Growth expectations, undue optimism, and short-run fluctuations*

Zeno Enders
University of Heidelberg
and CESifo

Michael Kleemann
Deutsche Bundesbank

Gernot J. Müller
University of Tübingen,
CEPR, and CESifo

January 31, 2020

Abstract

This paper proposes a two-step procedure in order to identify belief shocks—shocks to expectations about the current state of the economy. First, we use the Survey of Professional Forecasters to measure the nowcast error about contemporaneous output growth. Second, we extract belief shocks from the nowcast error, once by regressing it on existing measures of structural shocks, and once by imposing sign restrictions on a VAR model. Using both approaches, we estimate how macroeconomic variables respond to belief shocks and obtain very similar adjustment dynamics, notably a high degree of co-movement across variables. Belief shocks account for about one third of output fluctuations in the short run.

Keywords: undue optimism, belief shocks, noise shocks, nowcast error, survey of professional forecasters, business cycle

JEL classification: E32

*E-mails: zeno.enders@uni-heidelberg.de, michael.kleemann@bundesbank.de, and gernot.mueller@uni-tuebingen.de. The authors thank Jonas Arias, Olivier Blanchard, Olivier Coibion, Domenico Giannone, Lutz Kilian, Kristoffer Nimark, Franck Portier, Petr Sedláček, and participants in various seminars and conferences for helpful comments, as well as Kerstin Gärtner and Tobias Anger for research assistance. Part of this research was conducted while Enders was visiting the International Monetary Fund (IMF), the hospitality of which is gratefully acknowledged. The views expressed in this paper represent the authors' personal opinions and do not necessarily reflect those of the Deutsche Bundesbank, the IMF or their staff. The usual disclaimer applies.

1 Introduction

Do autonomous changes in expectations cause business cycle fluctuations? The question dates back to Pigou (1927), who discusses the possibility that "errors of undue optimism or undue pessimism" are a genuine cause of "industrial fluctuations". Keynes' notion of "animal spirits" is a related, but distinct concept.[1] More recently, Beaudry and Portier (2004) explore the possibility of "Pigou cycles" in a quantitative business cycle model. Similarly, Lorenzoni (2009) puts forward a model in which misperceptions regarding the current state of aggregate productivity are a source of cyclical fluctuations. In this paper, we use time-series data to identify shocks to expectations about the current situation of the economy. We refer to these shocks as "belief shocks."

Blanchard et al. (2013) show that identification constitutes a formidable challenge in this case because belief shocks are mistakes of market participants. If, roughly speaking, we were able to detect such mistakes at a given point in time, so should market participants. Hence, there should be an immediate correction and no mistake to speak of. In light of these difficulties, one may resort to estimating full-fledged dynamic general equilibrium models in order to achieve identification (Barsky and Sims, 2012; Blanchard et al., 2013). This approach, however, is fairly restrictive as it imposes a lot of specific structure on the data.

In this paper, as an alternative, we device a two-step procedure. First, we use data from the Survey of Professional Forecasters (SPF) to construct a time series for the nowcast error of output growth. Second, we impose restrictions on the nowcast error to identify belief shocks. The nowcast error is the difference between actual output growth and output growth as perceived in real time. As such it represents a measure of real-time misperceptions and becomes available to the econometrician ex post only. For this reason, it provides us with an informational advantage over market participants. This is the key to our identification strategy. We document below in some detail that the nowcast error can be quite sizable.

[1] Keynes' animal spirits are "a spontaneous urge to action rather than inaction", which drive economic decisions beyond considerations based "on nothing but a mathematical expectation" (Keynes 1936, pp. 161–162).

Importantly, because it is a reduced-form measure, the nowcast error is not necessarily the result of a belief shock but may be caused by all structural shocks.

This insight informs the second step of our procedure. Here, we pursue two complementary strategies in order to extract belief shocks from the time series of the nowcast error. The first strategy is to purify the nowcast error of the effect of non-belief shocks. For this purpose we regress the nowcast error on past and present values of the major shock series that have been identified in earlier work. We refer to the regression residual as the "purified nowcast error". It represents realizations of output growth that a) differ from what is expected in real time and b) cannot be accounted for by non-belief shocks. The idea is straightforward: a positive nowcast error means that beliefs have been too pessimistic. To the extent that the error cannot be accounted for by structural shocks, it represents an adverse belief shock. And indeed, we find that while output rises after a nowcast error, it declines after a purified nowcast error. More generally, we observe that the co-movement of output and the nowcast error changes from positive to negative once it is purified.

This is no coincidence but precisely what one would expect if the purified nowcast error represents a belief shock because such a shock—by its very nature—causes a negative co-movement of output and the nowcast error. Importantly, this is also what sets it apart from other structural shocks. An expansionary shock due to, for instance, an unexpected change in total factor productivity causes a positive nowcast error since the shock is generally not fully observed and/or its effect is not fully processed in real time (Coibion and Gorodnichenko, 2012). At the same time, it boosts output. This feature is not specific to productivity shocks, but a general property of non-belief shocks. Likewise, a generic contractionary shock induces a decline of output, and, at the same time, a negative nowcast error if the shock is not fully observed in real time. In contrast, a favorable belief shock causes perceived growth to overshoot actual growth (a negative nowcast error) and, at the same time, a boost to economic activity. By the same token, an adverse belief shock causes a positive nowcast error and a decline of output. Either way, belief shocks induce economic activity and the nowcast error to move in opposite directions (see also Lorenzoni, 2009)

Our second strategy—complementary to the purification approach—is to identify the effect of belief shocks by imposing ex ante the very nature of belief shocks as a restriction on the data. We do so on the basis of a parsimoniously specified vector autoregression (VAR) model that features time series for output and, again, the nowcast error. Specifically, we rely on sign restrictions and impose that belief shocks induce a negative co-movement of output and the nowcast error. In contrast, non-belief shocks are restricted to induce a positive co-movement. Importantly, we find that the impulse responses to a belief shock identified in the VAR model resemble closely the responses to the purified nowcast error. This result is remarkable because the two identification strategies are conceptually distinct.

It is also robust across alternative specifications, including a VAR model that allows for a third shock which does not impact the nowcast error. Across specifications, we find that the nowcast error increases only temporarily to belief and non-belief shocks. The response of output, instead, is much more persistent. Computing a forecast error variance decomposition on the basis of the estimated VAR model, we find that non-belief shocks account for the largest part of the short-run fluctuations of the nowcast error. They also dominate the short-run fluctuations of output. Still, belief shocks account for about one third of output fluctuations. Also this result is robust across a number of specifications and forecast horizons. Lastly, we use local projections to estimate the responses of a large set of macroeconomic variables to the belief shock (as identified in our VAR model) and find a high degree of co-movement across variables. Again, the responses to the belief shock and to the purified nowcast error are very similar throughout.

We stress that for belief shocks to be reflected in nowcast errors, we require them to pertain to current economic activity. However, we do not rule out that market participants expect changes in current economic conditions to be longer lasting or even permanent. In our analysis, we do not capture the effects of expectation shocks to the extent that they pertain to future fundamentals only. Including forecast errors in the VAR model is of little help in this regard, because they are also the result of fundamental innovations and policy reactions along the entire forecasting horizon. In order to identify "noisy news" rather than belief shocks, one may instead resort to a dynamic rotation of the VAR's reduced-form residuals (Forni et al., 2017).

Conceptually, our analysis relates to a number of recent studies on the role of exogenous shifts in expectations as a source of business cycle fluctuations. Angeletos and La'O (2013) develop a model where "sentiment shocks" arise because market participants are unduly but simultaneously optimistic about their terms of trade. Angeletos et al. (2017) rely on DSGE models in order to show that "confidence shocks", that is, autonomous variations in higher-order beliefs, can account for salient features of the data. In earlier work, Milani (2011) introduces "expectation shocks" in a New Keynesian model with near-rational expectation formation. The model is estimated on U.S. data, including expectations data from the SPF. Expectation shocks are found to account for about half of the volatility of output. More recently, Lagerborg et al. (2019), like us, use a VAR framework in which they identify the effect of sentiment shocks by relying on mass shootings as an instrument. Benhima and Poilly (2020) distinguish between supply noise and demand noise and find the latter more important in accounting for output fluctuations.

A number of contributions have focused on the distinction between unexpected and anticipated technology shocks. Evidence from Beaudry and Portier (2006) suggests that business cycles are largely driven by expected future changes in productivity (see also Beaudry et al. 2011, Schmitt-Grohé and Uribe 2012, Leduc and Sill 2013, Görtz and Tsoukalas 2017), while Barsky and Sims (2011) find the role of expected productivity innovations to be limited. Chahrour and Jurado (2018) clarify the relation between news and noise and find that the importance of beliefs has been understated in estimated business cycle models. Kozlowski et al. (2019), in turn, put forward a model of belief formation in which tail events have lasting effects because the true distribution of events is unknown.

Our analysis also relates to earlier studies that attempt to estimate the importance of optimism or sentiments for business cycle fluctuations. Blanchard (1993) provides an animal-spirits account of the 1990–91 recession, focusing on consumption. Carroll et al. (1994) show that consumer sentiment is a good predictor of consumption spending but suggest a "fundamental explanation" based on habits and precautionary saving motives. Oh and Waldman (1990) show that "false macroeconomic announcements", identified as measurement errors in early releases of leading indicators, cause future economic activity. They refrain from a structural interpretation, however. Mora and Schulstad (2007) show that, once announcements regarding current growth are taken into account, the actual growth rate has no predictive power in determining future growth.

Finally, there is recent work that uses survey-based expectations data in order to show that incomplete information, imperfectly rational expectations or confidence may impact macroeconomic outcomes not only as an autonomous source but also by altering decision-making more generally. Nimark (2014) and Melosi (2017) develop and estimate dispersed-information models on data sets that include inflation expectations as reported in the SPF. Both studies illustrate the potential of informational frictions in accounting for business cycle dynamics. Gennaioli et al. (2015) document that corporate investment is well explained by expectations data that, in turn, fail to satisfy a number of rationality tests. Bachmann and Sims (2012) show that consumer confidence amplifies the transmission of fiscal shocks in times of economic slack.

The remainder of the paper is structured as follows. In the next section we perform the first step of our procedure and construct the nowcast error. We also report descriptive statistics. In Section 3 we turn to the second step and extract belief shocks from the time series of the nowcast error based on two distinct strategies. We show that both deliver very similar results. In Section 4 we quantify the role of belief shocks in accounting for the business cycle. A final section concludes.

2 Nowcast errors

In our analysis we follow on a two-step procedure. Eventually, that is, in the second step, we aim to uncover the effects of *belief shocks*—shocks to expectations about the current state of the economy. In this section, as a first step towards this end, we consider a reduced-form measure of misperceptions by computing *nowcast errors* regarding current U.S. output growth. Nowcast errors can be the result of belief shocks, but they may also be due to fundamental, or "structural", innovations. Nowcast errors are the key to our identification strategy in Section 3 below. In what follows, we describe the construction of nowcast errors and compute a number of statistics in order to illustrate their scope, possible causes, and their relation to economic activity.

2.1 Data

Our main data source is the SPF, initiated by the American Statistical Association and the NBER in 1968Q4, now maintained at the Federal Reserve Bank of Philadelphia.[2] The SPF is a widely recognized measure of private-sector expectations regarding the current state and prospects of the U.S. economy. It is also a frequently used benchmark to assess forecasting models (see, for instance, Giannone et al., 2008).

The survey is conducted on a quarterly basis. We focus on the forecast for output growth in the current quarter, the nowcast. In this regard, it is important to note that panelists receive questionnaires at the end of the first month of the quarter and have to submit their answers by the second to third week of the following month. The results of the survey are released immediately afterwards. At this stage, no information regarding current output is available from the Bureau of Economic Analysis (BEA). At most, in order to nowcast output growth for the current quarter, forecasters may draw on the NIPA advance report

[2] Professional forecasters are mostly private financial-sector firms. The number of participating institutions declined from 50 to fewer than 20 in 1988. After the Philadelphia Fed took over in 1990, participation rose again; see Croushore (1993). Regarding our latest observation in 2014Q4, 42 forecasters participated in the survey.

regarding output in the previous quarter.[3] Predicted quarterly output growth is annualized and measured in real terms. Note that, within the SPF, output is initially measured by GNP, later by GDP.[4] For what follows, we use chained growth rates to construct a measure of output that is consistent with this convention. Last, we focus on the nowcast error for output because the nowcast for other variables is harder to evaluate to the extent that actual data becomes available within the quarter.[5]

Our analysis below is based on nowcast errors. We compute it as the difference between the actual value reported (later) by the BEA and the survey's median nowcast. We use the median nowcast error over all forecasters, as it is less prone to outliers than the mean error. Also, nowcast errors based on the mean rather than the median exhibit a somewhat higher variance. Our main results, however, are unchanged once we use the mean nowcast error. We compute two measures of the nowcast error based on the advance and the final estimate for actual output growth, which correspond to the BEA's first and third data release. We thereby address concerns that the assessment of nowcasts or, more generally, forecasts depends on what is being used as "actual" or realized values (see, for example, Stark and Croushore 2002).[6] Throughout we refer to nowcast errors as either "based on first-release" or "based on final-release" data. Note that our final-release-based measure is computed on data prior to further comprehensive and benchmark revisions of the data, which take place at a later date. We do not benchmark the nowcast against these data because these revisions may involve changes in national accounting standards.[7]

[3] In a robustness exercise below, we use monthly (Bloomberg-survey) data on expectations for industrial production. In this case, participants report their nowcast shortly before the official data release, which takes place after the end of the month.

[4] For the SPF forecasts of GNP/GDP, we use the series DRGDP2, which we obtain from the Real-time Data Research Center of the Philadelphia Fed. This series corresponds to the median nowcast of the quarterly growth rate of real output, seasonally adjusted at annual rates (real GNP prior to 1992 and real GDP afterwards). Prior to 1981Q3, the SPF asks for nominal GNP only. In this case, the implied forecast for real GNP is computed on the basis of the nowcast for the price index of GNP.

[5] In this case, since the deadline for submitting responses to the SPF differs over time, it is unclear whether part of the actual data is already available (e.g., in terms of a monthly release) when a response is submitted.

[6] In fact, the authors consider a set of alternative definitions of actuals and find statistically significant differences in forecast evaluations for real output. We show below, however, that our results hold irrespective of the choice of first or final-release data.

[7] Benchmark revisions take place approximately every five years. Comprehensive revisions are more frequent and may also be quite substantial concerning, for instance, the classification of R&D expenditure.

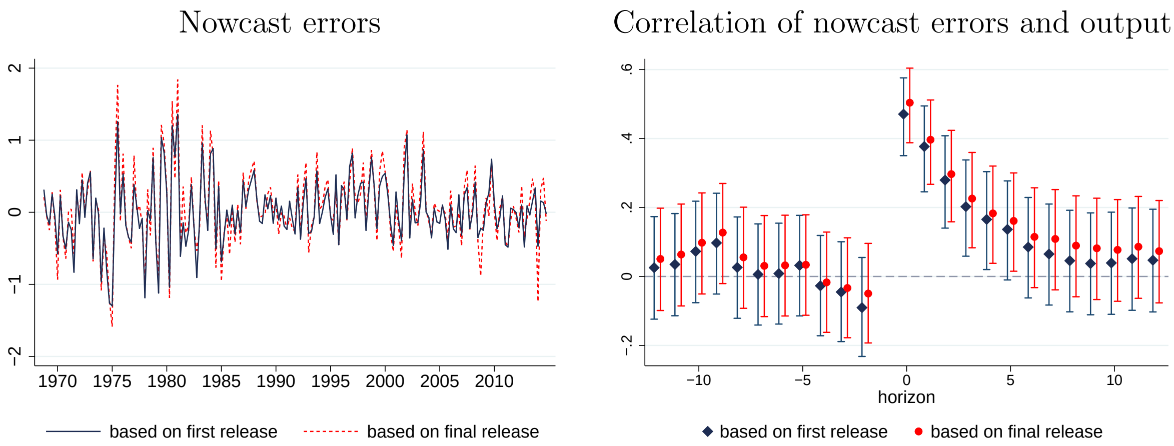

Figure 1: Nowcast errors. Left panel: series based on first-release data (solid lines) and final-release data (dashed lines). Errors are measured in annualized percentage points (vertical axis). Right panel: dynamic correlation of nowcast errors and output at different quarterly horizons. For horizon h, the figure displays the correlation of log(GDP(t+h))- log(GDP(t-1)) and Nowcast Error(t) together with 95%-confidence bounds; correlations for nowcast error series based on first-release data (diamonds) and final-release data (circles)

2.2 Basic properties

The left panel of Figure 1 shows the time series of nowcast errors, measured in annualized percentage points. The solid and dashed lines represent results based on first and final-release data, respectively. Although the two series comove strongly (correlation: 0.94), there are perceptible differences. For instance, only the measure based on the final release displays sizeable negative errors in the second half of 2008. Presumably, at the beginning of the Great Recession, the actual growth slowdown was larger not only relative to what professional forecasters predicted in real time but also to what initial data suggested. The same holds true for 2012Q4, as the U.S. economy approached its so-called fiscal cliff. Instead, during the first half of the sample, errors based on first-release data are shifted somewhat downward relative to those based on final-release data.

We provide summary statistics for both time series in Table 1. The mean of the nowcast error is not significantly different from zero. The standard error and the largest realizations of the nowcast error are somewhat larger in the case of final-release data. Finally, the last two columns of Table 1 report results of a Ljung–Box test, suggesting that there is no serial

Table 1: Summary statistics for nowcast errors

	N	Mean	SD	Min	Max	Ljung–Box test Q-stat.	p-value
Final release	185	0.058	0.539	-1.585	1.843	12.048	0.442
First release	185	0.022	0.457	-1.299	1.358	10.716	0.553

Notes: Nowcast errors computed on the basis of final release (top) and first release (bottom), measured in percentage points; sample: 1968Q4–2014Q4. Means are tested against zero based on a standard t-test. For both series, $H_0 = 0$ cannot be rejected at the 10%-level. The last two columns report Q-statistics and p-values for a Ljung-Box test assessing the null hypothesis of zero autocorrelations up to 12 lags.

correlation in neither series.[8] Hence, in this regard, nowcast errors differ markedly from forecast errors, which tend to exhibit considerable persistence as established in earlier work by Zarnowitz (1985).[9] Also, recent work by Coibion and Gorodnichenko (2012, 2015) shows that forecast errors adjust only sluggishly to new information. We discuss below how our results on nowcast errors relate to theirs.

Nowcast errors are positive surprises regarding current activity. They are also positively correlated with output growth. To explore systematically how current nowcast errors relate to economic activity, we compute the cross-correlation function for the nowcast error and output growth and report the result in the right panel of Figure 1, again for first and final release data. The vertical axis measures the correlation coefficient of the nowcast error with leads and lags of cumulative output growth, measured against the horizonal axis for horizons $\pm h$. The correlation of lagged output growth and the current nowcast error, shown in the left part of the panel, is not significantly different from zero. The contemporaneous correlation, instead, is strongly positive and remains so going forward for several quarters. This finding suggests that what gives rise to nowcast errors also has a lasting impact on GDP.

[8] We also reject the hypothesis that there is first or higher-order serial correlation if tested individually for each lag length up to 12 lags.

[9] He finds that serial correlation in forecast errors tends to increase with the forecasting horizon for many macroeconomic variables in the SPF. In addition, serial correlation seems to be most prevalent in inflation forecasts, generating a large body of literature on the topic, while evidence for GDP forecasts is rather mixed.

In order to provide some suggestive evidence as to what causes nowcast errors, we investigate how popular (and relatively uncontroversial) series of structural innovations impact the nowcast error. Specifically, we consider the following shock measures: innovations in total factor productivity compiled by Fernald (2014), neutral and investment specific technology shocks as identified in Fisher (2002, 2006), monetary policy shocks identified by Romer and Romer (2004) and extended by Coibion et al. (2017), or estimated by Nakamura and Steinsson (2018), uncertainty shocks of Bloom (2009), macroeconomic uncertainty measured by Jurado et al. (2015), defense spending shocks identified by Ramey (2014), tax shocks estimated by Romer and Romer (2010), and oil supply shocks from Kilian (2008).[10]

For each shock series, we regress the nowcast error on the contemporaneous realization of the structural shock while also including four lags of the nowcast error in the regression model. We standardize regression coefficients such that they represent the nowcast error associated with a one-standard-deviation innovation in the shock series. In a second experiment, we also assess whether there is Granger causality by estimating a bi-variate VAR on the nowcast error and, in turn, each shock series. In this case we determine the lag length according to the AIC criterion, except if residuals are autocorrelated, in which case we include additional lags until the autocorrelation vanishes, up to 8 lags. Finally, we compute Wald statistics and test the hypothesis that all coefficients on the lags of the known shock series are jointly zero in the equation for the nowcast error. The sample varies across shock series, since we use the longest overlapping sample in each case.

[10] TFP innovations correspond to the series (dtfp and dtfp_util) derived by Fernald (2014). We obtain the Fisher-type shocks by estimating a Fisher-type VAR based on the relative price of investment, per capita output, and hours per capita. For monetary policy shocks, we use the quarterly average of the monthly shock series (RESID) of Romer and Romer (2004), as extended by Coibion et al. (2017) and of the high frequency measure (FFR_shock) identified by Nakamura and Steinsson (2018), respectively. In case of the former, we also include dummies to control for the impact of the Volcker dis-inflation period, given that Coibion (2012) finds that the two episodes 1980Q2 and 1980Q4, which are showing the strongest shocks, are likely to be particularly noisy for measuring policy innovations and to drive the results of Romer and Romer (2004). Regarding uncertainty shocks, we rely on the quarterly average of the monthly series of stock-market-volatility shocks identified in the baseline VAR of Bloom (2009) and of macroeconomic uncertainty for $h = 3$ months ahead as measured by Jurado et al. (2015), respectively. The defense news identified by Ramey (2014) are the present-value changes in expected defense spending due to political events scaled by lagged nominal GDP. For tax shocks, we use the quarterly average of the monthly shock series "sum of deficit-driven and long-run tax changes" (EXOGENRRATIO) of Romer and Romer (2010). The series of exogenous oil-supply shocks (oilshock) is taken from Kilian (2008).

Table 2: Nowcast errors and fundamental innovations

Structural innovation	Sample	Nowcast error	Concurrent-effect Coeff.	SE	Granger-causality yes/no	prob. > F
Neutral technology	1968-2014	final	0.258***	(0.075)	yes	(0.041)
Fisher type		first	0.245***	(0.074)	yes	(0.089)
Investment-specific techn.	1968-2014	final	0.452***	(0.074)	yes	(0.005)
Fisher type		first	0.419***	(0.089)	yes	(0.000)
Total factor productivity (TFP)	1968-2014	final	0.499***	(0.065)	yes	(0.063)
Fernald (2014)		first	0.461***	(0.067)	yes	(0.061)
Utilization-adjusted TFP	1968-2014	final	0.090	(0.094)	no	(0.443)
Fernald (2014)		first	0.037	(0.094)	no	(0.576)
Monetary Policy	1968-2008	final	0.127	(0.137)	no	(0.643)
Romer & Romer (2004)		first	0.122	(0.138)	no	(0.709)
Monetary Policy	2000-2014	final	-0.089	(0.127)	no	(0.649)
Nakamura & Steinsson (2018)		first	-0.109	(0.101)	no	(0.873)
Uncertainty	1968-2008	final	0.039	(0.074)	no	(0.234)
Bloom (2009)		first	0.026	(0.077)	no	(0.250)
Macro Uncertainty	1968-2014	final	-0.093	(0.102)	no	(0.281)
Jurado et al. (2015)		first	-0.094	(0.102)	no	(0.216)
Defense Spending	1968-2013	final	0.086	(0.075)	no	(0.193)
Ramey (2014)		first	0.064	(0.064)	no	(0.186)
Taxes	1968-2007	final	-0.099	(0.123)	no	(0.427)
Romer & Romer (2010)		first	-0.089	(0.120)	no	(0.516)
Exogenous Oil Supply	1971-2004	final	-0.023	(0.059)	no	(0.633)
Kilian (2008)		first	-0.057	(0.061)	no	(0.574)

Notes: Impact effect on nowcast error obtained from regression of nowcast error on concurrent structural innovations (standardized regression coefficients); regressions include 4 lags of nowcast error; Newey-West standard errors robust for auto-correlation up to 4 lags in parentheses; Granger-causality tests based on bi-variate VARs of structural innovations and nowcast errors; number of lags according to AIC; in case of residual auto-correlation, more lags are included.

Table 2 reports results for both final and first-release data. Newey-West standard errors are displayed in parentheses. The rightmost column reports probability levels for the null hypothesis of no Granger causality based on small-sample F statistics. Overall we find mixed results. Technology shocks, in particular, impact nowcast errors significantly. This holds true for the technology shocks identified à la Fisher, but also for Fernald's (unadjusted) TFP measure. In these instances, there is also evidence for Granger causality.

Instead, for the other shocks we find that results are insignificant. Prima facie, there is also no evidence for Granger causality.[11] However, we note that the evidence presented in Table 2 is merely suggestive. In particular, we cannot entirely rule out that nowcast errors respond to non-technology shocks with a certain delay only. Overall, the results reported in Table 2 suggest a somewhat limited role of well-known shocks in accounting for the nowcast error. This leaves room for belief shocks as an alternative source of nowcast errors. We explore this possibility in the next section.

3 Belief shocks

We now turn to the second step of our approach. The first step was to construct the time series for the nowcast error (Section 2). Now we rely on this time series in order to identify belief shocks and to investigate how they propagate through the economy. We define belief shocks as shocks to expectations about the current situation of the economy. An example of belief shocks are noise shocks in the model of Lorenzoni (2009). In this model, agents observe the components that determine the current level of total factor productivity (TFP) only subject to noise. A noise shock induces agents to attach some probability to an increase of aggregate TFP, even as it has not actually changed. In the working paper version of this paper we develop a simplified version of the model and show in closed form that noise shocks induce a negative co-movement of the nowcast error and output, while other structural shocks induce a positive co-movement, unless they are fully observed in real time (Enders et al., 2017).

[11] Coibion and Gorodnichenko (2012) find that mean forecast errors of inflation respond persistently to shocks. In order to resolve an apparent conflict with our results regarding non-technology shocks, we make two observations. First, we are interested in output growth rather than inflation. In a related paper, Coibion and Gorodnichenko (2015) consider to what extent current forecast revisions predict forecast errors. In a univariate context, the contribution of forecast revisions (averages over all considered horizons; sample: 1968–2014) appears to be strongly significant for inflation, but not significant in the case of output growth. Second, we focus on nowcast rather than on forecast errors. It is thus important to recognize that professional forecasters tend to adjust forecasts rather smoothly (Nordhaus 1987). Indeed, Coibion and Gorodnichenko (2015) find that, while forecast revisions tend to predict forecast errors (averages over all considered variables), the effect is only marginally significant for nowcast errors.

Importantly, this feature is not specific to noise shocks, but the key characteristic of belief shocks more broadly defined. It sets them apart from other structural shocks. As explained in the introduction, any expansionary shock causes a positive nowcast error to the extent that the shock is either not fully observed and/or its effect is not fully processed in real time. A favorable belief shock, on the other hand, causes perceived growth to overshoot actual growth and, hence, a negative nowcast error. At the same time it provides a boost to economic activity. Hence, belief shocks—in contrast to non-belief shocks—induce economic activity and the nowcast error to move in opposite directions.

The nowcast error (as introduced above) is a reduced-form measure of agents' misperceptions and is likely to be caused by many structural shocks, not only by belief shocks. Still, it is key to the second step of our analysis as it provides us with an informational advantage over market participants. This is essential when it comes to identifying belief shocks in time-series data. In what follows we pursue two alternative strategies to implement the second step, that is, to achieve identification. The first strategy is to purify the nowcast error of the effect of non-belief shocks (Section 3.1). For this purpose we regress the nowcast error on past and present values of the major shock series that have been identified in earlier work. We then study how the residual of this regression, that is, the "purified nowcast error," impacts the economy by means of local projections. Second, we estimate a VAR model in which we achieve (set) identification through sign restrictions (Section 3.2).

3.1 Purification

Our first strategy to achieve identification is straightforward.[12] We regress the nowcast error on a set of structural shocks established in the literature. As discussed above, structural shocks are potentially important causes of the nowcast error—to the extent that they or their effect on output are not fully appreciated in real time. In addition, the nowcast error may be caused by a genuine belief shock and the regression residual, to which we refer as

[12]We are very grateful to the editor Olivier Coibion for suggesting this strategy.

the "purified nowcast error", provides us with a candidate of the belief shock. We test this hypothesis informally by contrasting how several key variables respond to the nowcast error on the one hand, and to the purified nowcast error on the other hand.

Figure 2 displays the time series for the purified nowcast error (solid lines) jointly with the original nowcast error series (dashed lines), both for first and final-release data in the left and right panels of the figure, respectively. In our baseline specification, for which we show the results in the upper row of the figure, we consider simultaneously all shock types available to us. Specifically, we use innovations in total factor productivity and utilization adjusted total factor productivity compiled by Fernald (2014), monetary policy shocks identified by Romer and Romer (2004) and extended by Coibion et al. (2017), uncertainty shocks of Bloom (2009), changes in macroeconomic uncertainty measured by Jurado et al. (2015), defense spending shocks identified by Ramey (2014), tax shocks estimated by Romer and Romer (2010), and oil supply shocks from Kilian (2008). We refer to this specification as the "full set of shocks". To purify the nowcast error we regress it on a constant and on the contemporaneous value and 8 lags of all shocks as well as on 8 lags of output growth. The number of lags is conservative because we want to allow for a considerable delay in the impact of structural shocks on the nowcast error. Given the available shock series and the number of lags in our specification, we obtain a time series for the purified nowcast error for the period from 1973Q1 to 2004Q3.

We also consider a limited set of shocks in order to cover a longer period, namely from 1970Q4 to 2013Q4. In this case we cannot consider monetary policy, uncertainty shocks of Bloom (2009), tax shocks, and oil supply shocks in the purification. We refer to this specification as "limited set of shocks". We show results in the bottom panels of Figure 2 and observe that while the purified series is substantially less volatile in both instances, the reduction in volatility is certainly more pronounced for the full set of shocks (upper panels). This is noteworthy because the analysis in Section 2 above shows that only technology shocks impact the nowcast error significantly, see Table 2. Still, since we consider the full and (to

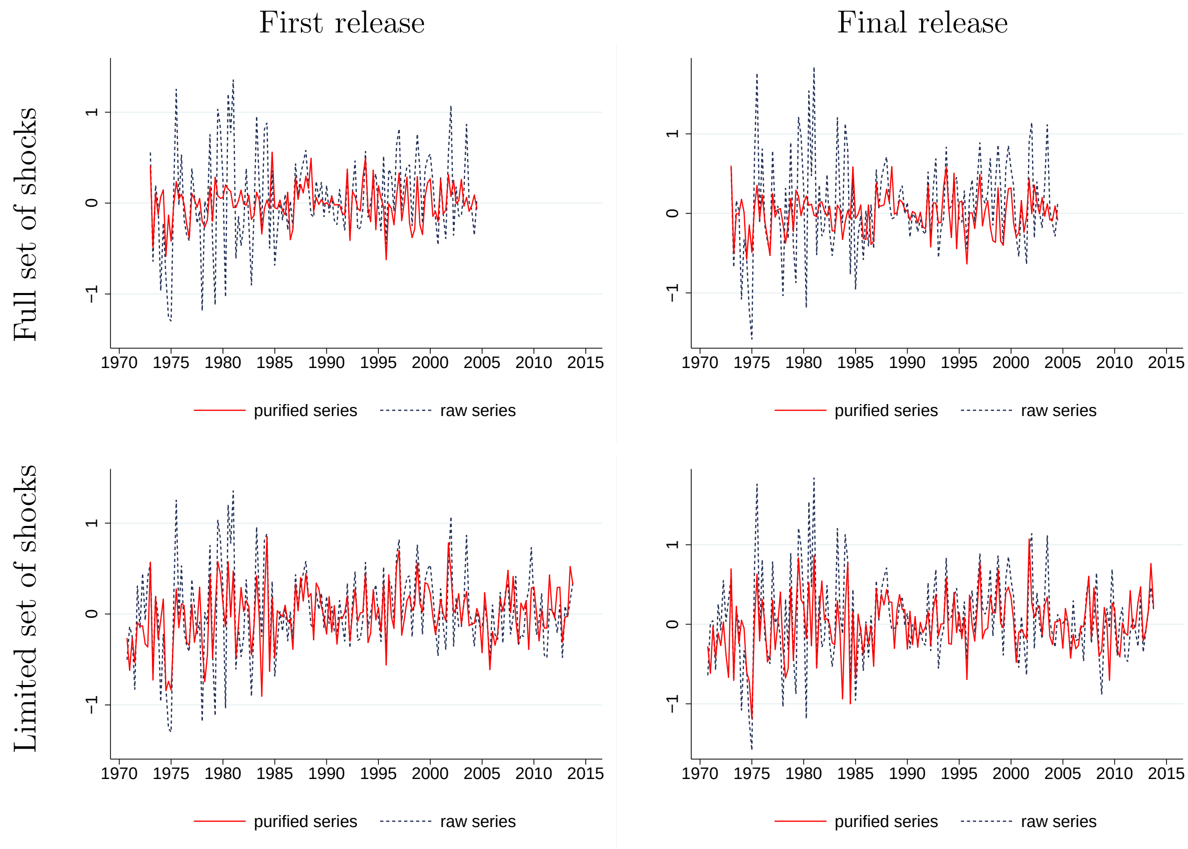

Figure 2: Nowcast error (dashed line) vs. purified nowcast error (solid line). Left and right columns show results based on first and final-release data, respectively; top row based on full set of shocks, bottom panel based on limited set of shocks (longer sample coverage). Errors are measured in annualized percentage points (vertical axis).

a lesser extent) the limited set of shocks, we pursue a deliberately conservative strategy as we permit all shocks at our disposal to cause nowcast errors. As the top row of Figure 2 makes clear, this means that the regression residual—our candidate for the belief shock—is considerably less volatile than the nowcast error. Table 3 provides summary statistics for both specifications. Comparing these statistics to those for the nowcast error reported in Table 1, we observe that the standard error of the purified nowcast error is about one half once we consider the full set of shocks and about two thirds if we consider the limited set.

The purified nowcast error is the part of the nowcast error that cannot be accounted for by structural shocks, neither contemporaneously nor through lags. In what follows we assess the hypothesis that it represents a belief shock. We do so by contrasting the impulse response

Table 3: Summary statistics for purified nowcast errors

	N	Mean	SD	Min	Max	Ljung–Box test Q-stat.	p-value
Purification with full set of fundamental innovations 1973:1-2004:3							
Final release	127	0.000	0.251	-0.635	0.594	12.020	0.444
First release	127	0.000	0.214	-0.622	0.560	6.530	0.887
Purification with limited set of fundamental innovations 1970:4-2013:4							
Final release	173	0.000	0.372	-1.185	1.071	20.303	0.063
First release	173	0.000	0.323	-0.905	0.845	15.613	0.210

Notes: Purified nowcast errors computed on the basis of final-release and first-release data, measured in percentage points. Purification based on alternative set of fundamental innovations. Rightmost columns report Q-statistics and p-values for a Ljung-Box test assessing the null hypothesis of zero autocorrelations up to 12 lags.

functions to the nowcast error and those to the purified nowcast error. For this purpose we estimate local projections (Jordà, 2005). Formally, letting u_t denote the (purified) nowcast error in period t, we estimate the following model:

$$x_{t+h} - x_{t-1} = c^{(h)} + \sum_{j=1}^{J} \alpha_j^{(h)} (x_{t-j} - x_{t-j-1}) + \sum_{k=0}^{K} \beta_k^{(h)} u_{t-k} + \varepsilon_{t+h}. \tag{3.1}$$

Using this specification, we estimate the effect of the nowcast error on a variable of interest, x_t, at horizon h relative to the pre-shock level. In this way we account for the possibility that nowcast errors may have permanent effects on the variables of interest (Stock and Watson, 2018). Our specification includes $J = 4$ lags of the dependent variable and $K = 8$ lags of the nowcast error. $\hat{\beta}_0^{(h)}$ provides an estimate for the impulse response, $c^{(h)}$ is a constant for horizon h, and ε_{t+h} is an iid error term with zero mean. We compute heteroscedasticity and autocorrelation-consistent standard errors as in Newey and West (1987).[13]

We show the results in Figure 3. In the left panels we show the response to the nowcast error, in the right panels the response to the purified nowcast error. Here, we consider the full set of shocks and report results on the basis of first-release data (baseline specification).

[13]Note also that the shocks u_t are generated regressors. Pagan (1984) shows that the standard errors on the generated regressors are asymptotically valid under the null hypothesis that the coefficient is zero; see also the discussion in Coibion and Gorodnichenko (2015).

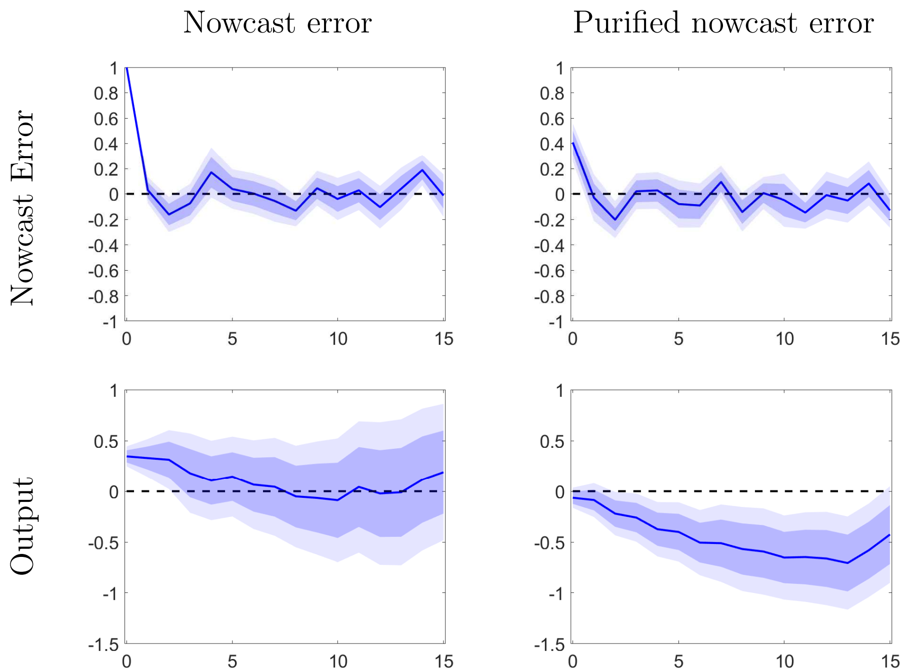

Figure 3: Impulse responses to one-standard deviation of nowcast error (left) and purified nowcast error (right). Notes: Nowcast errors based on first-release data. Purification based on full set of shocks (sample period: 1973Q1–2004Q4). Shaded areas indicate 68% and 90%-confidence bounds. Horizontal axes measure quarters. Vertical axes: deviations from pre-shock level in percentage points.

We normalize the size of the shock to one standard deviation in each instance. Here and in what follows, shaded areas indicate 68% and 90%-confidence bounds around the point estimate for the first release. The horizontal axis measures time in quarters, the vertical axis the deviation from the pre-shock levels in percentage points. The response of the nowcast error is shown in the top row. The increase is purely transitory. We observe no systematic effect after the impact period. The response of output is shown in bottom row. Here we observe an increase in response to a nowcast error, but a decrease in response to the purified nowcast error. The response is significant and persistent in both instances. The fact that the sign of the output response changes once we purify the nowcast error is a central result of our analysis. It is consistent with the notion that the purified nowcast error represents an (adverse) belief shock. It is because of undue pessimism that output contracts. At the same time, this pessimism is reflected in a positive nowcast error—actual output growth is higher than the nowcast, precisely because the latter is too pessimistic.

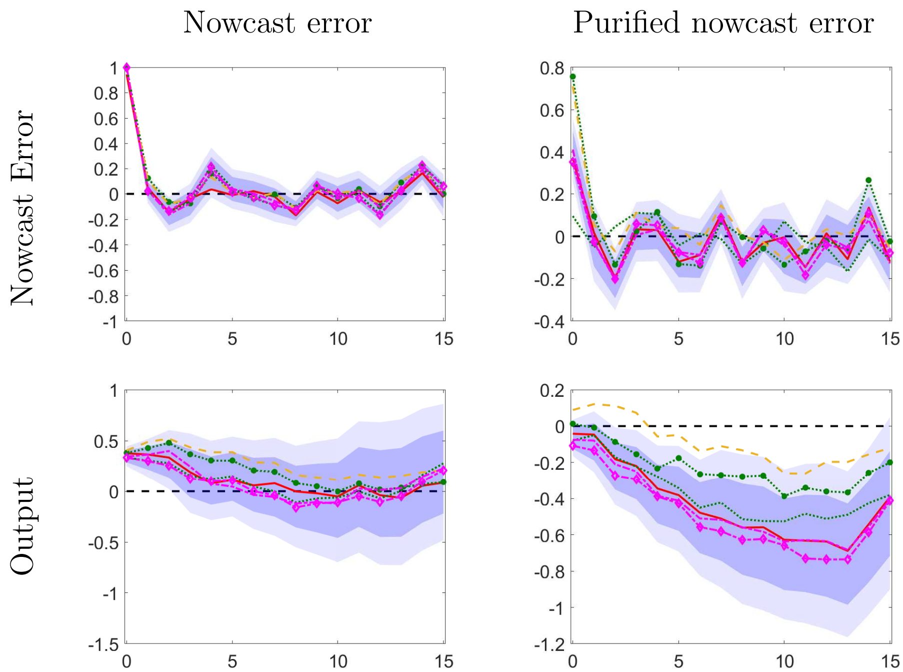

Figure 4: Impulse responses to one-standard deviation of nowcast error (left) and purified nowcast error (right) for alternative specifications. Notes: Shaded areas reproduce 68% and 90%-confidence interval of baseline estimate, shown in Figure 3. Solid line: final-release data; dotted lines with and without asterisk: purification based on 4 and 12 lags of structural innovations, respectively; dash-dotted lines: local projection with 4 additional lags of the dependent variable ($J=8$); dash-dotted lines with diamonds: local projection with 4 additional lags of the shock series ($K=12$); dashed lines: limited set of shocks on the longest possible sample period 1970:4–2013:4.

Figure 4 illustrates that this result is robust across a number of alternative specifications. The figure is organized in the same way as Figure 3 above. It also reproduces the confidence intervals for the baseline estimate. Robustness exercises include constructing the nowcast error based on final-release data (rather than the first release), considering more or less lags of structural innovations in the purification (baseline: 8 lags) or in the local projection (baseline: 4 lags). In all instances, responses are very similar to the baseline and stay largely within the baseline's confidence bounds. Only for the long sample, the output response is somewhat weaker. In this specification, we can only employ a limited set of shocks in the purification, rendering the purification less complete.

3.2 Sign restrictions

We now turn to an alternative identification strategy, based on a VAR model. As before, we seek to extract belief shocks from the nowcast error, constructed in the first step of our analysis (Section 2). In what follows, rather than relying on existing time series of structural shocks as in the previous subsection, we impose ex ante as an identification assumption the defining feature of belief shocks—that they induce a negative co-movement of the nowcast error and output. We work with a parsimoniously specified VAR model which features the nowcast error and output as endogenous variables.

Formally, following the textbook treatment by Kilian and Lütkepohl (2017), we collect both the nowcast error and output from top to bottom in the vector y_t and represent the VAR model in reduced form as follows:

$$y_t = A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t. \tag{3.2}$$

The model also includes a constant and a linear-quadratic time trend, which we omit in expression (3.2) to ease the notational burden. A_i are coefficient matrices, p is the number of lags, and u_t is a vector of potentially mutually correlated innovations with covariance matrix $\Sigma_u = \mathbb{E}(u_t u_t')$.

We estimate the model equation-by-equation with OLS. Given the estimated model, we identify two structural shocks, namely, the belief shock, denoted w_t^b, and a second structural shock, denoted w_t^n. The latter is a composite of all structural shocks insofar as they differ from the belief shock and, hence, we refer to it as the "non-belief shock". It differs from the belief shock because it causes the opposite co-movement of output and the nowcast error: while it is negative for belief shocks, it is positive for non-belief shocks.

We collect the structural shocks in the vector $w_t = \begin{pmatrix} w_t^n & w_t^b \end{pmatrix}'$. They are uncorrelated and their variance is normalized to unity such that $\mathbb{E}(w_t w_t') = I$. We need to find the matrix B which maps structural shocks into reduced form innovations: $u_t = Bw_t$. In line with the arguments put forward above, we restrict the sign of the elements of B as follows:

$$\begin{bmatrix} u_t^{ne} \\ u_t^{y} \end{bmatrix} = \begin{pmatrix} + & + \\ + & - \end{pmatrix} \begin{bmatrix} w_t^{n} \\ w_t^{b} \end{bmatrix}. \tag{3.3}$$

Imposing sign restrictions on matrix B, rather than, say, a zero restriction, means that the impact matrix is only set identified. That is, there are multiple matrices B rather than a single matrix that satisfies the restrictions implicit in (3.3). We find these matrices using Givens rotation matrices (for details, see Kilian and Lütkepohl, 2017).

Including the nowcast error in the VAR model (3.2) is essential in order to capture the effect of belief shocks. Conventional VAR models have difficulties to recover this type of shocks since they represent mistakes of market participants (Blanchard et al., 2013). By their very nature such mistakes are hard to observe in real time, since otherwise they would be corrected instantaneously. The nowcast error, however, provides us with an informational advantage over market participants. Importantly, this advantage becomes available to us only ex post. We establish formally that this allows us to recover belief shocks in a VAR setup in the working paper version of this paper, based on a stylized version of the dispersed-information model of Lorenzoni (2009), see Enders et al. (2017).

We also use the full dispersed-information model of Lorenzoni (2009) to verify that the VAR model (3.2) is able to recover belief (or noise) shocks by means of a Monte Carlo experiment. Specifically, we simulate the calibrated model as described in Section II of Lorenzoni (2009). It features technology and noise shocks, where the noise shocks correspond to belief shocks, as defined above. In the model, we construct the nowcast error as the difference between actual output growth and the real-time estimate of current output growth by the agents of the model. We simulate the model in order to generate time series for the nowcast error and output for 185 periods, corresponding to the number of observations that we use in our estimation below. We repeatedly estimate the VAR model (3.2) on simulated time-series data and compute classical confidence intervals for 1000 repetitions in order to

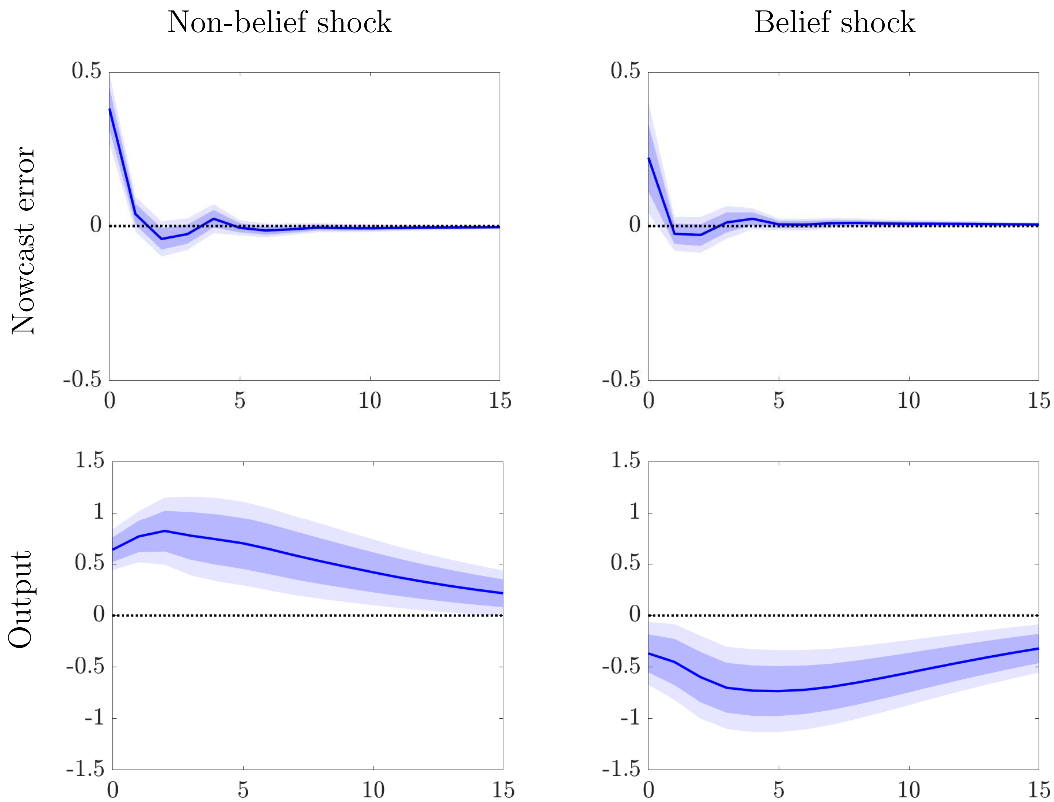

Figure 5: Impulse response to non-belief shock (left) and belief shock (right). Estimates based on VAR model (3.2), identification via sign restrictions. Shaded areas indicate 68% and 90%-confidence intervals around the median responses. Nowcast error computed on basis of first-release data. Horizontal axes measure quarters. Vertical axes: deviations from trend in percentage points.

capture uncertainty about both the identification of the model and the estimation. Next, we compute the median of the confidence interval limits across simulations. Overall, the VAR performs quite well since the true impulse responses are generally within the (median) confidence interval limits. Moreover, we find that the VAR captures the dynamics triggered by the belief shock particularly well. We show results in Figure A.1 in the appendix.

We turn to the actual estimation on quarterly data covering the period 1968Q4–2014Q4. Our sample is restricted because the nowcast error is available only since 1968Q4 (see Section 2). In order to account for possible structural breaks we also consider later starting periods and obtain similar results, as we show in our sensitivity analysis. In our baseline we use the nowcast error computed on the basis of first-release data and set $p = 4$. Figure 5 shows the results. It displays the impulse responses to the non-belief shock and to the belief shock in the left and right column, respectively. In the top row, we show the response of the nowcast error, in the bottom row the response of output. In each panel the vertical axis measures the effect of the shock on the outcome variable in percentage points, while the horizonal axis measures time after impact in quarters.

The nowcast error responds positively to both shocks. Instead, the sign of the output response differs on impact across shocks. This is a direct result of our identification assumption. Note, however, that this assumption restricts only the impact response, not the adjustment dynamics. In this regard, we observe that the response of the nowcast error to both shocks is very transitory, while the response of output is fairly persistent. The non-belief shock is expansionary. In fact, actual output overshoots real-time expectations and hence the nowcast error is positive. Instead, output declines in response to a belief shock. Here we are dealing with an adverse belief shock: pessimism about current output growth causes actual output to decline. Importantly, also in this case output exceeds expectations, such that the response of the nowcast error is positive. Put differently, expectations have been unduely pessimistic. Reassuringly, we find that the dynamic adjustment of output to the belief shock is fairly similar to what we observe in response to the purified nowcast error in Figure 3 above. This holds not only for the size, but also for the shape of the response.

We also verify that our results are robust across a number of alternative specifications. For this purpose, we consider the measure of the nowcast error based on final-release data (rather than first-release data as in the baseline), the mean nowcast (rather than the median), eight lags in the VAR model (rather than four), and a linear time trend (rather than a linear-quadratic trend). Moreover, we also consider a shorter sample for which we disregard observations prior to 1983 because the U.S. business cycle was subject to considerable changes in the early 1980s, possibly due to a change in the conduct of monetary policy (Clarida et al. 2000; McConnell and Perez-Quiros 2000). In all instances, we obtain impulse responses that are very similar to those in the baseline, as Figure B.1 in the appendix shows.

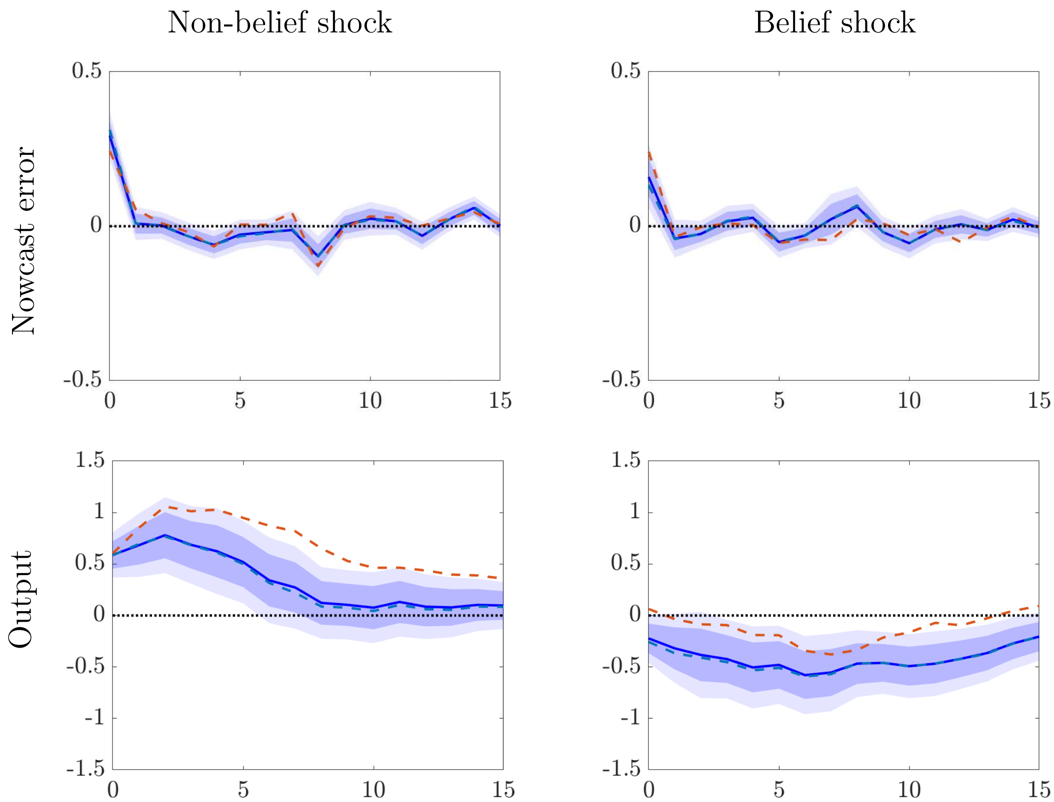

Figure 6: Impulse response to non-belief shock (left) and belief shock (right) based on trivariate VAR model. Identification via combination of size and sign restrictions (solid line), combination of zero and sign restrictions (blue dashed), or combination of long-run and zero restrictions (orange dashed). Shaded areas indicate 68% and 90%-confidence intervals around point estimate for combination of size and sign restrictions. Horizontal axes measure quarters. Vertical axes: deviations from pre-shock levels in percentage points.

In addition, we estimate the VAR model on monthly data (rather than on quarterly data) because our identification strategy hinges on the assumption that current output growth is not directly observed by market participants. This assumption is less restrictive the higher the data frequency. We construct the monthly nowcast error for industrial production from a survey of professional forecasters by Bloomberg.[14] We obtain results that are very similar to the baseline, despite considerable differences in the sample, data frequency, and the measure of economic activity, see Figure B.2 in the appendix.

[14]The Bloomberg survey forecasts are available since 1996M10. We consider data up to 2014M12. We estimate the VAR with 12 lags and a linear time trend.

Next, we consider an extended VAR model. In addition to the nowcast error and output, it includes labor productivity as a third variable. Estimating a trivariate VAR model allows us to consider a possibility which we have ruled out in our VAR analysis up to now: that there are distinct shocks to which the nowcast responds not at all or very little. Specifically, in terms of identification, we consider two possibilities. First, we impose—in addition to the sign restrictions on output and the nowcast error, given in expression (3.3) above—that the third structural shock has a smaller absolute initial impact on the nowcast error than the belief and the non-belief shock. We refer to this restriction as the "size restriction". Second, we alternatively impose that the third shock does not impact contemporaneously the nowcast error at all. We refer to this specification as the "zero restriction". Figure 6 shows the results. The solid blue line, with the corresponding confidence intervals, depicts the median response for the combination of sign and size restrictions, while the blue dashed line shows the median response for the combination of zero and sign restrictions. Reassuringly, results are fairly similar to those of Figure 5.

As a final exercise, we consider an alternative to the sign restrictions used so far. In particular, in line with the perception of a belief shock as a demand shock, we restrict the belief shock to have no impact on labor productivity in the long run. We refer to this restriction as the "long-run restriction". In addition, we maintain the restriction that the third shock has no effect on the nowcast error on impact (the zero restriction) and impose that it has no long-run effect on the nowcast error, rendering the identification just identified, see Enders et al. (2017) for more details.[15] We show point estimates with orange dashed lines in Figure 6. Despite the very distinct approach, we find the results to be again close to those of the baseline, shown in Figure 5. In particular, a positive nowcast error is followed

[15]In this case, we deviate from our baseline VAR specification in several dimension. We use a trivariate VAR including the first-release nowcast error, labor productivity (business Sector: Real Output Per Hour of All Persons) in first differences, and hours worked per capita (Business Sector: Hours of All Persons divided by Civilian Noninstitutional Population). This strategy takes advantage of the level information comprised in hours worked, while avoiding cointegration issues at the same time. We calculate the output response from those of labor productivity and hours worked. Lastly, because of a possible truncation bias when using long-run restrictions, we extend the lag length to 12. The remaining specification is as in our baseline VAR. Results are robust to changes in the specification, such as the lag length.

by a negative output response in case of the belief shock, and by a positive output response in case of the non-belief shock that affects labor productivity in the long run. Importantly, both responses are not restricted to achieve identification. This lends further credibility to our baseline identification scheme and to the view according to which belief shocks are best understood as demand shocks.

4 Belief shocks and the business cycle

Finally, we are in position to take up the question that motivates our analysis, namely whether (and to what extent) belief shocks cause business cycle fluctuations. We try to answer it in three ways. First, we compute a forecast error variance decomposition on the basis of our baseline VAR model. Second, we use the same model to compute a historical decomposition. Last, we use the shocks identified in our baseline model to shed more light on the transmission of belief shocks. Overall, we provide evidence in support of the notion that belief shocks indeed are a major source of the business cycle.

Consider the forecast error variance decomposition first, shown in Table 4 for for 4-quarter and 8-quarter forecast horizons. In the first row we report the results for our baseline VAR model. We find that belief shocks account for about a quarter of the variation in the nowcast error (NE) and about a third of the variation in output (Y) at a one-year horizon. The remainder of the variation is accounted for by non-belief shocks. It is noteworthy that the relative contribution of the non-belief shock is larger for the nowcast error than for output, suggesting that a nowcast error reflects to a large extent the effect of fundamental shocks. Overall, these numbers turn out to be quite robust across alternative specifications of the VAR model. We report them in the lines 2-6 of the table. At a two-year horizon we find the contribution of belief shocks to output fluctuations even higher. In this case they amount to about 40 percent and, again, this result is robust across specifications.

Table 4: Forecast error variance decomposition

	Horizon: 1 year				Horizon: 2 years			
	Non-belief shock		Belief shock		Non-belief shock		Belief shock	
	NC	Y	NC	Y	NC	Y	NC	Y
Baseline	74.88	65.69	25.12	34.31	74.74	55.60	25.26	44.40
8 lags	76.86	67.58	23.14	32.42	70.23	53.86	29.77	46.14
Short sample	66.31	63.09	33.69	36.91	66.23	60.61	33.77	39.39
Linear trend	75.74	69.79	24.25	30.21	75.81	62.72	24.19	37.28
Final release	75.84	67.28	24.16	32.72	75.75	56.07	24.25	43.93
Mean nowcast	75.25	65.46	24.75	34.54	75.13	56.98	24.87	43.02

Notes: contributions in percent, alternative VAR specifications.

In sum, the contribution of the belief shock to output fluctuations is fairly large. Blanchard et al. (2013) estimate a medium-scale DSGE model featuring "noise shocks". These shocks are structurally identical to belief shocks as defined in the present paper and are found to account for about 20 percent of short-run output volatility. Hürtgen (2014) reports similar values for output growth. Barsky and Sims (2012), estimating a fully specified DSGE model by means of indirect inference methods, find that "animal spirit" shocks do not contribute to the volatility of output.[16]

These are, however, model-based estimates and thus likely to understate the role of belief shocks (Chahrour and Jurado, 2018). Against this background, we note that Forni et al. (2017), using a VAR-based approach, find that "noisy news" account for about 40 percent of output fluctuations at a two-year horizon. Likewise, Lagerborg et al. (2019), in their VAR-based study, find that sentiment shocks account for a quarter of the fluctuations in industrial production at a one-year horizon.

[16] While their animal spirit shock is conceptually closely related to our belief shocks, it is restricted to pertain to future productivity (growth) only. Moreover, their analysis is centered around innovations in consumer confidence as reported by the Michigan Survey of Consumers. They find these innovations to reflect correctly anticipated future output growth, that is, according to their estimates, confidence innovations represent news rather than undue optimism. We also explore whether including their shock series in the purification above alters the output response to the purified nowcast error. We find that it does not and conclude that their shock is quite distinct from the belief shock that we identify. We also note that their shocks and the purified nowcast error (baseline) are negatively correlated.

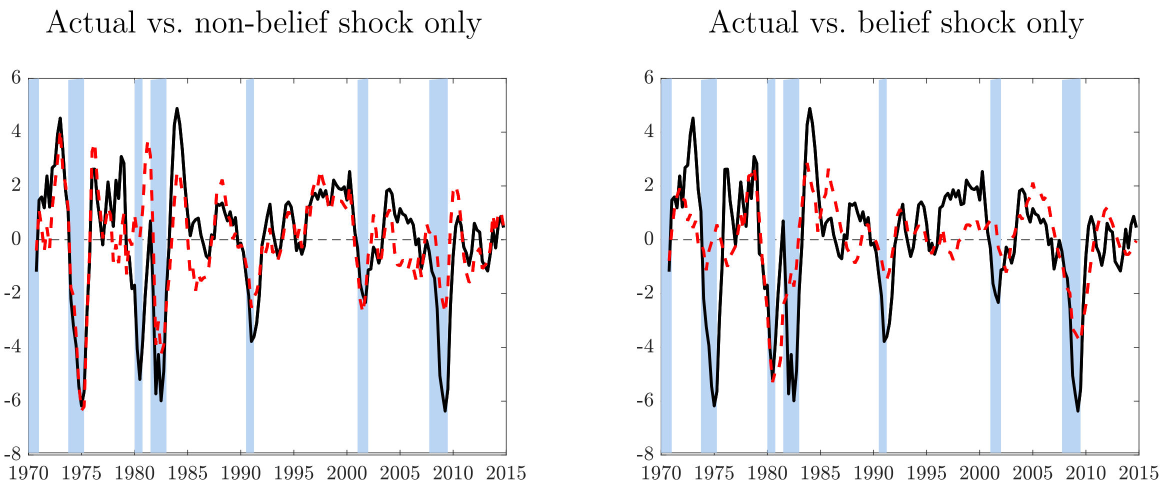

Figure 7: Historical decomposition of output growth (yoy) based on bivariate VAR model. Solid line: actual output growth, dashed line: counterfactual assuming either only non-belief shock (left) or only belief shock (right) occurred. Shaded areas indicate NBER recessions.

Next, we use the estimated VAR model to measure the historical contribution of belief shocks to actual output fluctuations and contrast it to the role of non-belief shocks. We show the result in Figure 7. The panels show actual output growth on a year-on-year basis (solid line) and output growth for a counterfactual assuming that either only non-belief shocks (left panel) or only belief shocks (right panel) had materialized (dashed lines). Shaded areas indicate NBER recessions. It turns out that the majority of recessions are mostly caused by non-belief shocks and some are exclusively due to non-belief shocks. This is particularly so for the 1973–75 recession, which is typically attributed to oil-price shocks, for the 1981–82 recession, which is typically attributed to the Volker disinflation, and for the 2001 recession, which is often associated with the earlier decline of the stock market. However, there are also recessions to which belief shocks contributed strongly according to our estimates. A case in point is the relatively short-lived recession in the first half of 1980. Belief shocks also contributed to the 1990-91 recession, in line with findings of Blanchard (1993). Finally, we observe that belief shocks account for a sizeable part of the growth collapse during the great recession. In fact, here their contribution exceeds those of non-belief shocks.

In order to shed some light on the transmission of belief shocks, we estimate the impulse responses of a number of macroeconomic variables of particular interest. For this purpose, we rely again on the local projections as introduced in equation (3.1) above and use as shock measure u_t the belief shock identified in the VAR.[17] By way of comparison, we also estimate the responses to the non-belief shock obtained from the VAR, as well as those to the purified nowcast error.

Figure 8 shows the results. In the left column we report impulse responses to the non-belief shock, in the middle column the responses to the belief shock (both as identified in the VAR), and in the right column the responses to the purified nowcast error (as constructed in Section 3.1). In each row we consider different variables of interest, namely, from top to bottom, the nowcast error, consumption, investment, durable consumption and residential investment (all measured in real terms). Results are based on the nowcast error constructed using first-release data. We show the response of the nowcast error in the top row as a reminder that all three shocks induce output growth to overshoot real-time expectations.

Overall the figure provides a coherent picture. First, the sign of all responses (except for the nowcast error in the top row) flip, once we move from the non-belief shock (left column) to the belief shock and the purified nowcast error. Second, the responses to the belief shock and the purified nowcast error are very similar. While we have made both observations already with respect to the response of output, we now establish that this is a broader phenomenon which holds for all macroeconomic indicators that we consider. Third, and relatedly, we observe that the belief shock triggers a strong co-movement across GDP and its components. This provides a rationale for why it features so prominently in our account of business cycle fluctuations, see Table 4 above.

Last, we observe that investment, residential investment, and durable consumption react more strongly to belief shocks than nondurable consumption or output (shown in Figure 5 above). This is consistent with the notion that these variables are to a large extent de-

[17] To extract the structural shocks from the (set) identified VAR model, we construct an impact matrix B that contains the median impact responses of both variables across all rotations. We use this matrix to compute the structural shocks, given the reduced-form innovations.

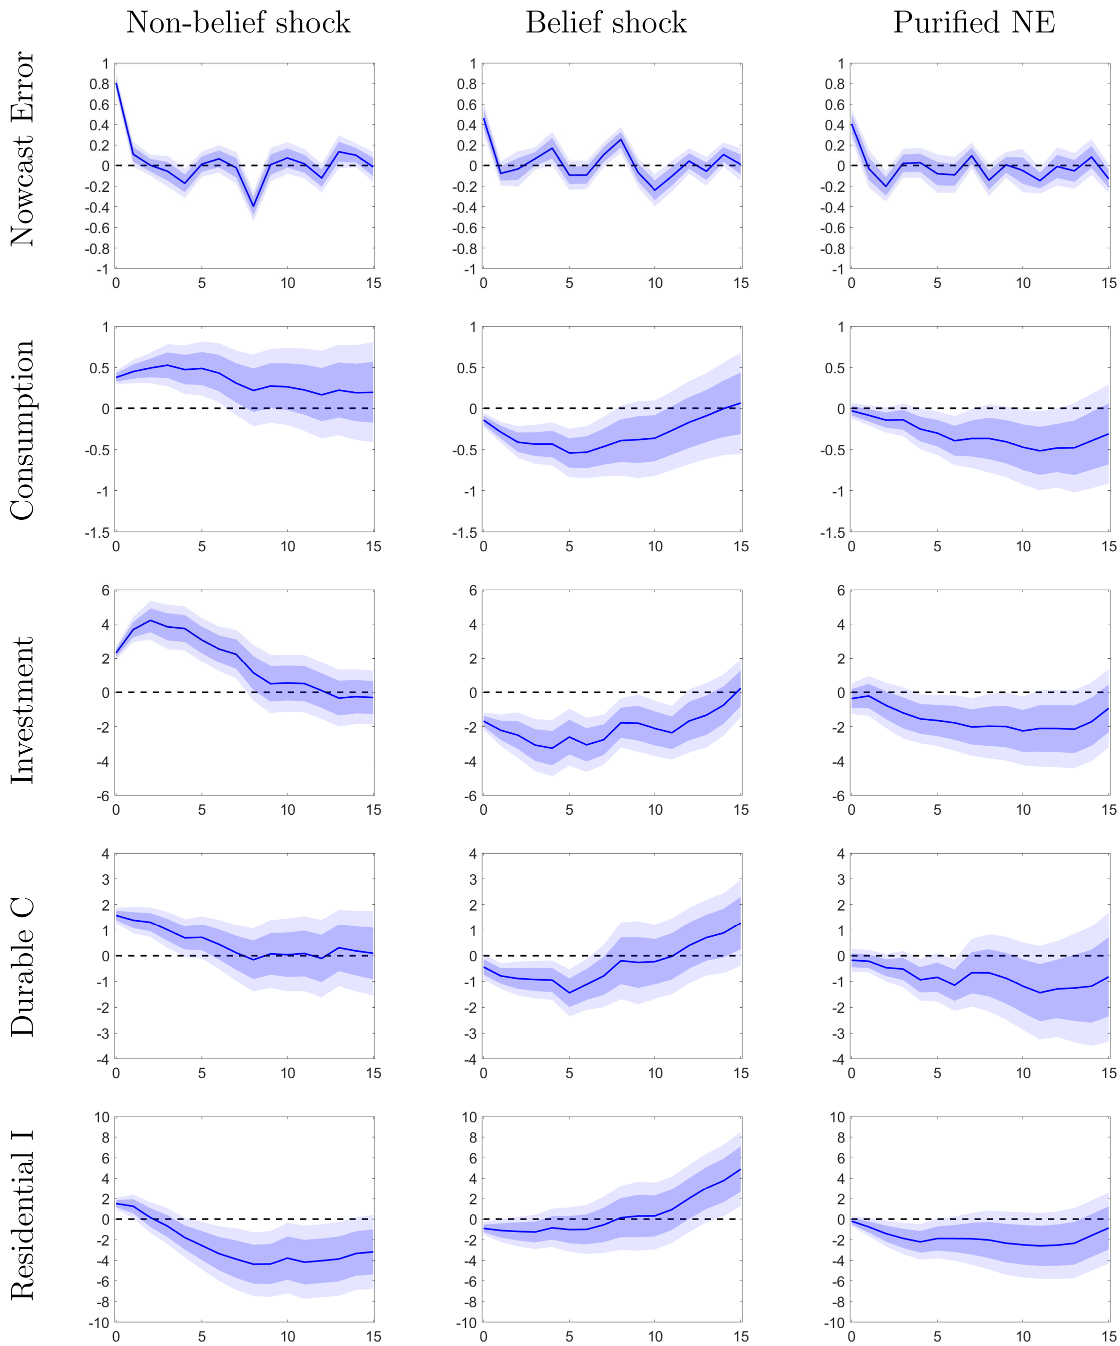

Figure 8: Impulse responses to non-belief shock, belief shock, and purified nowcast error. Notes: Estimates based on local projections (3.1), size of shock normalized to one standard deviation, nowcast error based on first-release data. Solid line represent point estimates, shaded areas indicate 68% and 90%-confidence bounds. Horizontal axes measure quarters. Vertical axes: deviations from pre-shock levels in percentage points.

termined by expectations and hence prone to be affected by belief shocks, see also Enders et al. (2019) for recent evidence based on firm-level data. Also, for durable consumption and

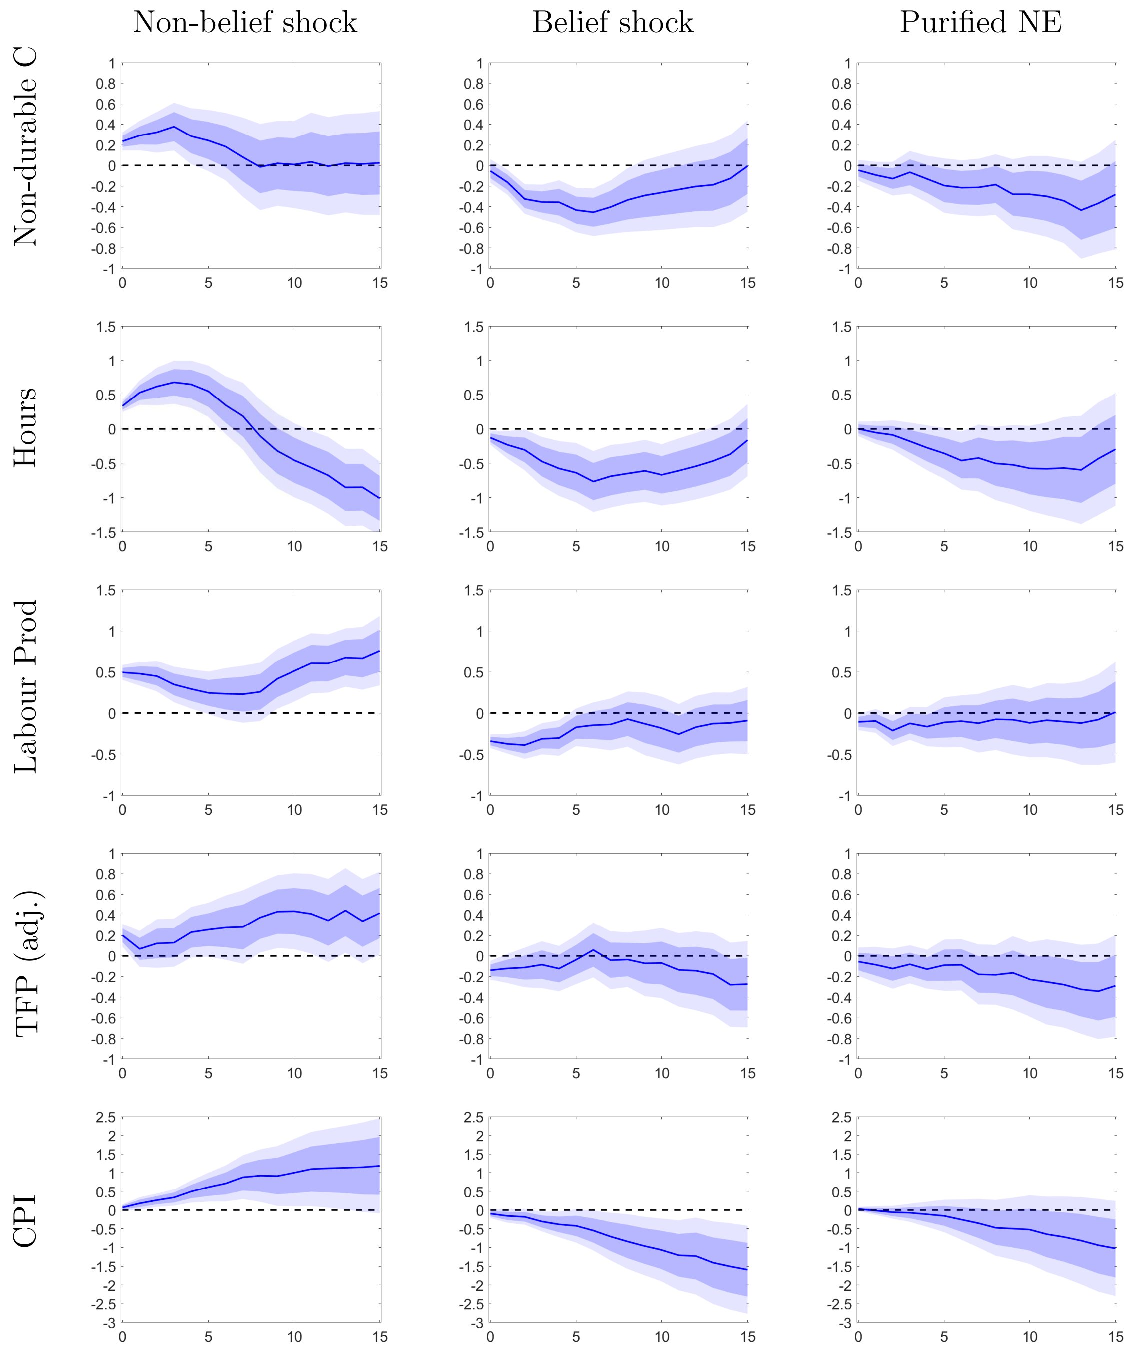

Figure 9: Notes: Impulse responses to non-belief shock, belief shock, and purified nowcast error. Notes: see Figure 8 for details.

residential investment, we observe a reversal of the initial decline following a belief shock: after a period of two to three years, the responses turn positive. Hence, market participants tend to correct their decisions to the extent that they have been caused by belief shocks.

Figure 9 shows the impulse responses of additional variables. It is organized in the same way as the previous figure and the same basic patterns emerge here as well, namely the reversal in the sign of the response as we move from the non-belief shock to the belief shock as well as the similarity in the responses to the belief shock and the purified nowcast error. More specifically, the top row shows the response of non-durable consumption. In reaction to all three shocks, it is less strong than the response of durable consumption shown in Figure 8 above and the rebound after the belief shock is less pronounced. The response of hours is shown in the second row. Remarkably, also here we find a strong degree of co-movement with the components of GDP.

We report the response of labor productivity in the third row of the figure. It increases strongly following a non-belief shock. In this case, we observe a permanent effect—in line with the idea that the non-belief shock reflects to a large extent innovations in TFP. And indeed, we find that utilization-adjusted TFP, shown in the forth row of the figure, increases permanently in response to the non-belief shock. We observe a very short-lived decline of labor productivity in response to the belief shock and the purified nowcast error. This can be rationalized by labor hording in the face of declining demand. In line with this interpretation, we note that TFP does not respond significantly to either the belief shock or the purified nowcast error. Lastly, in the bottom row of the figure, we report the response of the consumer price index (CPI). Here we find that prices increase strongly and permanently following a non-belief shock. Note that this is not inconsistent with a positive TFP shock, if it is highly persistent. More importantly, we observe a strong decline of the CPI in response to the (adverse) belief shock and also in response to the purified nowcast error. This, again, is what we expect to the extent that the belief shock represents a genuine demand shock (Lorenzoni, 2009).

5 Conclusion

Are business cycle fluctuations caused by undue optimism and pessimism and, if so, to what extent? In this paper, we pursue a new approach to address this question. Barsky and Sims (2012) and Blanchard et al. (2013), among others, estimate fully specified DSGE models to quantify the importance of "noise". This approach is fairly restrictive as it imposes a lot of specific restrictions on the data. Moreover, both studies reach quite different conclusions as to the quantitative importance of belief shocks. We therefore pursue an alternative, less restrictive approach based on a structural VAR model. Yet, as shown by Blanchard et al. (2013), identifying the effect of belief shocks within VAR models constitutes a formidable challenge.

We devise a two step-procedure to confront this challenge. First, we construct a time series for the nowcast error of output growth on the basis of the survey of professional forecasters. Nowcast errors are a reduced-form measure of misperceptions and possibly caused by all sorts of structural shocks. They may, however, also reflect belief shocks, that is, shocks to expectations about the current situation of the economy. Still, they are the key to our identification strategy because they provide us with an informational advantage over market participants and allow us to recoup undue beliefs about the economy. This is what we do in the second step.

In this step, we consider two alternative identification schemes. One scheme is to "purify" the nowcast error of the effect of structural shocks. For this purpose, we regress the time series of the nowcast error on existing time series of structural shocks. We refer to the regression residual as the "purified nowcast error". The other scheme consists of imposing sign restrictions on an estimated VAR model. We then compare the impulse responses to the belief shock (as identified in the VAR) and the responses to the purified nowcast error.

We find that a wide range of variables respond very similarly to both shocks. For this reason we are confident that we are indeed able to capture the effect of belief shocks on the macroeconomy. Moreover, we find that belief shocks induce a high degree of co-movement not only of GDP and its components, but also of output and hours as well as prices. This result is consistent with the notion that belief shocks are an important source of the business cycle. We confirm this notion through a forecast error variance decomposition. It shows robustly that belief shocks account for about one third of output fluctuations at an one-year forecasting horizon.

References

Angeletos, G.-M., Collard, F., and Dellas, H. (2017). Quantifying confidence. *Econometrica*, forthcoming.

Angeletos, G.-M. and La'O, J. (2013). Sentiments. *Econometrica*, 81:739–779.

Bachmann, R. and Sims, E. R. (2012). Confidence and the transmission of government spending shocks. *Journal of Monetary Economics*, 59:235–249.

Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58:273–289.

Barsky, R. B. and Sims, E. R. (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review*, 52:1343–1377.

Beaudry, P., Nam, D., and Wang, J. (2011). Do mood swings drive business cycles and is it rational? NBER Working Paper 17651.

Beaudry, P. and Portier, F. (2004). An exploration into Pigou's theory of cycles. *Journal of Monetary Economics*, 51:1183–1216.

Beaudry, P. and Portier, F. (2006). Stock prices, news, and economic fluctuations. *American Economic Review*, 96:1293–1307.

Benhima, K. and Poilly, C. (2020). Does Demand Noise Matter? Identification and Implications. *Journal of Monetary Economics, forthcoming.*

Blanchard, O. (1993). Consumption and the recession of 1990–91. *American Economic Review*, 83:270–274.

Blanchard, O., L'Huillier, J.-P., and Lorenzoni, G. (2013). News, noise, and fluctuations: An empirical investigation. *American Economic Review*, 103:3045–70.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Carroll, C. D., Fuhrer, J. C., and Wilcox, D. W. (1994). Does consumer sentiment forecast household spending? If so, why? *American Economic Review*, 84(5):1397–1408.

Chahrour, R. and Jurado, K. (2018). News or noise? the missing link. *American Economic Review*, 108(7):1702–36.

Clarida, R., Galí, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115:147–180.

Coibion, O. (2012). Are the Effects of Monetary Policy Shocks Big or Small? *American Economic Journal: Macroeconomics*, 4(2):1–32.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

Coibion, O. and Gorodnichenko, Y. (2015). Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review*, 105(8):2644–78.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent Bystanders? Monetary policy and inequality. *Journal of Monetary Economics*, 88(C):70–89.

Croushore, D. (1993). Introducing: the survey of professional forecasters. *Business Review*, 6.

Enders, Z., Hünnekes, F., and Müller, G. J. (2019). Firm expectations and economic activity. CESifo Working Paper 7623.

Enders, Z., Kleemann, M., and Müller, G. J. (2017). Growth expectations, undue optimism, and short-run fluctuations. Discussion Papers 11/2017, Deutsche Bundesbank.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco Working Paper 2012-19.

Fisher, J. D. M. (2002). Technology shocks matter. Working Paper Series WP-02-14, Federal Reserve Bank of Chicago.

Fisher, J. D. M. (2006). The Dynamic Effects of Neutral and Investment-Specific Technology Shocks. *Journal of Political Economy*, 114(3):413–451.

Forni, M., Gambetti, L., Lippi, M., and Sala, L. (2017). Noisy news in business cycles. *American Economic Journal: Macroeconomics*, 9(4):122–52.

Gennaioli, N., Ma, Y., and Shleifer, A. (2015). Expectations and Investment. In *NBER Macroeconomics Annual 2015, Volume 30*, NBER Chapters, pages 379–431. National Bureau of Economic Research, Inc.

Giannone, D., Reichlin, L., and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676.

Görtz, C. and Tsoukalas, J. D. (2017). News and financial intermediation in aggregate fluctuations. *The Review of Economics and Statistics*, 99(3):514–530.

Hürtgen, P. (2014). Consumer misperceptions, uncertain fundamentals, and the business cycle. *Journal of Economic Dynamics and Control*, 40:279–292.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95:161–182.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.

Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money.* Edition Harvest/Harcourt Brace 1964, Orlando, Florida.

Kilian, L. (2008). Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *The Review of Economics and Statistics*, 90(2):216–240.

Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis.* Themes in Modern Econometrics. Cambridge University Press.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2019). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy, forthcoming.*

Lagerborg, A., Pappa, E., and Ravn, M. O. (2019). Sentimental business cycles. mimeo.

Leduc, S. and Sill, K. (2013). Expectations and economic fluctuations: An analysis using survey data. *The Review of Economics and Statistics*, 95:1352–1367.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99:2050–2084.

McConnell, M. M. and Perez-Quiros, G. (2000). Output fluctuations in the United States: What has changed since the early 1980s? *American Economic Review*, 90:1464–1476.

Melosi, L. (2017). Signaling effects of monetary policy. *Review of Economic Studies*, 84:853–884.

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *Economic Journal*, 121:379–401.

Mora, J. V. R. and Schulstad, P. (2007). The effect of GNP announcements on fluctuations of GNP growth. *European Economic Review*, 51:1922–1940.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*. forthcoming.

Newey, W. and West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–08.

Nimark, K. P. (2014). Man-Bites-Dog Business Cycles. *American Economic Review*, 104(8):2320–67.

Nordhaus, W. D. (1987). Forecasting efficiency: Concepts and applications. *The Review of Economics and Statistics*, 69(4):667–674.

Oh, S. and Waldman, M. (1990). The macroeconomic effects of false announcements. *The Quarterly Journal of Economics*, 105:1017–1034.

Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, 25(1):221.

Pigou, A. C. (1927). *Industrial Fluctuations*. Macmillan.

Ramey, V. (2014). Defense news shocks, 1889-2013: Estimates based on news sources. Mimeo.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.

Romer, C. D. and Romer, D. H. (2010). The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *American Economic Review*, 100(3):763–801.

Schmitt-Grohé, S. and Uribe, M. (2012). What's news in business cycles? *Econometrica*, 80:2733–2764.

Stark, T. and Croushore, D. (2002). Forecasting with a real-time data set for macroeconomists. *Journal of Macroeconomics*, 24(4):507–531.

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. (1985). Rational expectations and macroeconomic forecasts. *Journal of Business & Economic Statistics*, 3(4):293–311.

Appendix A: Monte Carlo assessment of the VAR

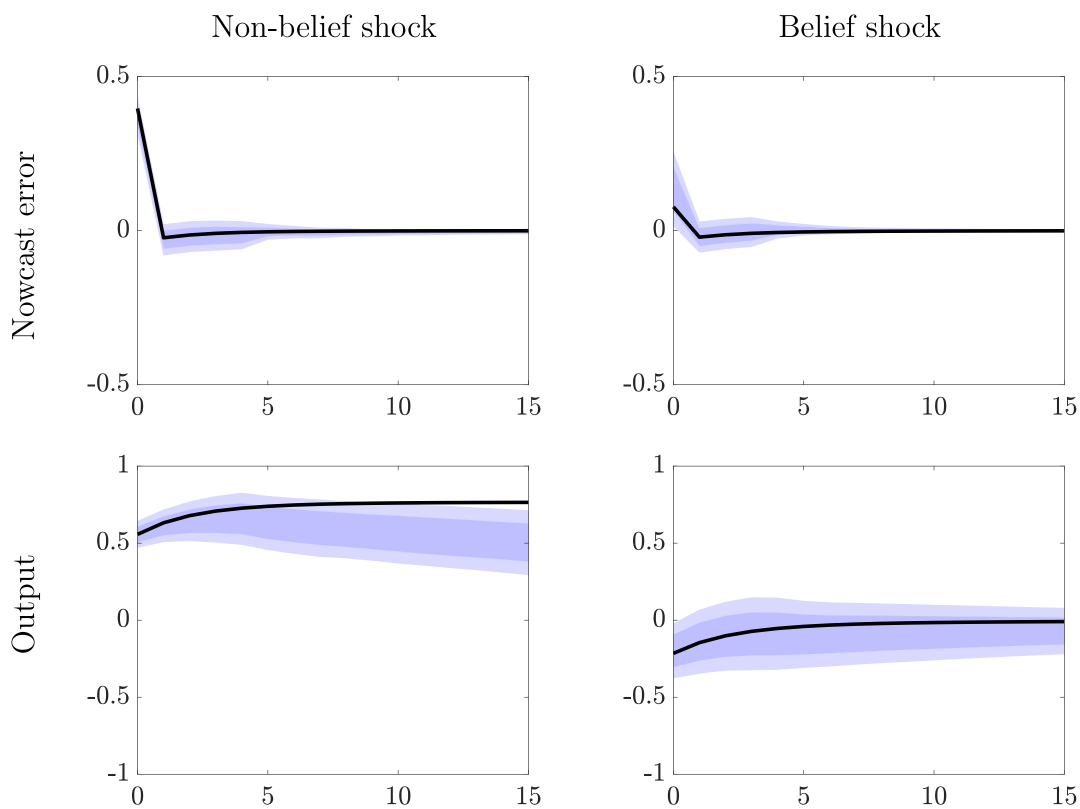

Figure A.1: Impulse responses in the model of Lorenzoni (2009) and according to the VAR, estimated on simulated data. Notes: solid line represents true response. Shaded areas represent 68% (90%) confidence bounds, median over 1000 estimates. Sample comprises 185 observations, Section 3.2 provides details on VAR specification.

Appendix B: Further VAR evidence

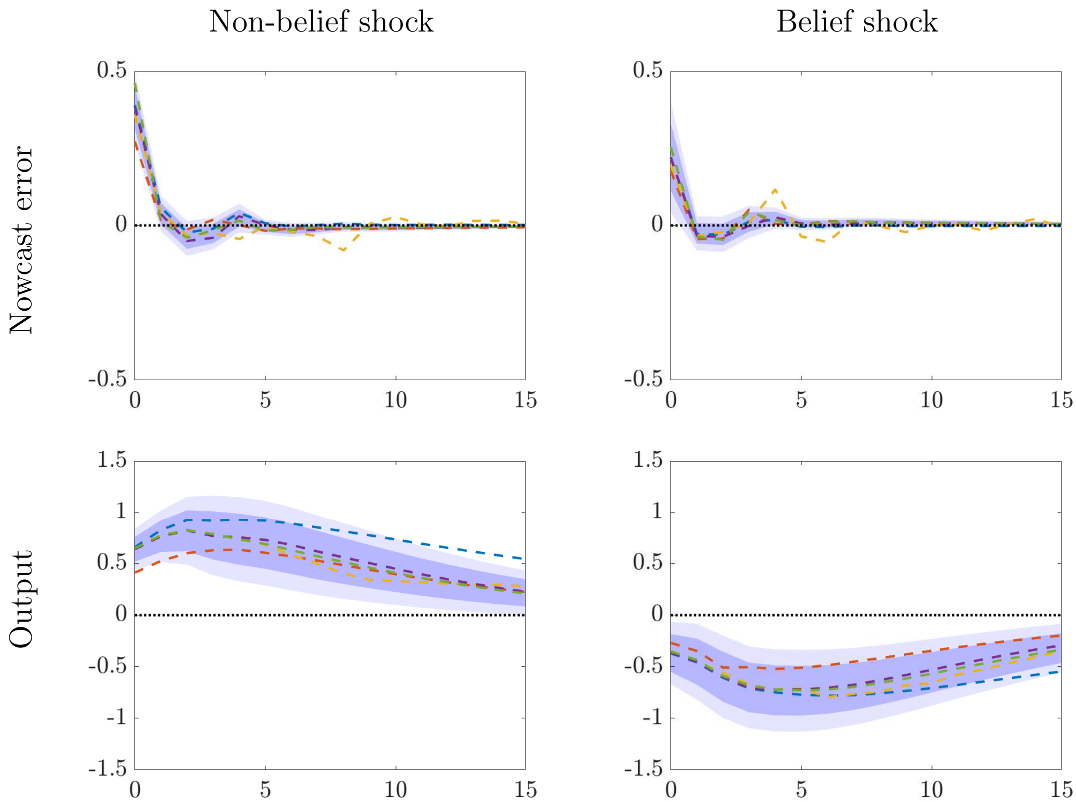

Figure B.1: Sign restrictions: robustness. Blue: linear trend; orange: short sample; yellow: 8 lags; purple: mean nowcast; green: final-release nowcast. Shaded areas indicate 68% and 90% confidence intervals of baseline specification.

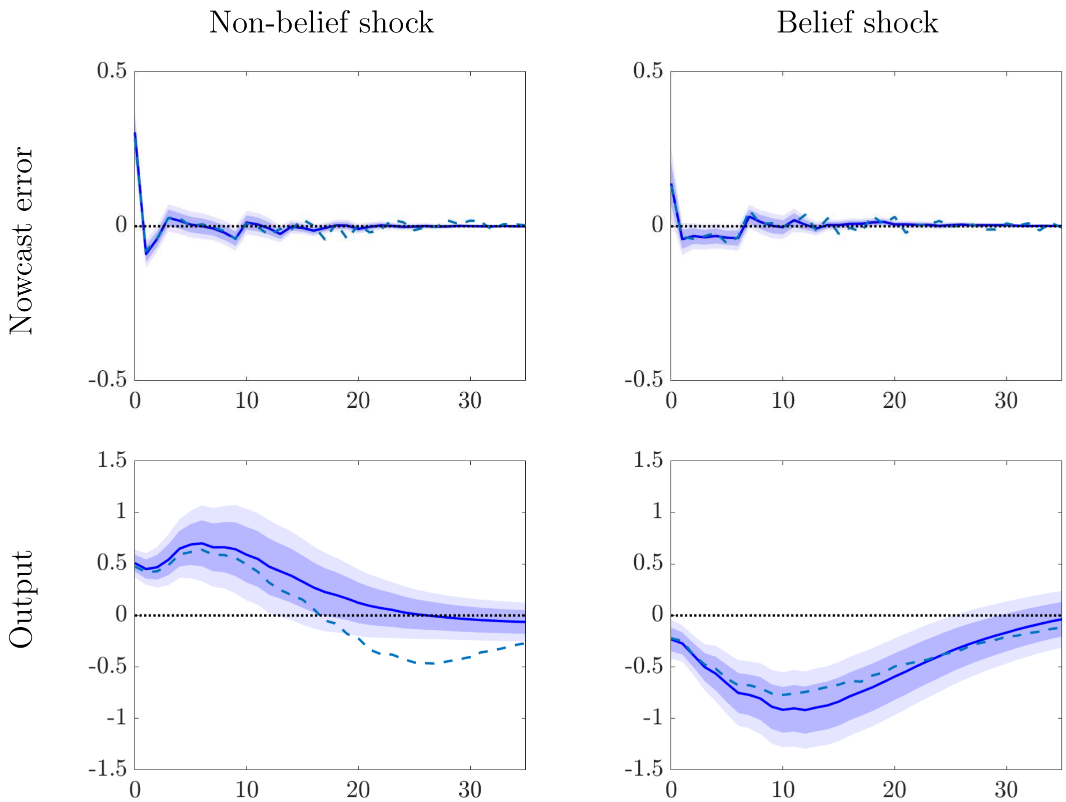

Figure B.2: Sign restrictions: monthly VAR. Blue solid lines: linear trend and 12 lags; turquoise dashed lines: 24 lags. Shaded areas indicate 68% and 90% confidence intervals of first specification.