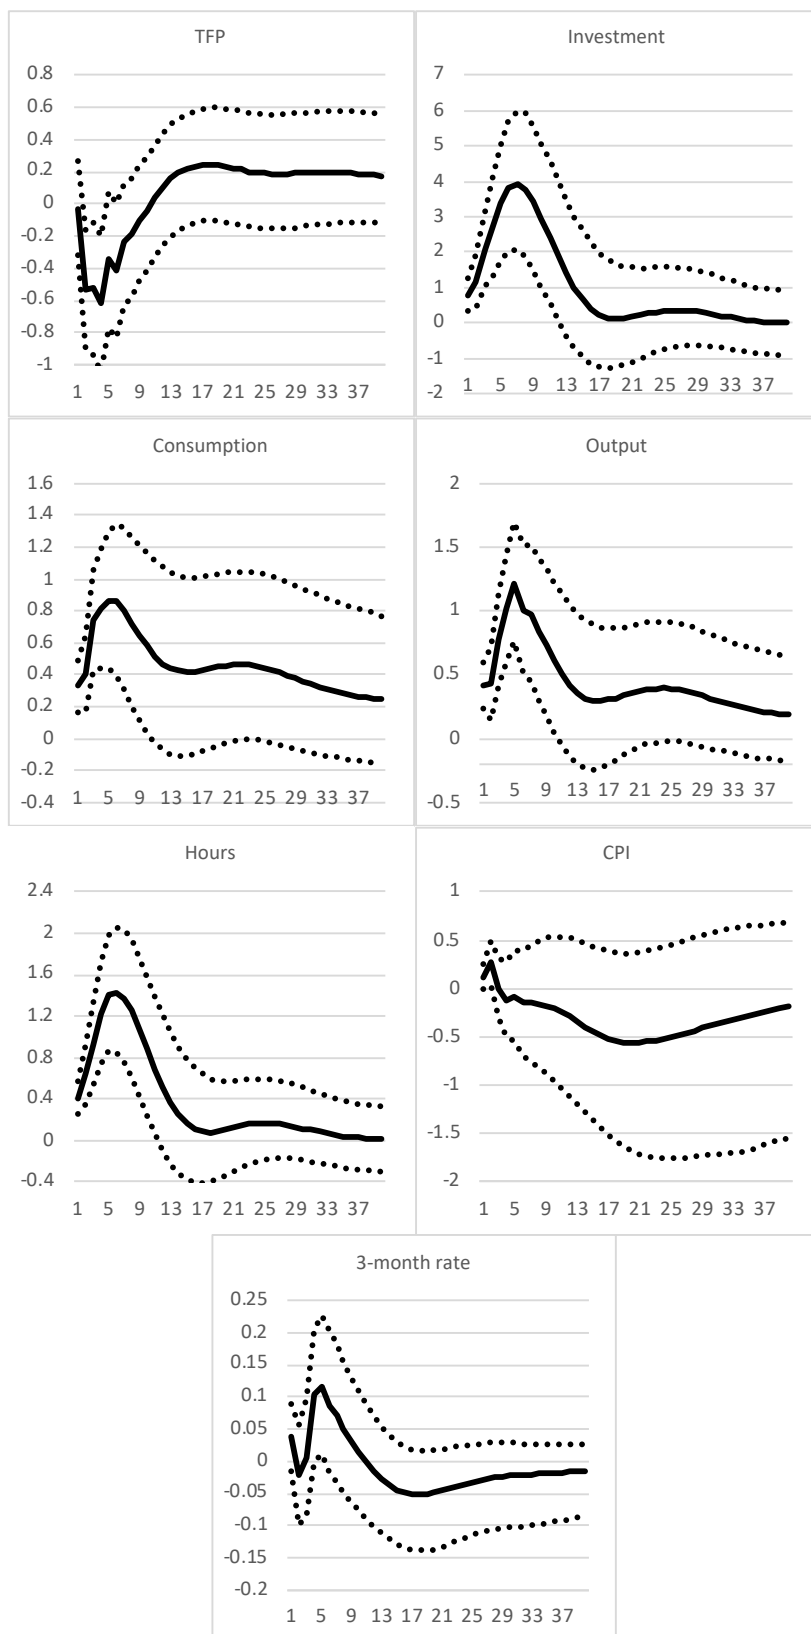
Note: The black lines are responses computed using the instrumented Macroeconomist's VAR using expectation shocks estimated with the real-time mixed-frequency VAR ($\hat{u}_{t|t}^{exp}$) with 5 monthly variables. Dotted line are 68% bands. The red lines are posterior mean responses computed using an alternative estimate of the expectations shocks that employs additional monthly variables in the real-time mixed frequency VAR. The variables included are unemployment, oil prices, long-term rate, corporate bond spread, USS/CAN exchange rate and stock market volatility. The responses are computed using the instrumented Macroeconomist's VAR.

Figure 10: Responses to Expectations Shocks: $\hat{u}_{t|t}^{exp}$ vs $\hat{u}_{t+2|t}^{exp}$



Note: The black lines are responses computed using the instrumented Macroeconomist's VAR using expectation shocks estimated with the real-time mixed-frequency VAR ($\hat{u}_{t|t}^{exp}$) using forecasters' updates to nowcasts. Dotted line are 68% bands. The red lines posterior mean responses computed with the instrumented Macroeconomist's VAR using expectation shocks estimated with the real-time mixed-frequency VAR ($\hat{u}_{t+2|t}^{exp}$) using forecasters' updates to two-quarter-ahead forecasts, that is, $(Y_{t+2|t} - Y_{t+2|t-1})$.

Figure 11: Responses to Expectations Shocks (purged of confidence and news shocks) with the instrumented Macroeconomist's VAR



Notes: Values are the mean response computed for 20,000 draws from the posterior distribution of the parameters. Dotted lines are 68% confidence bands. Sample period: 1970Q1-2012Q4.

Table 1: Characteristics of the Expectations Updates, 1968Q4-2016Q3

	$Y_{t t} - Y_{t t-1}$	$Y_{t+1 t} - Y_{t+1 t-1}$	$Y_{t+2 t} - Y_{t+2 t-1}$
Mean	-0.318	-0.210	-0.260
Median	-0.259	-0.066	-0.110
Std Dev	1.292	0.961	0.771
Corr with $Y_{t t} - Y_{t t-1}$	1		
Corr with $Y_{t+1 t} - Y_{t+1 t-1}$	0.479	1	
Corr with $Y_{t+2 t} - Y_{t+2 t-1}$	0.076	0.396	1

Table 2: Narrative for Large Expectations Shocks from 1978

Date	Business Cycle Phase	Sign	Possible Link
1978Q2	Expansion	Positive	Anticipation of 1978Q2 rate (16.4%)
1979Q3	Expansion	Negative	Change Monetary Policy regime
1980Q1	Recession – Peak	Positive	
1984Q2	Expansion	Negative	End of high growth recovery
1996Q2	Expansion	Positive	End of Government Shutdown
2001Q4	Recession – Trough	Negative	09/11 Terrorist Attack
2006Q1	Expansion	Positive	Anticipation of 2006Q1 rate (5.1%)
2009Q1	Recession	Negative	Major financial crisis

Note: We draw on the online appendix to Romer and Romer (2004) available at: <https://eml.berkeley.edu/~cromer/Data/Shocks/RomerandRomerNarrativeAppendix.pdf>

Table 3: Proportion of the Variance Decomposition of Output explained by Expectations Shocks in the RT-VAR, the Macroeconomist’s VAR, and the Instrumented Macroeconomist’s VAR, 1970Q1-2012Q4.

h	lower		upper
Real-time VAR – first-released output			
1	0.03	0.057	0.09
2	0.02	0.042	0.07
3	0.03	0.053	0.08
4	0.03	0.052	0.08
8	0.03	0.051	0.08
40	0.03	0.049	0.08
Macroeconomist VAR			
1	0.02	0.044	0.08
2	0.01	0.039	0.08
3	0.02	0.057	0.11
4	0.04	0.083	0.15
8	0.06	0.121	0.21
40	0.03	0.090	0.18
Instrumented Macroeconomist VAR			
1	0.02	0.045	0.08
2	0.02	0.040	0.08
3	0.02	0.060	0.11
4	0.04	0.086	0.15
8	0.06	0.119	0.20
40	0.03	0.076	0.15

Note: Entries are the mean proportion computed for 20,000 draws of the posterior distribution of the parameters. Lower and upper denote the lowers of 68% bands.

Table 4: Correlation between expectations shocks and alternative shocks in the literature

	$u_{t t}^{exp}, s_t$
$s_t = tfpnews_t$	0.158 [1.524]
$s_t = consconf_t$	0.179** [2.321]
$s_t = u_{t t}^{exp,final}$	0.845*** [19.59]

Note: t -statistics of the significance of s_t in the regression for the expectations shock are shown in brackets, computed with the Newey-West HAC estimator. Sample period for the regression is 1970Q1-2016Q3. These alternative shocks – news (tpfnews) and consumer confidence (consconf) are computed as described in section 4.

Table 5: Proportion of the Variance Decomposition explained by Expectations Shocks (purged from confidence and news shocks) in the Instrumented Macroeconomist's VAR for 1970Q1-2012Q4.

h	$u_{t t}^{exp}$			$u_{t t}^{exp*}$		
	lower	upper		lower	Upper	
Output						
1	0.02	0.045	0.08	0.01	0.031	0.06
2	0.02	0.040	0.08	0.01	0.028	0.06
3	0.02	0.060	0.11	0.02	0.045	0.09
4	0.04	0.086	0.15	0.02	0.065	0.12
8	0.06	0.119	0.20	0.03	0.084	0.16
40	0.03	0.076	0.15	0.02	0.051	0.11
Investment						
1	0.00	0.017	0.04	0.00	0.018	0.04
2	0.01	0.024	0.06	0.00	0.021	0.05
3	0.01	0.038	0.08	0.01	0.033	0.07
4	0.02	0.050	0.10	0.01	0.042	0.09
8	0.04	0.095	0.18	0.02	0.075	0.15
40	0.03	0.079	0.15	0.02	0.062	0.13
Consumption						
1	0.01	0.029	0.06	0.01	0.023	0.05
2	0.01	0.037	0.08	0.01	0.028	0.06
3	0.03	0.066	0.12	0.02	0.050	0.10
4	0.04	0.086	0.15	0.02	0.060	0.12
8	0.05	0.107	0.19	0.02	0.064	0.14
40	0.02	0.072	0.16	0.01	0.043	0.11
Hours						
1	0.01	0.025	0.05	0.01	0.037	0.07
2	0.01	0.033	0.07	0.02	0.043	0.08
3	0.01	0.045	0.09	0.02	0.054	0.10
4	0.02	0.064	0.12	0.03	0.070	0.13
8	0.05	0.107	0.19	0.04	0.100	0.18
40	0.03	0.080	0.15	0.03	0.070	0.14
TFP						
1	0.00	0.009	0.03	0.00	0.003	0.01
2	0.00	0.013	0.03	0.00	0.016	0.04
3	0.01	0.014	0.03	0.01	0.023	0.06
4	0.01	0.016	0.04	0.01	0.031	0.07
8	0.01	0.019	0.04	0.01	0.034	0.08
40	0.02	0.041	0.09	0.02	0.038	0.08

Note: Entries are the mean proportion computed for 20,000 draws of the posterior distribution of the parameters. Lower and upper denote the lowers of 68% bands.