Deconstructing Uncertainty

Xiaohan Ma[*] Roberto Samaniego[†]

July 11, 2018

Abstract

We develop aggregate and industry-level measures of uncertainty using the absolute values of median forecast errors drawn from a large firm-level dataset. First, we confirm that our aggregate measure of uncertainty (using all firms) behaves much like other uncertainty measures in the literature. Then we capture uncertainty within particular sectors of the economy using selected subsets of the forecasts. We find that industry uncertainty measures share a common factor that closely follows aggregate uncertainty, while also containing sector-specific information. We explore the economic impact of various sector-specific measures, finding that uncertainty measured among financial firms appears to precede uncertainty measured in the rest of the economy, and has greater economic impact.

Keywords: Uncertainty, business cycles, forecast errors, financial uncertainty, industry uncertainty.

JEL Codes: D80 E32 E37 E47 E71 G17.

[*]Texas Tech University, Department of Economics, P.O. Box 41014, Lubbock, Texas 79409-1014 U.S.A. Email: xiaohan.ma@ttu.edu.

[†]Department of Economics, The George Washington University, 2115 G Street, NW Monroe Hall 340, Washington, DC 20052, U.S.A. Email: roberto@gwu.edu.

1 Introduction

The macroeconomic impact of uncertainty remains a topic of active research. At the same time, several important research questions remain unanswered. Is uncertainty a macroeconomic phenomenon? Or is uncertainty something that is mainly specific to particular industries or sectors?[1] If so, is aggregate uncertainty just a reflection of what is happening (on average) across different industries? Alternatively, is uncertainty something that arises in certain sectors and then spreads to the rest of the economy through specific channels, such as input-output links?

Answering these questions require industry-level measures of uncertainty. The purpose of this paper is to address these questions by developing a new approach to measuring uncertainty, one which allows us to measure both aggregate uncertainty and uncertainty in selected subsectors of the economy in a similar fashion. Specifically, we develop a measure of uncertainty using millions of *analyst stock forecast errors*, drawn from the I/B/E/S database.

First, using forecasts across all economic sectors, we obtain a measure of aggregate uncertainty. We find that this measure behaves very similarly to other measures of aggregate uncertainty in the literature:[2] in particular, in a vector autoregression estimation (VAR) with uncertainty and other macroeconomic variables, an increase in the aggregate uncertainty generates a downturn in GDP, as well as inflation and aggregate inputs.

As our benchmark, we use 12 month-ahead earnings-per-share (EPS) forecast errors. We use these for several reasons. One reason is that this is the most complete and the most widely available series: we are able to develop a continuous monthly series starting in the early 1980s. We use the 12 month forecasts because, as well as analyst- or stock-specific factors of uncertainty, these measures should reflect uncertainty at business cycle frequencies. We use EPS forecasts because EPS ratios are a basic indicator of the profitability of a share, and are thus widely understood

[1] The business cycle literature finds evidence that sectoral shocks can have an important aggregate effects, e.g. Horvath (1998, 2000) and Garín et al (2018). Whether this is the case for uncertainty is unknown.

[2] See for example Bloom (2009), Bachmann et al (2013), Jurado et al (2015) and Scotti (2016).

and followed both by financial analysts and their clients. EPS forecast errors that are due to *idiosyncratic* factors should be unrelated to each other, and presumably should not vary much over time, representing a sort of background noise. On the other hand, a significant change in this background level of noise indicates an increased level of uncertainty overall. We focus on the absolute value of the *median* forecast error within any given month, to avoid being swayed by individual outliers. Besides, it ignores the direction in which analysts typically made an error. We also check several variants of our basic measure for robustness – including measures that are scaled by share prices, so that they are interpretable as forecasts of price-earnings ratios, and the median forecast errors *one, two or three quarters* ahead, which can serve as potential leading indicators of the impact of uncertainty on future aggregates.

We find that our baseline measure of uncertainty has several useful properties. First, the series is fairly flat over time, but subject to occasional large spikes. This is what we would expect to see in an environment with occasional uncertainty shocks. Moreover, these spikes correspond well to NBER recessions as well as other notable events of economic significance. When we include our measure in a recursively identified VAR, we find that uncertainty shocks have a significant and persistently negative impact on industrial production as well as other aggregate variables. The behavior of our baseline index is similar to that of several other measures of uncertainty found in the literature, suggesting that we are indeed capturing similar macroeconomic impulses in a new way. In particular, by leading to a decrease in output, inputs and the price level, an increase in uncertainty behaves like a negative demand shock[3]. The contribution of our measure of uncertainty shocks to aggregates is also among the strongest of the measures considered, especially when normalized by the share price.

Then, by using only subsets of the forecasts, we can obtain an uncertainty measure for that specific industry or sector. This is possible by merging I/B/E/S forecast data with CRSP, which contains SIC codes for most of the firms in the database. We perform 3 exercises, to assess in general whether industry specific uncertainty

[3]The idea that uncertainty shocks behave as negative aggregate demand shocks is explored in Leduc and Liu (2016), and Basu and Bundick (2017), among others.

measures themselves have significant macroeconomic impact, and to see whether any particular industries are key for generating or transmitting uncertainty to the rest of the economy:

1. measuring uncertainty inside and outside the financial sector. This is to examine whether the financial sector – where profits are made and lost based on forecasts about the future profitability of economic actors – is an early warning indicator of aggregate uncertainty, or whether it is more related to the stock market and less related to real economic activity;[4]

2. measuring uncertainty inside and outside "hub" industries, as identified using input-output tables – specifically, wholesale trade, iron and steel, and management. This is to investigate whether uncertainty might arise from these industries and spread through input-output relationships;

3. measuring uncertainty among 13 non-overlapping sectors of the private economy, roughly corresponding to the 1-digit SIC level of aggregation. This allows us to study whether there are any *common factors of uncertainty* measured across the different parts of the economy, and to see whether these measures *on their own* have any macroeconomic impact.

First, we find that financial uncertainty has its main impact on the stock market, whereas uncertainty measured outside the financial sector mainly affects real variables rather than financial variables. The impact of financial uncertainty on industrial production is also greater. A cross-correlogram also finds that financial uncertainty precedes non-financial uncertainty. This complements the finding of Ludvigson et al (2018), who find that financial uncertainty shocks cause both economic downturns and increases in aggregate measures of uncertainty.

When we look at hub-industries, we find that hub uncertainty does not always have a significant independent impact on output, suggesting that these industries are

[4]Some authors have studied the impact of financial uncertainty such as Shin and Zhong (2018), Carriero et al (2018), Ludvigson et al (2018), among others. However, the measures they use are based on macroeconomic series that are thought of as being related to financial information, whereas ours are based on data regarding firms that operate in the financial sector.

not key sources of uncertainty and that in general uncertainty likely does not spread through input-output relationships.

Finally, when we decompose the economy into 13 sectors, we find that there is a common factor which strongly co-moves with uncertainty measured in the economy as a whole. At the same time we also find that there are some sectors where measured uncertainty behaves quite differently – particularly agriculture, construction and utilities. In general, with the notable exception of the financial sector, we do not find that these sector-specific uncertainty measures have weaker impact on the macroeconomy. We conclude that there is such a thing as aggregate uncertainty, interpretable as a common factor of uncertainty across sectors, and that it is related to uncertainty about (or arising from) firms in the financial sector.

Section 2 describes our measure of uncertainty and its relation to the literature. Section 3 describes in detail the data that we use, and some basic properties of the measure. Section 4 describes the impact of our measure of uncertainty on macroeconomic dynamics, and compares it to other measures of uncertainty, as well as variants of our basic measure. Section 6 examines industry-specific measures of uncertainty. Section 7 concludes with a discussion of potential future work. The robustness of results is explored in Appendix.

2 The Measure

The premise behind our measurement strategy is that changes in uncertainty, and thus the predictability of the economic environment – at the aggregate or at the industry level – will be reflected in the analyst forecasts being of lower accuracy than usual. We use this idea to develop a measure of uncertainty, and study its relationship to macroeconomic aggregates to examine whether it is indeed a useful measure of uncertainty for purposes of understanding the macroeconomic impact of uncertainty. Then we adapt it to measuring uncertainty at the industry level.

Time is discrete and divided into days which are collected into months: specifically, let $M \subset \mathbb{N}$ be the set of months, numbered consecutively, and and let $t \in M$ be a month. Then, define $D_t \subset [t, t+1)$ as the set of days in the month t, so that

$d \in D_t$ represents a day in month t. Let $S_{i,t}$ be a statistic about a firm i observed on day d, and let $F\left[S_{i,d^*}|I_{j,d}\right]$ be the forecast about the realization of statistic S at firm i on a future day d^*, using information set $I_{j,d}$ available to them on day d – so that d^* minus d is the forecast horizon. Note that d^* will not be in the same month if the forecast horizon is longer than a month, as will be the case in our data. We define the firm-level forecast error as the difference between the forecast made on day d about statistic S at date d^*, and the actual realization of the statistic on day d^*:

$$FE_{i,d} = F\left[S_{i,d^*}|I_{j,d}\right] - S_{i,d^*}. \tag{1}$$

In our benchmark measure, the forecast period is a year, but we also look at quarterly forecasts. If more than one analyst makes a forecast about firm i on day d, we define $F\left[S_{i,d^*}|I_{j,d}\right]$ as the average forecast error made about firm i on day d.[5]

There are thousands of forecasts made every day about different firms. To measure uncertainty U_t in month t, we will focus on the uncertainty experienced by a typical firm. In particular we look at the absolute *median* forecast error across all firms within the month. We focus on the median in order to avoid being swayed by individual outliers, which a large data set of forecasters will inevitably have.[6] In addition, we define uncertainty based on the *absolute value* of the median forecast error. This way uncertainty is measured as lack of accuracy – regardless of the direction. Not doing so would lead to a measure of relative optimism or pessimism compared to the realization, not uncertainty. As discussed below, we try several other approaches for robustness. In practice, all our measures will be monthly, the highest frequency for which we have data on industrial production. Thus, on each date within month t, we compute the absolute median value of $FE_{i,d}$ within the month pooling all firm-day forecasts within the month, which gives our baseline uncertainty measure for

[5] In practice, only about 14 percent of all firm-day combinations have more than one analyst making a forecast about it, ranging from 2 to 5 analysts. For robustness, we also repeat this procedure looking at the average forecast about a particular firm *within a month*, in which case about 87 percent of firm-month combinations have more than one forecast, ranging up to 42. Results are similar, see Appendix.

[6] Indeed, we found that uncertainty measured using the mean rather than the median was extremely volatile and had no meaningful properties.

month t:[7]

$$U_t = \| median \{FE_{i,d} : \forall d \in [t, t+1)\} \| .$$

The specific statistics that we look at are forecasts of the earnings-per-share ratios (EPS) of individual companies. We use EPS forecasts because they are the most widely available in our database, and also because EPS ratios are a basic indicator of the profitability of a share, and are thus widely understood and followed both by financial analysts and their clients.

One might wonder whether realized error necessarily implies that an agent was uncertain at the date the forecast was made – although naturally uncertainty and ex-post errors would be expected to be correlated. This concern, however, is not important for our measurement strategy. Our measure is based on millions of forecasts about thousands of firms by thousands of analysts. It is not based on the assumption that when a particular forecaster makes a mistake this indicates the forecaster is uncertain. It is based on a large number of forecasters making *systematic* mistakes. Since the analysts form a large decentralized system of interacting agents, we think of our measure as being related to the literature on swarm intelligence.[8]

2.1 Relationship to other measures

Before moving on to the statistical analysis of our uncertainty measure and its economic impact, we discuss its relationship to other measures. Naturally our aim is to complement other measures of uncertainty, explore their robustness to other measurement approaches and, in particular, develop a measure that can be used to study uncertainty in different subsectors of the economy. Still, our strategy has certain conceptual and practical advantages.

The literature contains several approaches to measuring uncertainty. One approach is to measure uncertainty using the second moment of macroeconomic variables, such as GDP growth, or stock price as in Bloom (2009). A downside of this

[7]An alternative would be to take median over absolute forecast errors, in other words to define $FE_{i,d} = \| F[S_{i,d^*}|I_{j,d}] - S_{i,d^*} \|$. Another alternative would be to take means instead of medians. We discuss these alternative measures later and in Appendix E.

[8]See Beni and Wang (1989).

approach is that the observed variability of variables does not necessarily correlate with the extent of *uncertainty*.[9]. Thus, a second approach has been to measure uncertainty using measures of *unpredictability* – essentially using ex-ante or ex-post forecast errors or functions thereof, as in Bachmann et al (2013), Jurado et al (2015), Scotti (2016). A third approach has been to use textual analysis or surveys to determine whether agents at a given point in time feel uncertain about current or future economic conditions, as in Baker et al (2016) and Bachmann et al (2017) .

Our measures are first moments, not second moments, so they are distinct from the first approach. Instead, our work straddles both the second and third approaches. We follow the second approach by measuring uncertainty as the unpredicted component in forecasts made by agents. In this sense we are also related to the third approach in that we use the behavior (in our case, the predictions) of economic agents to measure uncertainty. The third approach, including survey- and text-based measures, depends on the self-reported views of economic agents or the views of journalists about future economic conditions, and may suffer from various types of response bias, including the agent's response to spoken or unspoken cues from the surveyor or framing effects that could contaminate the measures.[10] In our case, however, there is no survey, only a datum to forecast, so there is little scope for framing effects.[11] In addition, the agents we use are professional forecasters employed in the financial sector. These are analysts who are making stock recommendations to investors, and thus have a financial incentive to make accurate forecasts: their reputations and livelihood are on the line[12].

[9]For the first example, a change in a variable does not necessarily mean it was not predicted. For the second example, in a world of perfect foresight but different cyclicality of certain activities or companies, dispersion in forecasts will be observed before a recession because of this asymmetry, even with perfect foresight.

[10]See Tversky and Kahneman (1974) regarding how people develop or report subjective probabilities of events in potentially biased ways, and Kalton and Schuman (1982) for a survey of response bias in self-reported surveys.

[11]So (2013) argues that analysts may have an incentive to bias their forecasts. However, as long as this bias does not systematically co-move with uncertainty this poses no problem for our measurement strategy, as it is equivalent to analysts adding a constant to their forecasts. In addition, such bias should not affect the results of our VAR either if it does not cause movements in output.

[12]Hillary and Hsu (2013) find that analysts whose forecasts are more informative are more likely

Although our measures of uncertainty follow closely the second approach, there are several differences compared to previous studies in this strand. According to our strategy, no analyst is making a forecast of aggregate uncertainty: rather, the measure of uncertainty we develop is immanent in their collective behavior. Nonetheless, as we shall see, our uncertainty measure turns out to behave quite similarly to the Jurado et al (2015) measure that is based on a variety of aggregate series, even though the measurement strategy is totally different and even though there is no overlap in the data inputs used to construct the measures. Put another way, unpredictability in aggregate time series and unpredictability in analyst forecasts about particular firms turn out to have similar behavior. This is informative to the literature on uncertainty as it implies that microeconomic uncertainty and uncertainty about aggregate fundamentals are linked.[13]

In addition, an issue with measures that use the second approach is that forecasters are known to converge towards a consensus forecast. This herding behavior among forecasters is well known and may affect any environment where many forecasters are trying to forecast the same thing: they may be ignoring private information as they herd, or may be converging (in informational terms) towards a local rather than a global optimum. In our case, the forecasters are forecasting different things, so they cannot herd regarding aggregate uncertainty.

Moreover, the details of the forecasts used vary in important ways between ours and other papers that follow the second approach. Bachmann et al (2013) use forecasts of business conditions that are qualitative ("up" or "down") that may introduce some imprecision, whereas our forecasts are quantitative and continuous. They construct uncertainty measures using both ex-ante disagreement among forecasts and ex-post forecast errors and show that these measures can have similar properties. In contrast, we focus on the median forecast error, not dispersion, as the measures of disagreement based on our data did not have useful statistical properties. Finally, Jurado et al (2015) and Scotti (2016) construct uncertainty measures using the de-

to be rewarded and/or promoted.

[13] Jurado et al. (2015) also show that microeconomic uncertainty, measured as the deviation of unpredicted profit growth of firms from the mean, behaves in a similar manner to the macroeconomic uncertainty measure in their paper.

viation (conditional variance) of forecasts from the sample mean of some series, such as a plethora of time series. In contrast, we impose no statistical structure on any moment of our uncertainty measures U_t.

3 Data

Our forecasts are drawn from the Institutional Brokers' Estimate System or I/B/E/S, available through a WRDS subscription and managed by Thomson Reuters. It contains analyst forecasts of several measures of interest to investors and researchers, the most widely-available being earnings per share (EPS) forecasts. I/B/E/S also reports realizations of the forecast data, collected from a variety of public data sources. Companies are included in the database as long as at least one analyst provides a forecast for that company. Forecasts are not included unless they are confirmed within 6 months.[14] Forecasts are collected each day as they are released by analysts.[15]

We focus on US firms. This yields about 4.5 million forecasts issued by $1,523$ different brokers, who make forecasts about many firms over time. For each firm on each day we compute the average forecast error.[16] Then we take the median forecast error across firms within each month, starting in September 1981.[17] Thus it is the median forecast error by firm-day pair. We deflate these measures by the monthly CPI in order to ensure our measures are reported in real terms.

This is our baseline measure: as discussed later, we also examine the behavior

[14] For further details, see https://wrdsweb.wharton.upenn.edu/wrds/support /Data/_001Manuals%20and%20Overviews/_003I-B-E-S/Release%20Notes/, last checked 3/20/2018.

[15] Later we also look at EPS forecasts made over different horizons. No forecasts for different horizons are made about the same firm on the same day by the same analysts. However, all analysts that make an annual forecast make a quarterly forecast about a given firm sometime that month. About 46 percent of forecasters who make an annual forecast about a given firm make a 2-quarter ahead forecast the same month, and about 39 percent for 3-quarter ahead forecasts.

[16] 86 percent of them are single forecasts about a firm on a given day. The rest have 2 forecasters making forecasts about a firm on the same day, except for 0.29 percent of the sample which has $3-5$ forecasts. Averaging when there are multiple forecasters yields about 2.9 million day-firm observations.

[17] This is the first month after which continuous series may be computed. The date is based on the month and year of the variable *anndats*.

of a variety of similar measures for robustness. In addition, we later develop some uncertainty measures that are scaled using the share price of each firm, or which focus on firms from a particular industry. For some purposes we use the raw series, for others we use the HP-detrended series using $\lambda = 129,600$ as recommended for monthly data by Ravn and Uhlig (2002).

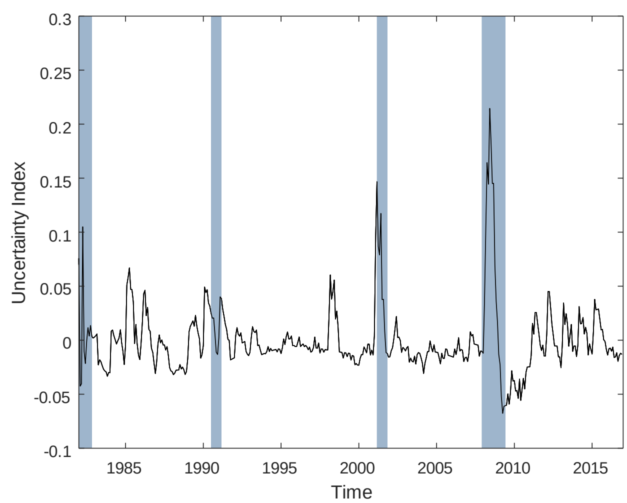

Figure 1 – Uncertainty and the business cycle, 1981-2016. Bands represent NBER recession dates. The measure is the absolute value of the median forecast error from I/B/E/S by month, HP-filtered. The forecast error is the difference between the 12-month EPS forecast and the realized EPS. The measure in the Figure is deflated using the CPI and HP-filtered following Ravn and Uhlig (2002).

One might ask whether we should *scale* our forecast errors. This is because the number of shares of any particular firm may not bear much relation to the size of the firm and therefore the magnitude of EPS forecasts and realizations. Not scaling our measures means that we are looking at a typical firm in terms of share size, but we may also wish to use share prices as a way of scaling the EPS forecast errors.

For robustness, we also scale our measures by the price of a share on the date that the forecast was issued ($FE_{i,d} \div P_{i,d}$, where $P_{i,d}$ is the share price), or by the price of a share on the date that EPS are realized ($FE_{i,d} \div P_{i,d^*}$). This is similar to using a forecast of the (inverse) price-earnings ratio rather than the EPS ratio – except that of course the price is a realized price, not a price forecast. These measures would make an adjustment for the fact that shares of particular firms vary in size relative to operations or value. While our baseline measure does not make this adjustment, we do examine these measures in Appendix A, finding that they generally behave much like our standard measure but with some differences. We obtain these prices and industry codes from CRSP.[18]

We also explored two alternative approaches to measuring uncertainty using our forecast database. One was to take the absolute values of forecast errors *before* computing medians. Another was to use means rather than medians. We found that both these approaches had shortcomings. The first approach displayed some variation that was systematically related to certain months of the year. The second suffered from the impact of large outliers. The existence of large outliers cannot be overcome easily by just dropping unusually large or small values, however, as outliers could be concentrated in episodes of high uncertainty (and thus a procedure to clean outliers would in fact remove the shocks we are trying to identify). As a result we do not use these approaches in the analysis that follows. That said, when we used monthly dummy variables to remove monthly effects (and, in the case of the mean-based measure, eliminated the top and bottom 5 percent of observations to remove outliers), we obtained time series that were highly correlated with the series in Figure 1, as well as having similar macroeconomic implications.[19]

Figure 1 displays the detrended baseline series. Several observations stand out. First, the series appears to have a more or less stable level of uncertainty, punctuated by sharp spikes. This is consistent with the notion that there is a background level of systemic uncertainty which is subject to occasional shocks. Second, these spikes

[18] CRSP reports the NAICS and SIC codes of these firms. We use SIC codes because NAICS codes did not exist early in our sample.

[19] See Appendix E.

tend to coincide with the beginning of NBER recessions (or, in the case of the 1990 recession, to precede them). Notably, all the recessions coincide with increases in our uncertainty measure. Also several notable economic events the literature tends to think of as being significant and unpredicted are clearly visible, including the monetary uncertainty of $1981-1982$, the Asian Crisis + LTCM collapse of $1997-8$, the tech bust of 2001 and the financial crisis of late 2008.

4 Macroeconomic Impact of Uncertainty

Existing empirical research on uncertainty has often found important dynamic relationships between real activity and various uncertainty proxies. In particular, these proxies are often countercyclical and VAR estimates suggest that they have a significant impact on output and employment in the months after an innovation in these measures. A key finding is that a rise in some proxies of uncertainty – notably stock market volatility – depresses real activity in the short run, consistent with the predictions of some theoretical models where uncertainty is a driving force of macroeconomic fluctuations.

We now use VARs to investigate the dynamic responses of key macro variables to innovations in our uncertainty measures. For brevity in discussing the results, we will often refer to these innovations to uncertainty as *uncertainty shocks*. As is the case in all VAR analyses, the impulse responses and variance decompositions depend on the identification scheme, which in our case is based on the ordering of the variables.

For the benchmark analysis, we choose a specification similar to that studied in Bloom (2009), as to which variables to include in the VAR and how to order the variables. Following Bloom (2009), we use 12 lags of monthly data of the log S&P 500 index, federal funds rate, log wages, log CPI, log hours worked in the manufacturing sector, log employment for the manufacturing sector, and log industrial production between 1981 and 2016. The macroeconomic dynamics of these variables have been extensively studied in the literature. In addition, the variables are ordered as follows in the benchmark VAR:

$$\begin{bmatrix} \text{log (S\&P 500 Index)} \\ \textit{uncertainty} \\ \text{federal funds rate} \\ \text{log (wages)} \\ \text{log (CPI)} \\ \text{log (hours)} \\ \text{log (employment)} \\ \text{log (industrial production)} \end{bmatrix}$$

Placing the stock market index first guarantees that the effect of stock market changes is already taken into account when examining the effect of uncertainty. Unlike Bloom (2009), we do not HP filter the variables before estimating the VAR, but rather use all the variables in levels.

4.1 Results

Below we report the results of the VAR with the baseline measure, . In Appendix *A* we report results using some changes to the VAR specification, including a VAR with uncertainty placed last, a VAR with data spanning from 1983 to 2016, and a VAR with real variables only.

Figure 2 plots the impulse responses of federal funds rate, price level, industrial production, manufacturing employment, hours worked, and stock market index to innovations in our baseline uncertainty measure. The innovation is a one-standard deviation shock to uncertainty. The shaded area is +/- one standard error confidence bands estimated from the VAR system. A positive uncertainty shock leads to a significant decrease in macroeconomic activity. The maximum decline in production is 0.6 percent, occurring roughly 14 months after the initial shock. There are also declines in the intensive and extensive margins of manufacturing labor following an increase in uncertainty. Manufacturing employment falls with the peak impact of 0.5 percent occurring about 16 months after the initial shock. Hours worked declines to a trough of 0.15 percent below trend 10 months after the initial shock. After the initial decline, hours worked rebound back to the pre-shock path after about 2 years, whereas industrial production and employment both remain lower than the pre-shock level even after 5 years. The price level declines for roughly 24 months following the uncertainty shock. The fact that uncertainty reduces output – as well as prices – is what one would expect from an unexpected negative demand-side shock. The falls in production and inflation lead the monetary authority to decrease the federal funds rate by up to 0.18 percentage points within 12 months.

We do not find "volatility overshoot" in our results, which is shown in Bloom (2009) as a rebound in output and employment following the initial decline from an uncertainty shock measured using stock market volatility. This finding is on par with Bachmann et al. (2013), Jurado et al. (2015), and Scotti (2016)[20]. In Bloom (2009), the impulse responses are generated by a 15-point shock to the error in the stock market volatility equation, which is roughly equivalent to a four standard deviation shock. If we calculate the impulse responses to a four standard deviation shock to our measure of uncertainty, the maximum decline in production would become 2.5 percent, very large. This impact is more protracted than that suggested in Bloom (2009) and similar to the finding in Jurado et al (2015).

[20] One explanation is that the "overshooting" effect in Bloom (2009) is sensitive to the sample period used. See, for example, Choi (2013).

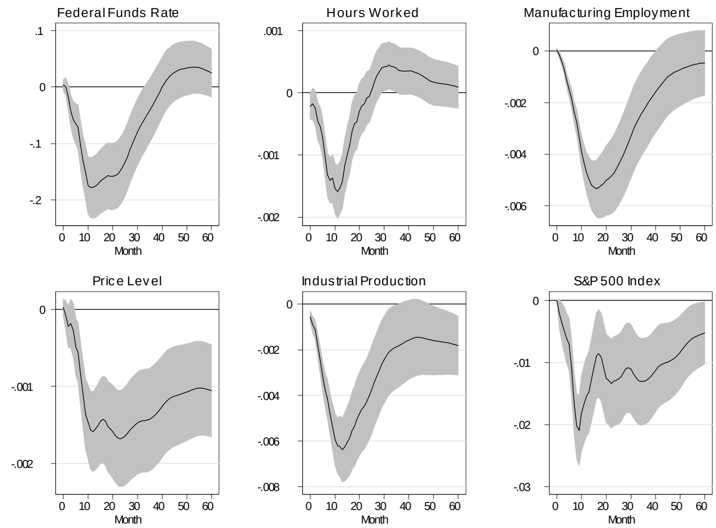

Figure 2 – Impulse response of federal funds rate, price level, hours worked, industrial production, manufacturing employment, and stock index from estimation of benchmark VAR, with uncertainty measured using the 12-month absolute forecast error. The shock is one standard deviation, and the error bands are +/- one standard error confidence bands. The sample period is 1981m9 to 2016m12.

4.2 Comparison with other uncertainty indices

Figure 3 plots the response of industrial production to six different uncertainty indices. These are: our baseline uncertainty measure; monthly stock market volatility as in Bloom (2009); the forecast disagreement index constructed using the Philadelphia Fed's Business Outlook Survey by Bachmann et al. (2013); the uncertainty index based on the implied forecast errors for real economic activity derived from economic and financial series in Jurado et al. (2015); the real-time real-activity uncertainty index in Scotti (2016); and the Google News uncertainty index measured as a weighted number of articles mentioning "uncertainty" by a number of articles containing the word "today", constructed by Bachmann et al. (2013) that uses the method in Baker et al. (2016). The impulse responses are obtained from separately estimating a VAR system as the benchmark system, but with different uncertainty measures. Qualitatively, the response of production to all six measures is similar, except for the measure in Scotti (2016)[21]: when hit by a positive uncertainty shock, output declines according to the other measures. Quantitatively, the response to our baseline uncertainty is more persistent and significant than that implied by the uncertainty indices constructed by Bloom (2009) and Google News, and comparable to that in Bachmann et al. (2013) and Jurado et al. (2015).

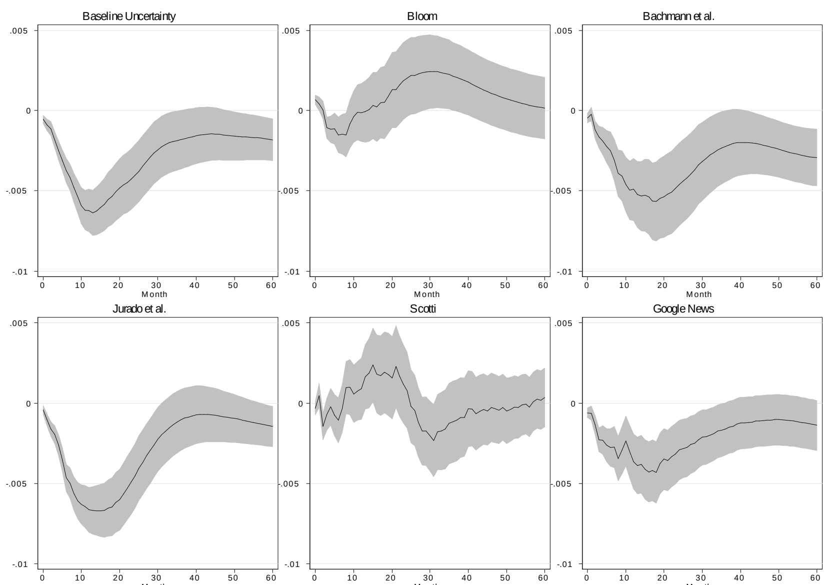

Figure 3 – Impulse response of industrial production with different uncertainty measures.

To study and compare the quantitative importance of uncertainty shocks as a

[21]Scotti (2016) constructs the uncertainty index based on real-time and real-activity related to the state of the economy. We take an average of the daily uncertainty index to get a monthly index.

driving force of business cycle fluctuations, we also calculate the forecast error variance decomposition for production based on each uncertainty measure. The variable h denotes the forecast horizon for the decomposition. Table 1 shows the fraction (as a percentage) of the VAR forecast error variance contributed by various uncertainty shocks over horizons of 3, 12, 36, and 60 months. The results suggest that innovations in our baseline uncertainty measure account for about 26 percent of the forecast error variance of production at horizons of $1-5$ years, between Bachmann et al. (2013) and Jurado et al. (2015). By contrast, the variance attributed to innovations in Google News is only about half, and that attributed to Bloom (2009) is only about one eighth.

Horizon	Baseline	Bloom	Bachmann	Jurado	Scotti	Google
$h=3$	3.33	0.71	1.87	5.19	3.96	2.90
$h=12$	27.06	1.28	11.07	34.44	1.88	9.79
$h=36$	26.27	3.22	21.18	33.64	4.59	16.28
$h=60$	25.53	3.63	21.24	30.71	3.40	12.97

Table 1 – Industrial Production Forecast Variance Due to Uncertainty (in percent)

Figure 4 also compares the time series for each of these measures over the period during which they overlap. For brevity we focus on 3 alternate series, one of each type of measure identified in Section 2.1: the Bloom (2009) series, the Jurado et al (2015) series and the Bachmann et al (2013) series. Visually it is clear that the Jurado et al (2015) series has the most comovement with baseline uncertainty: the correlation between them is 0.29 and is significant at the one percent level, whereas that between our measure of uncertainty and the Bachmann et al (2013) series is 0.24 (similar significance) and that with the Bloom (2009) series is only 0.04 and not statistically significant. It is interesting that all the measures rise around the beginning of the Great Recession (although the Bachmann et al 2013) series reacts only weakly), and our uncertainty measure reacts earlier than the others.

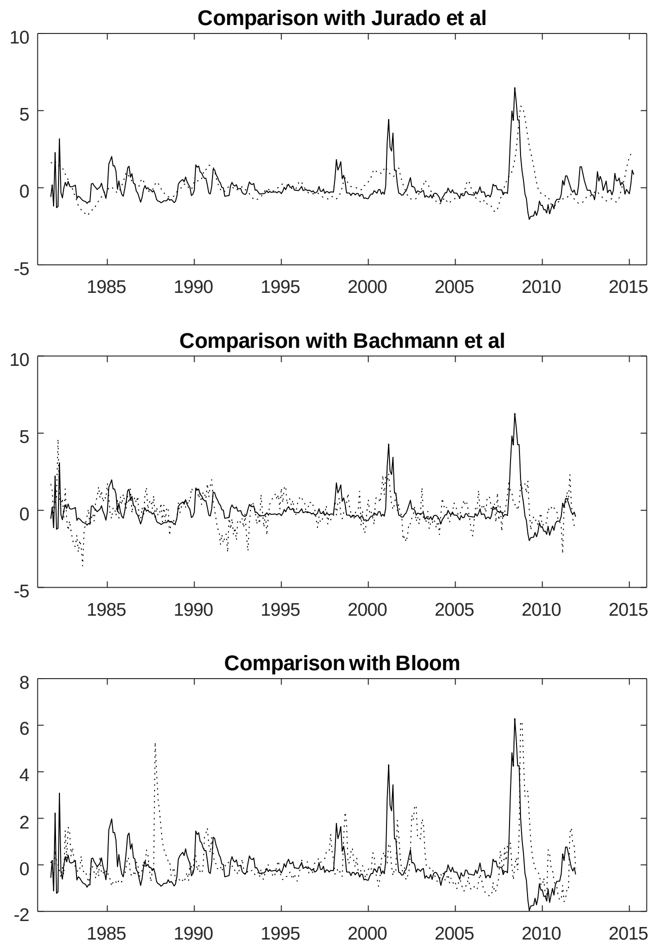

Figure 4 – Comparison of different uncertainty measures. In each panel the thick line is our uncertainty measure and the dotted line is an alternative measure, one of each type. All measures are Hodrick-Prescott filtered with λ set to $129,600$.

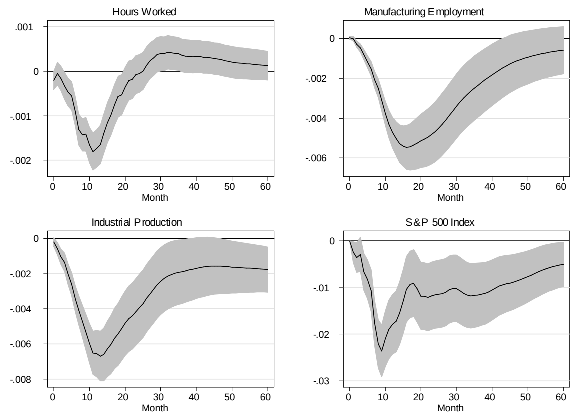

Figure 5 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with forecast errors measured by EPS/Price1

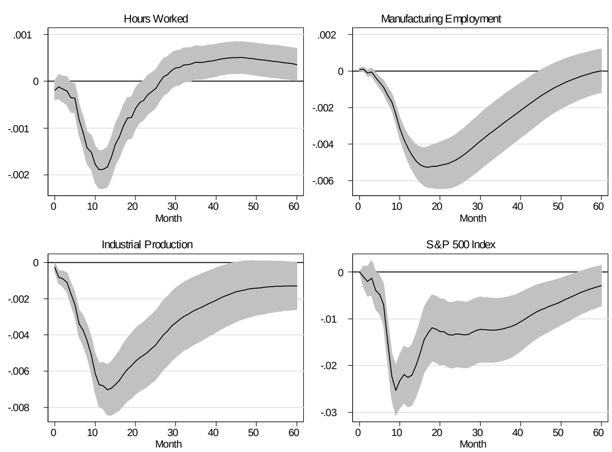

Figure 6 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with forecast errors measured by EPS/Price2

5 Robustness: Alternative measures

Scaling EPS ratios by share prices

A concern with the baseline measure is that some variation in EPS ratios could be due to the fact that firms have different scales – or rather that the granularity of their share size may vary. This suggests that we might wish to adjust our measures by *share prices*. Conceptually, such a measure would have the interesting property that it could be interpreted as a forecast of inverse price-earnings ratios. To produce such measures, we combine the I/B/E/S and CRSP data to create a new uncertainty measure, the EPS forecast error divided by the corresponding security prices. EPS/Price1 is the price-weighted forecast error calculated by using the security prices observed when the forecast is made, whereas EPS/Price2 uses the security prices observed when the forecasted EPS is realized. Figures 5 and 6 show the two scenarios. The impact of uncertainty shocks is still significant on production – indeed, the peak impact after about 14 months is quite a bit *larger* than with the baseline measure. It is interesting that an uncertainty shock by this measure does not affect industrial production on impact: rather, it appears to take one or two months.

Measures with various forecast horizons

Our uncertainty measures so far are useful for understanding the macroeconomic impact of uncertainty, but since they are ex-post statistics based on the realized EPS one year after the forecast is made, we cannot use the year-ahead measure as forecasting tools for the immediate future, i.e., during the 12 months between when the forecast is made and when the forecasted variable is realized. It is interesting, therefore, to look at whether a similar uncertainty measure using I/B/E/S survey can serve as a leading indicator of the impact of uncertainty on business cycle fluctuations. To this end, we construct some alternative uncertainty measures, using the median of the *2 and 3 quarters ahead* absolute forecast error.[22] That is, the uncertainty about the current month's economic situation is measured as the absolute difference

[22] We also have 1-quarter ahead forecast errors. These do not turn out to have any useful properties.

between the forecasted EPS and the realized EPS 2 quarters and 3 quarters later. If these measures have an impact on industrial production then it is something that can be detected farther in advance, as it takes less than one year to define the measure[23].

Figure 7 shows the impulse responses to a one standard deviation positive shock to these uncertainty measures. Although quantitatively weaker than the baseline measure, an increase in the uncertainty based on 2- and 3-quarter ahead absolute forecast errors does lead to a decrease in the production, hours worked, employment, and the stock index. This suggests that these measures could potentially serve as forecasting indicators.

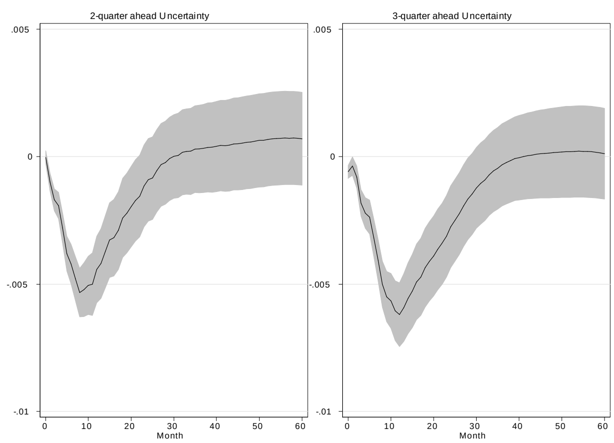

Figure 7 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with quarterly uncertainty measures

[23] Using quarter-ahead measures as forecasting tools allow us to incorporate more recent information. For example, suppose we are in 2018m4, if we want to forecast the industrial production in 2018m5, the last available data point of the year-ahead uncertainty measure would be *2017m4* given the definition of our uncertainty measure (the uncertainty in 2017m4 equals to the absolute difference between the forecasted EPS *made in 2017m4* for 2018m4 and the realized EPS in 2018m4), whereas the last available data point of the two-quart-ahead uncertainty measure is *2017m10* (the uncertainty in 2017m10 equals to the absolute difference between the forecasted EPS *made in 2017m10* for 2018m4 and the realized EPS in 2018m4).

To illustrate the forecasting performance of this two-quarter-ahead EPS forecast error, we reestimate the VAR system using data up until December 2013. We calculate the n-step ahead (up to 36 months) dynamic forecasts of the movement in output starting January 2014. Figure 8 shows the result. The solid line is the observed industrial production, and the dashed line denotes the forecasted series. The gray area corresponds to +/- one standard deviation (68% confidence level) of the forecast error. We find that the forecasted production lies within the same confidence level with the observed one during a majority of the periods.

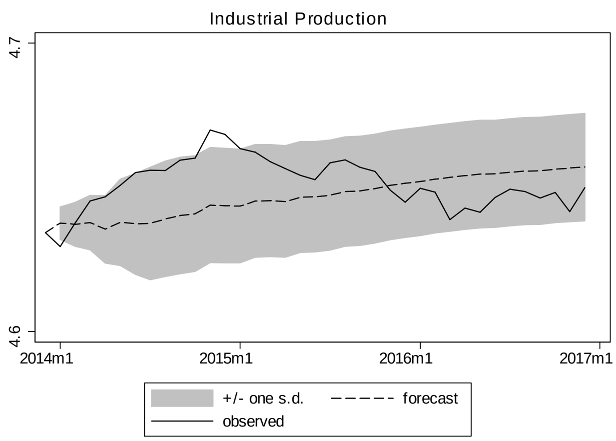

Figure 8 – 36-step ahead dynamic forecasts of industrial production based on the VAR system with two-quarter-ahead forecast error as the uncertainty measure

6 Industry Specific Uncertainty

Since our uncertainty measure is constructed using firm level observations, we can construct *industry specific* measures too, if the observations in the industry are numerous enough. Since analysts tend to observe the forecasts of other agents who follow the same stocks, and since some analysts specialize in particular sectors, this means analysts will be better informed about uncertainty in certain industries. Thus, it is interesting to explore whether the uncertainty shock to a specific industry has similar implications for macroeconomic fluctuations as demonstrated by the benchmark uncertainty. In this section, we examine the industry specific uncertainty by constructing uncertainty measure using the same methodology as in the benchmark, except that the EPS ratio is industry specific.

Specifically, we generate the following types of uncertainty measures. First, we develop an uncertainty measure for the financial sector and for the non-financial sector, to see whether the financial sector might be an origin for uncertainty shocks. Second, we distinguish between industries with high or low asset tangibility. Asset tangibility – typically defined as the share of fixed assets – is viewed as an important determinant of the sensitivity to financial shocks because tangible assets are easier to pledge as collateral and because the returns to tangible assets may be more easily secured by the owner – see Claessens and Laeven (2003) and Braun and Larrain (2005). As a result, the dynamics of uncertainty regarding firms with different levels of asset tangibility might be informative as to whether financial channels are important for the impact or propagation of uncertainty: if financial channels are important, we would expect uncertainty in low-tangibility industries to have a stronger link with aggregates. Third, we look at uncertainty inside and outside certain "hub" industries, identified according to input-output tables. This is to see whether input-output linkages might be important for the spread of uncertainty, something implied if input price uncertainty is important as in the models of Oi (1966), Hartman (1972) and Abel (1983). Finally, we split the private economy into 13 non-overlapping sectors and ask whether uncertainty as measured in each sector has common factors or not.

6.1 Financial sector

We explore the uncertainty measure based on the financial sector in a detailed manner. We are particularly interested in the financial sector because interactions between the financial sector and macroeconomy have received a lot of attention in the recent literature on uncertainty and in the literature that aims to understand the origins of the Great Recession of 2008. We define the financial sector as SIC codes $6000 - 6411$.

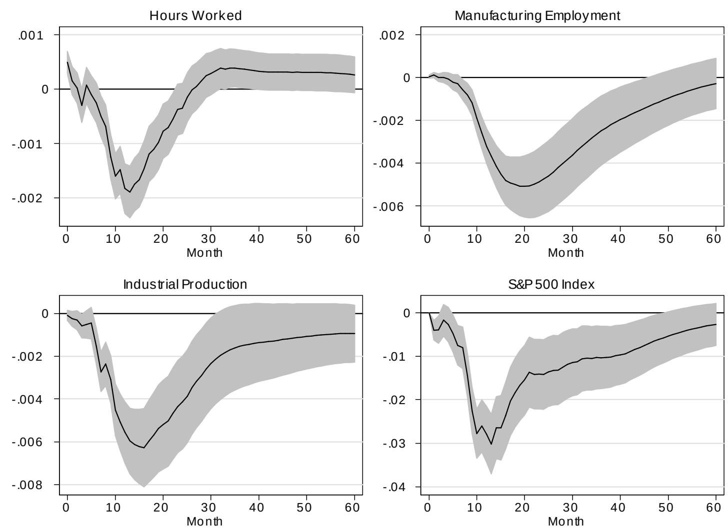

Figure 9 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with financial uncertainty

To quantitatively investigate the impact of financial uncertainty shocks, we estimate an alternative VAR system by replacing the benchmark uncertainty measure with financial uncertainty while keeping other variables the same. Figure 9 displays the impulse responses of hours worked, manufacturing employment, industrial production, and stock market to a one standard deviation innovation in the financial uncertainty measure. Like the benchmark uncertainty measure analyzed before, the macroeconomic variables decline following higher financial uncertainty. The magnitude of impact is also similar. However, financial uncertainty has more significant impact on the stock market, generating 50 percent more variation in the S&P 500 Index at the peak level. Another interesting result is that, immediately after impact, the responses of macroeconomic variables are significant for overall benchmark uncertainty, whereas *financial* uncertainty does not spread to the whole economy until it significantly influences the stock market.

We also calculate the forecast variance decomposition. Table 2 displays the VAR forecast error variance decomposition for industrial production and stock market over horizons of 3, 12, 36, and 60 months under the benchmark uncertainty measure and the financial uncertainty measure. Not surprisingly, since the overall uncertainty measure includes the financial sector and the non-financial sector, on average it conveys more information and therefore explains more variation in output, although financial uncertainty can explain as much variation in the medium run (36 months) as the overall measure. The stock market reaction, however, shows that financial uncertainty accounts for more variation in the change in stock prices than does the overall uncertainty measure, almost at all horizons. These results show that financial uncertainty has important implications for macroeconomic fluctuations, and is especially informative for studying the behavior of variables closely linked with the financial sector, such as the stock market.

	Industrial Production		S&P 500 Index	
Horizon	Baseline	Financial	Baseline	Financial
$h = 3$	3.33	0.22	0.37	0.73
$h = 12$	27.06	10.66	10.91	15.95
$h = 36$	26.27	21.70	14.28	25.49
$h = 60$	25.53	19.40	16.77	24.45

Table 2 – Industrial Production and Stock Market Forecast Variance Due to Baseline Uncertainty and Financial Uncertainty (in percent)

The conclusion is similar if we instead estimate VAR that includes both financial uncertainty and the overall uncertainty *excluding* the financial sector, which we term *non-financial* uncertainty. This VAR has the same variables as the benchmark VAR except that, (i) financial uncertainty replaces the benchmark overall uncertainty, and (ii) the non-financial uncertainty measure is added and ordered *after* financial uncertainty in the VAR system. This exercise gives us insight into whether uncertainty spreads from the financial sector to the rest of the economy.

Figure 10 shows the responses of industrial production and stock market to financial and non-financial uncertainty, respectively. In spite of the presence of the other uncertainty measure, both financial and non-financial uncertainty still have non-trivial effects on industrial production, which suggests that they contain independent information that has important implications for economic activity.

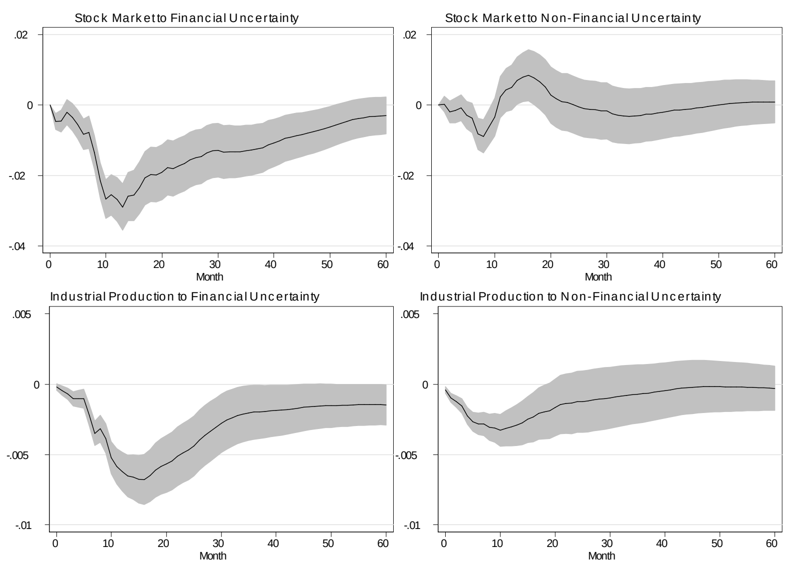

Figure 10 – Impulse response of stock market and industrial production from estimation of VAR with financial uncertainty and non-financial uncertainty. The VAR includes both forms of uncertainty, with financial uncertainty ordered first.

Furthermore, as financial uncertainty is placed before non-financial uncertainty, the effects of non-financial uncertainty on the stock market are measured *after* we have accounted for all the variation in non-financial uncertainty that is attributable to shocks to financial uncertainty. That the impact of non-financial uncertainty shocks on the stock market becomes almost insignificant after about 12 months reinforces the conclusion that uncertainty shocks to the financial sector play a more important role associated with stock market variation than uncertainty shocks to the non-financial sector. We obtain a similar conclusion if we change the ordering of financial and non-financial uncertainty, as illustrated in Appendix B.[24] This finding is similar to Carreiro et al (2018), Ludvigson et al (2018) and Shin and Zhong (2018), who also find that financial uncertainty has a more significant impact on financial markets and/or aggregates than overall uncertainty.

The results can also be shown by a forecast variance error decomposition exercise, similar with that displayed in Table 2, but under the financial uncertainty measure and the non-financial uncertainty measure. Table 3 shows that financial uncertainty explains more variation in output in almost all horizons compared to non-financial uncertainty. The stronger impact due to financial uncertainty is more so on the stock market, whereas non-financial uncertainty accounts for a relatively negligible variation.

Finally, when we compute the cross-correlogram of financial and non-financial uncertainty, we find that financial uncertainty leads non-financial uncertainty. This is seen in that the highest correlation is when financial uncertainty leads financial uncertainty by one month. In general the correlations of leads of financial uncertainty are higher than the similar lags. See Table 4 and Figure 11.

[24] In an alternative VAR system, we order non-financial uncertainty second, and financial uncertainty third. This implies that the effects of financial uncertainty on stock market are measured after we have removed all the variation in financial uncertainty that is attributable to shocks to non-financial uncertainty. The estimation results that the stock market still significantly responds to financial uncertainty shocks but not to non-financial uncertainty shocks suggest that financial uncertainty nontrivially contributes to the volatility in stock market.

	Industrial Production		S&P 500 Index	
Horizon	Financial	Non-Financial	Financial	Non-Financial
$h = 3$	1.04	3.93	1.03	0.10
$h = 12$	17.04	11.60	15.47	1.61
$h = 36$	26.96	5.72	27.43	1.70
$h = 60$	24.47	4.81	26.15	1.51

Table 3 – Industrial Production and Stock Market Forecast Variance Due to Financial Uncertainty and Non-Financial Uncertainty (in percent)

Lags of non-financial uncertainty	-5	-4	-3	-2	-1	0	1	2	3	4	5
Correlation	.15	.20	.22	.27	.32	.41	.44	.32	.33	.31	.26

Table 4 – Cross correlations of lags of financial and non-financial uncertainty

This finding is noteworthy, as the literature generally views uncertainty shocks and financial shocks as being different things, e.g. Caldara et al (2016), or views financial frictions as being a propagation mechanism for uncertainty shocks, as in Alfaro et al (2016). By measuring uncertainty independently inside and outside the financial sector, we find evidence that the former contains information that is not captured by uncertainty among non-financial firms. In particular, the time series of financial uncertainty shows that it precedes real sector uncertainty visually. Besides, a Granger causality test based on the VAR system with both measures reveals that the null hypothesis that lags of financial uncertainty do not affect non-financial uncertainty is rejected, whereas we cannot reject that lags of non-financial uncertainty do not affect financial uncertainty. This implies that, when we look at 12 months lags, innovations in financial uncertainty statistically significantly Granger-cause variation in non-financial uncertainty, but not vice versa. This evidence suggests that uncertainty itself may arise in the financial sector, before spreading to the real sector.

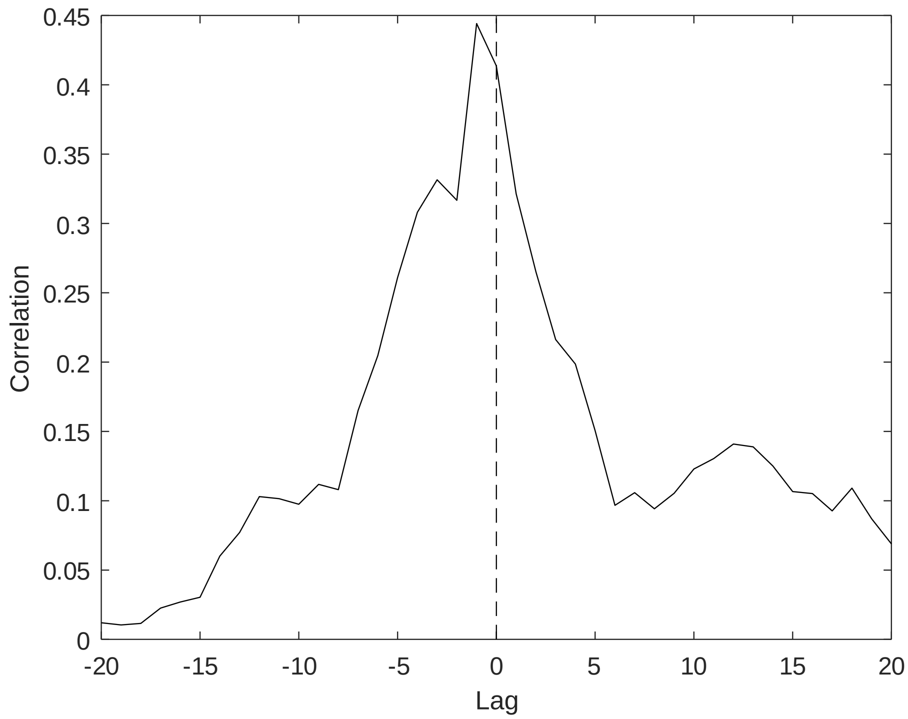

Figure 11 – Cross correlation between financial and non-financial uncertainty. The peak is to the left (one lag of financial uncertainty relative to non-financial uncertainty) and the area under the curve on the left is clearly greater than on the right.

6.2 Asset Tangibility

Another way of understanding whether uncertainty interacts with financial considerations is to distinguish firms based on whether they would likely be sensitive to financial conditions. Braun and Larrain (2005) measure asset tangibility in Compustat using the share of fixed assets out of total assets. We draw on the data of Samaniego and Sun (2015), who compute asset tangibility over the period 1970-2000.

We compute the tangibility measure for 41 industries.[25] Then, we compute two measures of uncertainty: a high tangibility measure (based on firms in industries above or equal to the median tangibility value) and a low tangibility measure (based on firms in industries below the median value of tangibility).

Figures 12 and 13 display the behavior of these two measures in the VAR separately. The dynamics are generated by estimating the same VAR system as for the benchmark uncertainty measure, except that the uncertainty measure is now high tangibility uncertainty or low tangibility uncertainty. We find that both measures have a smaller and less persistent impact on aggregates than the benchmark measure – particularly the high-tangibility measure. What is more striking is when we include both of them in one VAR system, where we order low-tangibility before high-tangibility when doing the estimation – see Figure 14. We find that the impulse responses of macroeconomic variables to the high-tangibilitiy uncertainty shock show not only a weaker response but indeed a response of the opposite sign after about 12 months. In contrast, the low-tangibility measure has the same impact as in Figure 13, indeed its impact on the stock market becomes more persistent. Switching the order does not affect these findings significantly. We conclude that uncertainty affecting firms that have difficulty raising external funds is important, consistent with there being important financial propagation channels or a financial origin for uncertainty shocks.

[25] We use the industry classification in Samaniego (2010).

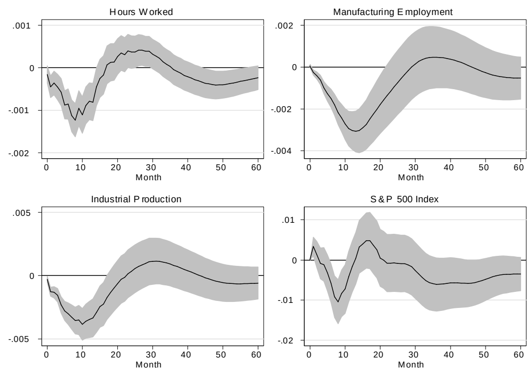

Figure 12 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with high-tangibility uncertainty.

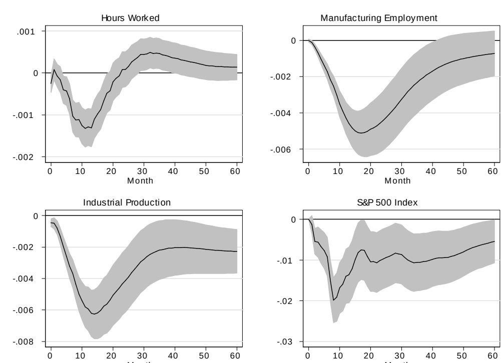

Figure 13 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with low-tangibility uncertainty.

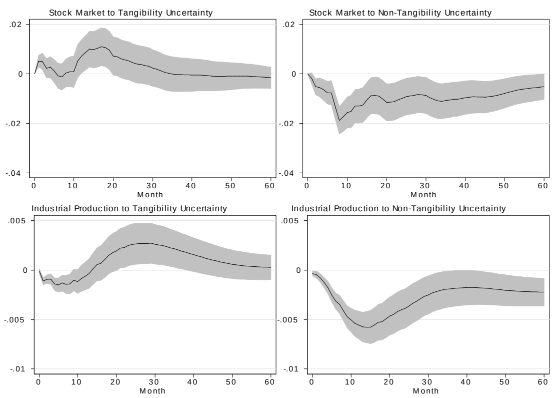

Figure 14 – Impulse response of stock market and industrial production from estimation of VAR with high-tangibility uncertainty and low-tangibility uncertainty. The VAR includes both forms of uncertainty, with low-tangibility uncertainty ordered before high-tangibility uncertainty.

6.3 "Hub" sectors

We also measure uncertainty for the wholesale+trucking sector, iron sector, and management sector by constructing the absolute forecast error of the EPS ratio for the corresponding industries only. We choose these three sectors because they are "hub" sectors according to input-output tables (see Jones 2013), and thus could potentially play an important role in either generating or transmitting uncertainty to other sectors.

Figure 15 plots the impulse responses of industrial production, manufacturing employment, hours worked, and stock market index to innovations in wholesale+trucking uncertainty measure. The results are obtained from the estimation of an alternative VAR system by replacing the benchmark uncertainty measure with wholesale and trucking uncertainty, while keeping other variables the same. A positive uncertainty shock to the wholesale and trucking sector leads to significant decrease in industrial production immediately. The maximum decline in production is 0.4 percent, smaller than the decline in production due to an increase in benchmark uncertainty. This is not surprising as the wholesale and trucking uncertainty may not capture all the uncertainty in the economy. There are also declines in hours worked and in manufacturing employment. Again, the declines are not as large as those generated by the baseline uncertainty shock. The contrast is more dramatic if we compare the responses in the stock market: the S&P 500 index does not significantly respond to the wholesale and trucking uncertainty shock at almost all horizons, whereas the baseline uncertainty shock can lead to 2 percent decline at its peak impact. Figure 16 shows the impulse responses of macroeconomic variables and stock market to iron uncertainty. The responses of labor market and production are non-trivial, but smaller and less persistent than those under the benchmark uncertainty. The results in Figure 17 suggest that management uncertainty almost does not lead to any significant changes in output at all horizons. See Appendix D for further robustness using hub uncertainty.

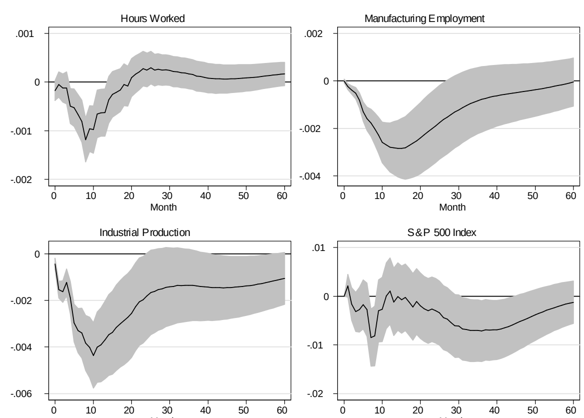

Figure 15 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with wholesale+trucking uncertainty.

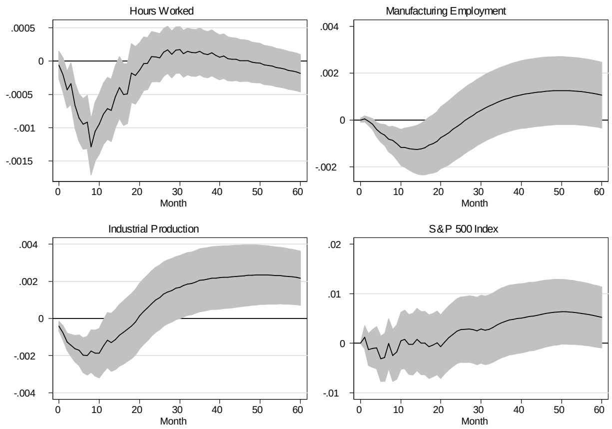

Figure 16 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with iron uncertainty.

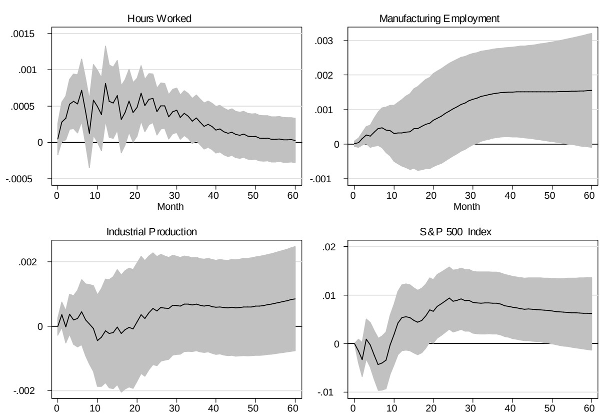

Figure 17 – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with management uncertainty.

We also introduce each of these hubs into a VAR along with an uncertainty measure for uncertainty at firms outside the hub, as we did for financial uncertainty earlier. We find that in general the hub uncertainty measures do not have much impact on the stock market or on output, whereas "non-hub" uncertainty does. We conclude that these hub industries are not important for generating nor channeling uncertainty, at least not through the stock market.

6.4 Decomposing the economy into 13 sectors

So far we have investigated the impact of uncertainty shocks to specific sectors of interest, such as financial sector, on the aggregate economy. Our approach allows to examine other industries as well. In this session, we construct 13 industry specific uncertainty measures that together cover the entire economy. This exercise allows us to study the properties of sector specific uncertainty for all the industries in the economy.

First, we must define our sectors. We began with the 1 digit SIC classification. This results in 9 sectors, most containing 1000 observations or more. We merged Wholesale Trade and Retail Trade, however, because separately they had few observations compared to other sectors. Then, we split certain sectors because they had significantly more than 1000 observations and because the sectors included products or services of arguably different nature. Specifically, we split codes $4000 - 4999$ into Transport, Communication and Utilities; codes $6000 - 6799$ into Finance/Insurance and Real Estate/Holdings; and codes $7000 - 8999$ into Personal Services, Business Services and Computer services. The resulting 13 sectors are Agriculture, Mining, Construction, Manufacturing, Transport, Communications, Utilities, Trade, Finance, Real Estate and Holdings, Personal services, Business services and Computer services.

Then, we estimate the VAR system with the same variables as the benchmark VAR model, except that the uncertainty measure is industry-specific. We perform the estimation for each of the 13 sectors, the impulse responses of which are shown in Appendix C. The impulse response results indicate that Computer Services uncertainty does not affect macroeconomic activities, whereas all the other industry-specific measures have significant impact on the output, at least at a short horizon. Among them, Agriculture, Transport, Real Estate/Holdings, and Business Services have less significant and less persistent impact than the other industry uncertainty measures. Interpreting the sector measures as reflecting sector-specific uncertainty shocks, these results suggest that uncertainty shocks to different sectors in the economy generate different dynamics, at least in some cases, and that these are mostly weaker than the impact of aggregate uncertainty.

We also perform principal-component factor analysis with all of them.[26] The objective of this exercise is to see whether there are important common factors of uncertainty contributing to all (or most) of the sectorial measures.

The results suggest that most of the measures are strongly correlated amongst themselves. Pairwise correlations are significant at the 5 percent level for each sector in 9 to 11 pairs in all cases. There are three exceptions however: Agriculture (3), Construction (4) and Computer services (6). This suggests that there may be common factors of uncertainty across sectors, but that some either are not sensitive to aggregate uncertainty or are exposed to strong sector-specific uncertainty factors.

Factor	Eigenvalue	Difference	Proportion	Cumulative
1	4.17	2.39	0.320	0.320
2	1.78	0.60	0.137	0.457
3	1.18	0.15	0.091	0.549
4	1.04	0.09	0.080	0.628

Table 5 – First four factors extracted using principal-component analysis. The other 9 have eigenvalues below one.

Table 5 reports the first four factors, which we retain based on the Kaiser (1960) criterion of having eigenvalues larger than one. However, observe that the eigenvalues

[26] There are other approaches to factor analysis: we use principal-component factorization because it requires no a-priori assumptions regarding the structure of the data.

drop rapidly, so that the Cattell (1966) screen test would select 2 or 3 factors at most. The first factor accounts for about a third of the variation in the sectorial data, and together the first four factors account for a bit under two thirds.

Table 6 displays the correlation between each of the four factors and sectorial uncertainty. Notice that the first factor is positively correlated with all of them, in some cases quite strongly. This suggests that Factor 1 is interpretable as a common factor of uncertainty, or *aggregate uncertainty*. This is not true of any of the other three factors. Factor 2 appears mainly related to the financial sector and the Real Estate and Holdings sector, suggesting that it is interpretable as a factor of financial uncertainty. Factor 3 is mainly related to the Agricultural sector (possibly indicating uncertainty in international commodities markets), and Factor 4 is mainly related to Utilities (possibly indicating uncertainty in international energy markets).

Sector	Factor 1	Factor 2	Factor 3	Factor 4	Uniqueness
Ag	0.161	-0.128	0.819	-0.068	0.282
Min	0.699	-0.289	0.026	-0.021	0.427
Con	0.186	0.451	0.166	-0.456	0.526
Man	0.800	-0.308	-0.050	-0.112	0.250
Tns	0.470	-0.447	0.392	0.169	0.397
Com	0.670	0.368	0.356	0.049	0.287
Uti	0.326	-0.066	-0.038	0.632	0.489
Trd	0.697	0.052	-0.215	-0.391	0.313
Fin	0.635	0.612	0.004	-0.079	0.217
RH	0.587	0.458	-0.082	0.320	0.336
Per	0.606	0.219	-0.276	0.184	0.474
Bus	0.614	-0.318	-0.120	0.065	0.503
Cmp	0.472	-0.554	-0.241	-0.286	0.331

Table 6 – Correlations between the factors and uncertainty by sector.

We also examine whether the main factor is similar to the aggregate uncertainty measure. We find that the correlation between aggregate uncertainty and the first factor is fully 0.93. Thus, all told, the first factor is interpretable as a common aggregate uncertainty factor. We remind the reader that, given how the measures were constructed independently, there is no particular reason why the sectorial measures should aggregate to or average out to aggregate uncertainty. In contrast, the correlation between aggregate uncertainty and the next three factors is 0.05, -0.07 and -0.03 respectively, again underlining the interpretation of them having mainly to do with one or two sectors.

7 Concluding Remarks

We construct new measures of economic uncertainty based on forecast errors of a large number of financial analysts. The measure we develop turns out to behave like other measures of uncertainty, even though its measurement strategy is very different. Variants of the baseline measure can be used for macroeconomic forecasting, or to look at uncertainty within particular industries or sectors. When we do this, we discover a key role for financial uncertainty. We do not find that input-output links are particularly important for propagation, and find evidence that there is indeed a common factor of uncertainty that influences most of the private economy.

There is scope for future work in several directions. The ability to measure uncertainty within sectors could be used more broadly to explore whether uncertainty shocks tend to start in particular sectors before spreading to others, or to study other channels whereby uncertainty might spread, for example through ownership structures, patent citation networks, etc.

Finally, we find that an increase in our uncertainty measure behaves like a negative demand-side shock that is not observed by forecasters at the moment when they make their forecasts, but that negatively affects current and future output. In future work it would be interesting to develop a model where uncertainty spreads in the ways we uncover. Embedded in a quantitative DSGE framework, this could be useful for understanding the spread of uncertainty to develop implications for optimal fiscal and monetary policy.

8 References

Abel, Andrew B. 1983. Optimal Investment Under Uncertainty. American Economic Review 73(1), 228 –33.

Alfaro, Ivan, Nicholas Bloom and Xiaoji Lin. 2016. The Finance-Uncertainty Multiplier. Mimeo: Stanford University.

Angeletos, George-Marios, Fabrice Collard and Harris Dellas. 2016. Quantifying Confidence. Mimeo: Massachusetts Institute of Technology and University of Bern

Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims. 2013. Uncertainty and Economic Activity: Evidence from Business Survey Data. American Economic Journal: Macroeconomics, 5(2): 217-49.

Bachmann, Rüdiger, Benjamin Born, Steffen Elstner, and Christian Grimme. 2017. Time-Varying Business Volatility and the Price-Setting of Firms. NBER working paper.

Baker, Scott R, Nicholas Bloom and Steven J. Davis. 2016. Measuring Economic Policy Uncertainty. Quarterly Journal of Economics 131 (4), 1593-1636.

Basu, Susanto, and Brent Bundick. 2017. Uncertainty Shocks in a Model of Effective Demand. Econometrica, 85(3): 937-958.

Beni, Gerardo and Jing Wang. 1989. Swarm Intelligence in Cellular Robotic Systems, Proceed. NATO Advanced Workshop on Robots and Biological Systems, Tuscany, Italy.

Bloom, Nicholas. 2009. The Impact of Uncertainty Shocks. Econometrica 77, 623-685.

Braun, Matias and Borja Larrain. 2005. Finance and the Business Cycle: International, Inter-Industry Evidence. Journal of Finance 60(3), 1097-1128.

Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek. 2016. The Macroeconomic Impact of Financial and Uncertainty Shocks. European Economic Review 88, 185-207.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino. 2018. Measuring Uncertainty and Its Impact on the Economy. Review of Economics and Statistics, Forthcoming.

Cattell, Raymond B. 1966. The Scree Test for the number of factors. Multivariate Behavioral Research, 1(2), 245-276.

Choi, Sangyup. 2013. Are the Effects of Bloom's Uncertainty Shocks Robust? Economics Letters, 119(2): 216-220.

Claessens, Stijn and Luc Laeven. 2003. Financial Development, Property Rights, and Growth. 58(6), 2401-2436.

Clement, Michael B. and Tse, Senyo Y. 2005. Financial Analyst Characteristics and Herding Behavior in Forecasting. Journal of Finance 60(1), 307-341.

Fudenberg, Drew and David Levine. 1998. The Theory of Learning in Games. MIT Press: Cambridge, MA.

Garín, Julio, Michael Pries and Eric Sims. 2018. The Relative Importance of Aggregate and Sectoral Shocks and the Changing Nature of Economic Fluctuations. American Economic Journal: Macroeconomics 10(1), 119–148.

Hartman, Richard. 1972. The Effects of Price and Cost Uncertainty on Investment. Journal of Economic Theory 5(2), 258–66.

Hillary, Gilles and Charles Hsu. 2013. Analyst Forecast Consistency. Journal of Finance 68(1), 271–297.

Horvath, Michael. 1998. Cyclicality and Sectoral Linkages: Aggregate Fluctuations from Independent Sectoral Shocks. Review of Economic Dynamics 1, 781-808.

Horvath, Michael. 2000. Sectoral Shocks and Aggregate Fluctuations. Journal of Monetary Economics 42, 69-106.

Jones, Charles I. 2013. Misallocation, Economic Growth, and Input-Output Economics. In D. Acemoglu, M. Arellano, and E. Dekel, Advances in Economics and Econometrics, Tenth World Congress, Volume II, Cambridge University Press, 419-455.

Jurado, Kyle, Sydney C. Ludvigson and Serena Ng. 2015. Measuring Uncertainty. American Economic Review 105(3), 1177-1216.

Kaiser, Henry F. (1960). The application of electronic computers to factor analysis. Educational and Psychological Measurement, 20, 141–151.

Kalton, Grahan and Howard Schuman. 1982. The Effect of the Question on Survey Responses: A Review. Journal of the Royal Statistical Society, Series A, 145(1), 42-73.

Leduc, Sylvain, and Zheng Liu. 2016. Uncertainty Shocks are Aggregate Demand Shocks. Journal of Monetary Economics, 82(C): 20-35.

Ludvigson, Sydney, Sai Ma and Serena Ng. 2018. Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response? Mimeo: NYU.

Milani, Fabio. 2017. Sentiment and the US Business Cycles. Journal of Economic Dynamics and Control 82, 289-311.

Oi, Walter Y. 1961. The Desirability of Price Instability under Perfect Competition. Econometrica 29(1), 58–64.

Ravn, Morten and Harald Uhlig. 2002. On Adjusting the Hodrick–Prescott Filter for the Frequency of Observations. The Review of Economics and Statistics. 84 (2): 37.

Scotti, Chiara. 2016. Surprise and Uncertainty Indexes: Real-Time Aggregation of Real-Activity Macro Surprises. Journal of Monetary Economics 82, 1-19.

Shin, Minchul and Molin Zhong. 2018. A New Approach to Identifying the Real Effects of Uncertainty Shocks. Forthcoming, Journal of Business and Economic Statistics.

So, Eric C. 2013. A New Approach to Predicting Analyst Forecast Errors: Do Investors Overweight Analyst Forecasts? Journal of Financial Economics 108(3), 615-640.

Trueman, Brett. 1994. Analyst Forecasts and Herding Behavior. Review of Financial Studies 7(1), 97-124.

Tversky, Amos and Daniel Kahneman. 1974. Judgement under Uncertainty: Heuristics and Biases. Science 185, 1124-1131.

Appendix: for online publication only

A Robustness: comparing different measures

We can also compare the quantitative importance of different measures by examining the forecast error variance decomposition for industrial production under alternative specifications for the measure of uncertainty discussed in the robustness and discussion sections. Similar to Table 1, we compute the VAR forecast error variance over horizons of 3, 12, 36, and 60 months. The 2- and 3-quarterly measures explain more variation than the benchmark measure within a short horizon, whereas the EPS/Price1 and EPS/Price2 measures account for a larger portion of variation than the benchmark in the medium run.

Horizon	Baseline	1 Qtr	2 Qtr	3 Qtr	NoFin	EPS/P1	EPS/P2
$h = 3$	3.33	0.63	5.70	1.72	3.86	2.04	2.20
$h = 12$	27.06	0.56	29.35	29.37	19.93	26.06	25.78
$h = 36$	26.27	0.98	15.30	25.92	17.02	26.31	30.90
$h = 60$	25.53	1.35	10.49	16.13	17.15	25.39	28.86

Table $A1$ – Industrial Production Forecast Variance Due to Uncertainty (in percent). 1 Qtr, 2 Qtr, and 3 Qtr are uncertainty measures based on 1-, 2-, and 3-quater ahead absolute forecast errors. NoFin is uncertainty without using financial firms.EPS/P1 and EPS/P2 are normalized using share prices. See text for details.

In this appendix we study some variations in the structure of the VAR including our baseline measure.

Change of order for uncertainty measure

In the benchmark system, we order the uncertainty in the second order in the VAR. In order to check whether the impulse response results are driven by the ordering of our uncertainty measure in the VAR system, we redo the estimation by placing the uncertainty last in the VAR. Given this structure, the effects of uncertainty shocks on other variables in the system are measured after we have removed all the variation in uncertainty that is attributable to shocks to the other endogenous variables. Figure $A1$ shows that the effects of uncertainty shocks on production, employment, hours worked, and stock market are still significant, which is consistent with the view that uncertainty has important implications for economic activity.

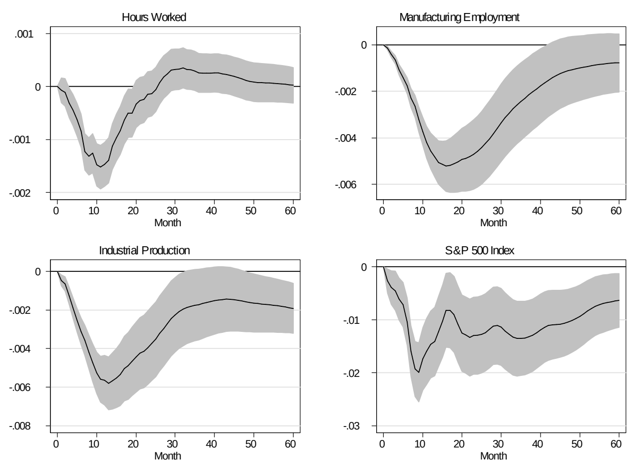

Figure $A1$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR *with uncertainty* placed last

Drop pre-1983 data (very volatile inflation).

Since our benchmark measure is defined as the forecast error of EPS deflated by CPI, a very volatile inflation could significantly affect the results. We test whether the impulse responses to our uncertainty measure are driven by the pre-1983 data, when the inflation rate was very volatile. Figure $A2$ shows the impulse responses to the same VAR system as the benchmark system, except that the data used for the estimation starts at 1983 rather than 1981. The results are qualitatively and quantitatively similar to the benchmark, i.e., a positive uncertainty shock reduces production and employment.

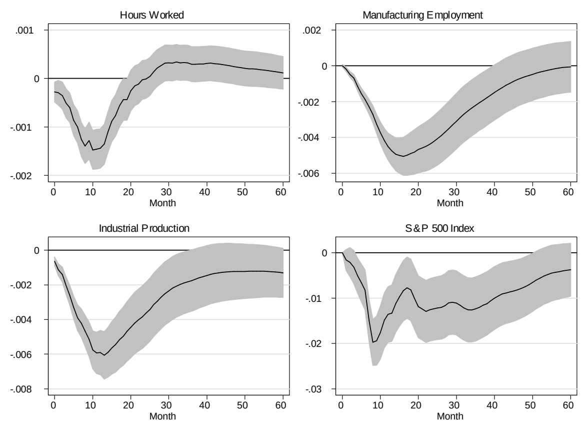

Figure $A2$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR excluding pre-1983 data

Excluding nominal variables in the VAR system

We estimate the system without nominal variables, CPI and federal funds rate. Qualitatively, the responses of production, labor market, and stock index are similar to those in the benchmark system. Quantitatively, however, the magnitude of the responses are smaller. This could be due to that monetary policy endogenously respond to a positive uncertainty shock, which is independently captured by the benchmark results. In other words, the impact of uncertainty shock is larger without the intervention of monetary policy. See Figure $A3$.

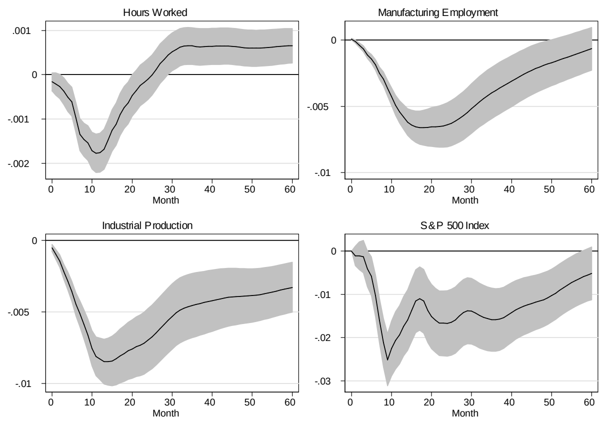

Figure $A3$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR without nominal variables

B Financial v.s. non-financial uncertainty

In the main text, we compare the performance of uncertainty measure constructed as the median forecast errors of EPS in financial and non-financial sectors. When doing that exercise, we order financial uncertainty before non-financial uncertainty. We find that both financial and non-financial uncertainty play a significant role in accounting for variations in macroeconomic variables. However, only financial uncertainty contributes significantly to the fluctuations in stock market. In this robustness check, we explore whether the conclusion is related to the ordering of the financial and non-financial uncertainty. We estimate an alternative VAR system by ordering non-financial uncertainty second and financial uncertainty third. This implies that the effects of financial uncertainty on stock market are measured after we have removed all the variation in financial uncertainty that is attributable to shocks to non-financial uncertainty.

The impulse responses of industrial production and stock market to financial and non-financial uncertainty are shown in Figure $C1$. As before, increases in financial uncertainty and non-financial uncertainty still generate significant impact on industrial production. The ordering is irrelevant when analyzing the contribution of these two uncertainty measures to macroeconomic dynamics. However, as in Figure 11, stock market only significantly responds to financial uncertainty shock but not to non-financial uncertainty shock. This suggests that financial uncertainty plays a more important role associated with the volatility in stock market.

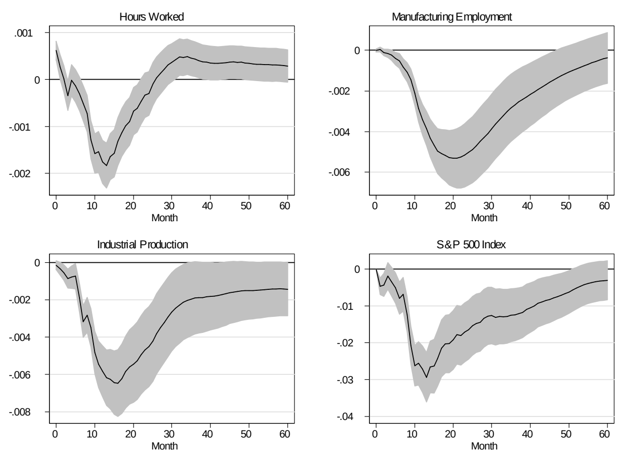

Figure $B1$ – Impulse response of stock market and industrial production from estimation of VAR with financial uncertainty and non-financial uncertainty. The VAR includes both forms of uncertainty, with non-financial uncertainty ordered before financial uncertainty.

C Industry specific uncertainty

In the remainder of the Appendix we report impulse responses from estimating the VAR with the uncertainty measures for each of the non-overlapping 13 sectors.

First, Figure 9 shows the financial uncertainty index over time. Compared with the benchmark uncertainty in Figure 1, the financial uncertainty is relatively more smooth in most of the time but is more volatile for certain historical periods, such as the Great Recession, which originated from the financial market.

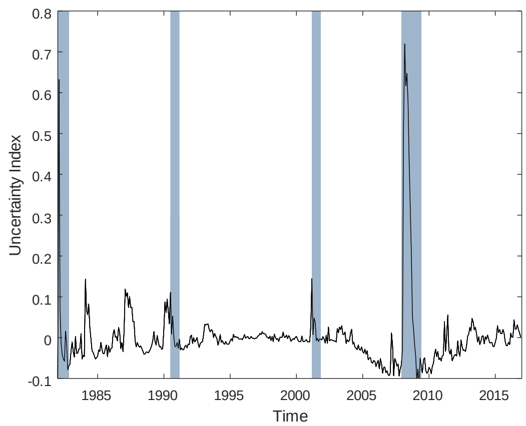

Figure $C1$ – Uncertainty measured only for financial firms. Financial firms are defined as firms with SIC codes 6000-6411 as reported in CRSP.

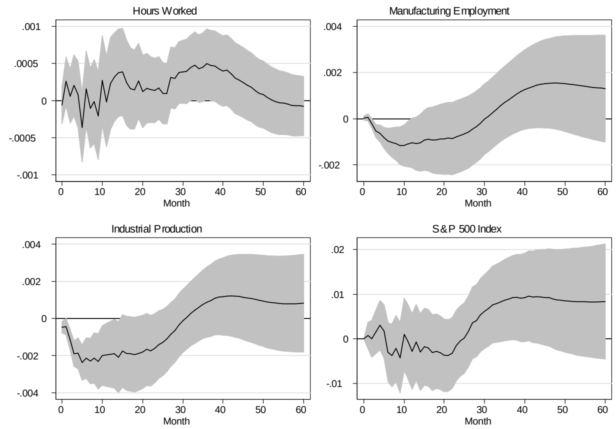

Figure $C2$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with AG uncertainty.

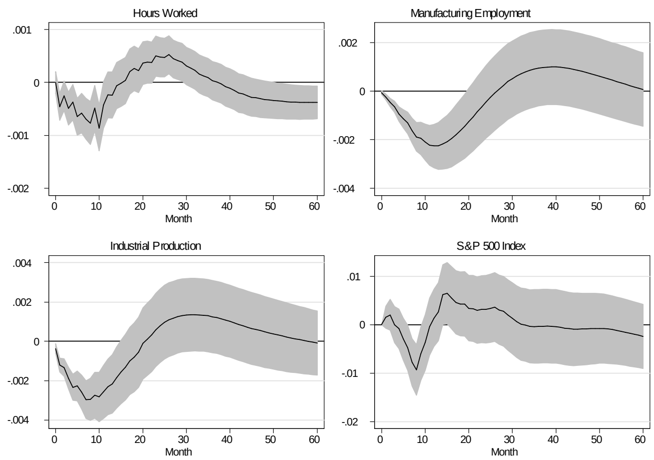

Figure $C3$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with MIN uncertainty.

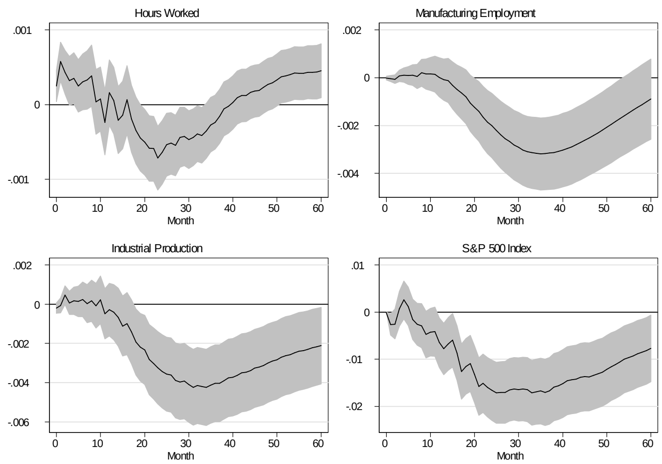

Figure $C4$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with CON uncertainty.

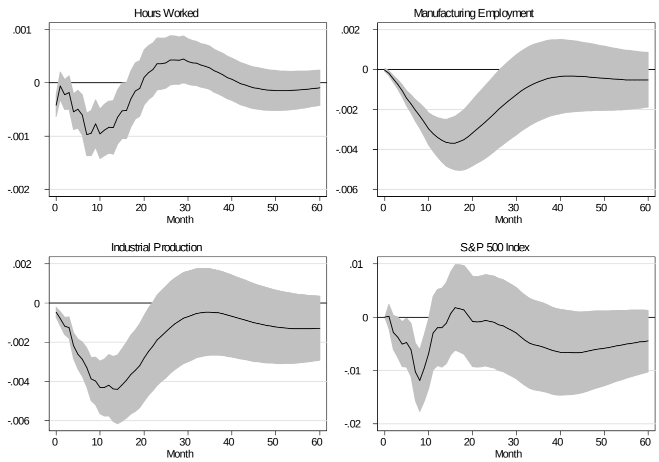

Figure $C5$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with MAN uncertainty.

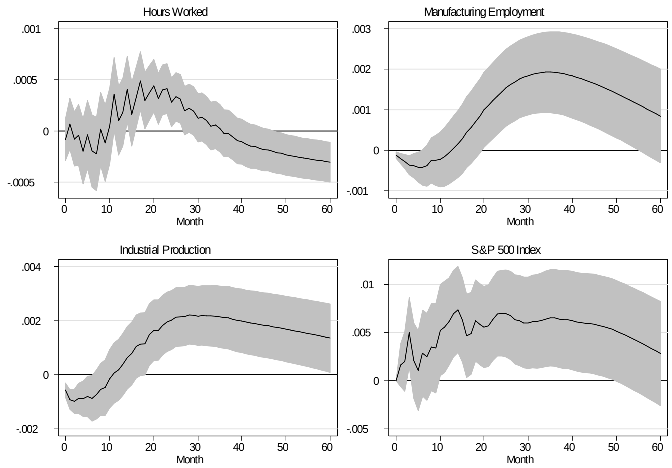

Figure $C6$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with TNS uncertainty.

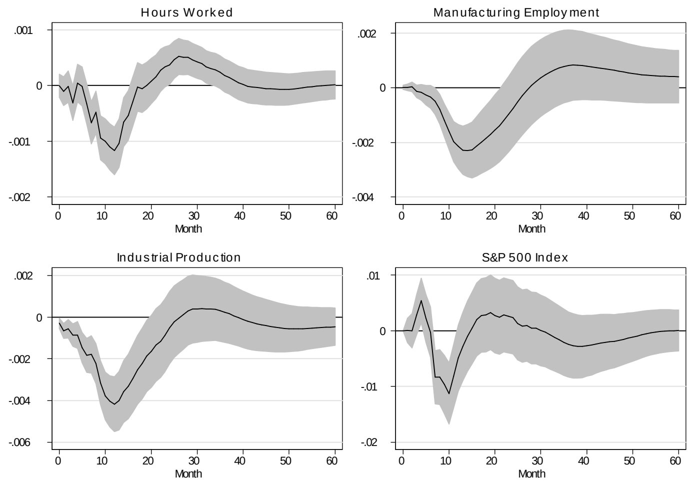

Figure $C7$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with COM uncertainty.

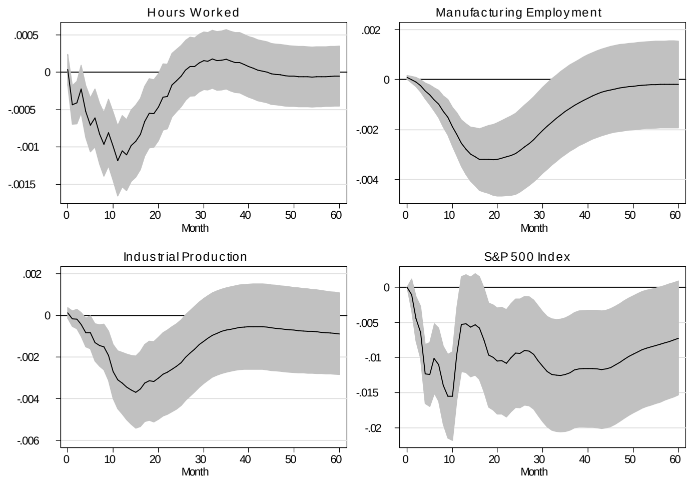

Figure $C8$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with UTI uncertainty.

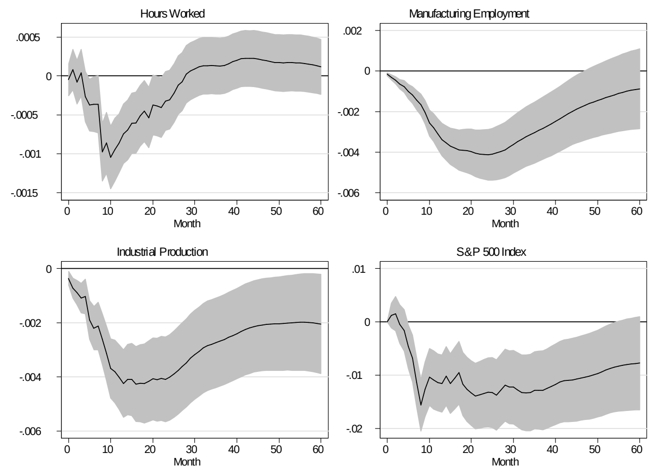

Figure $C9$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with TRD uncertainty.

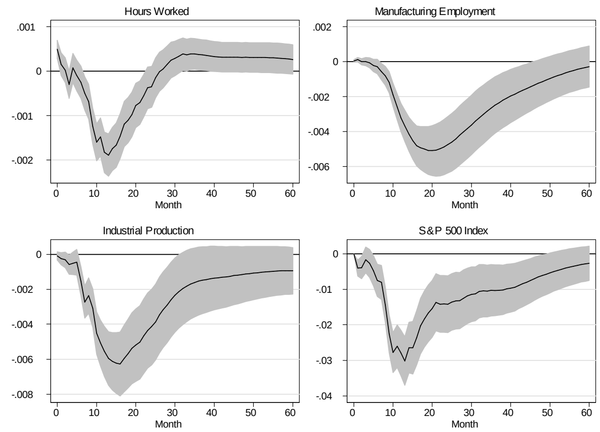

Figure $C10$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with FIN uncertainty.

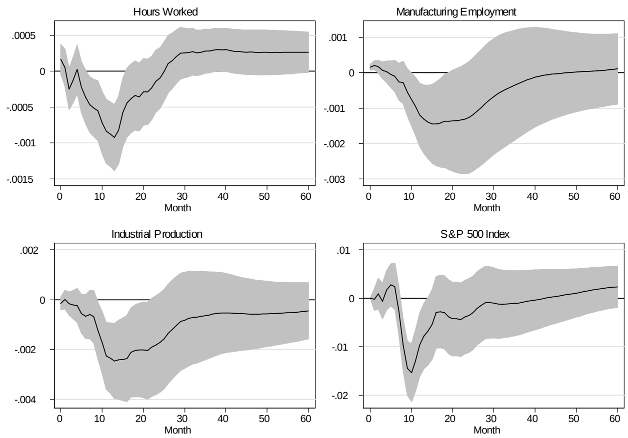

Figure $C11$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with RH uncertainty.

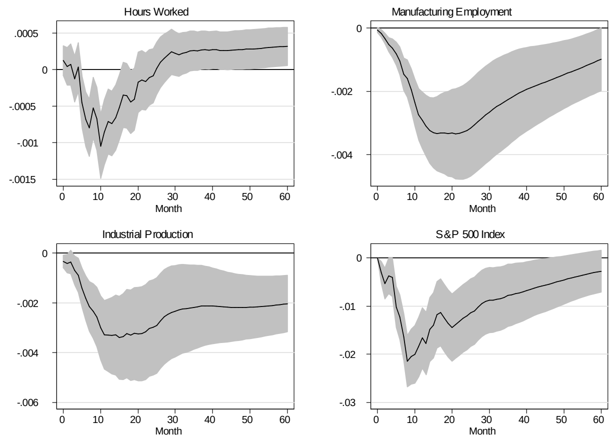

Figure $C12$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with PER uncertainty.

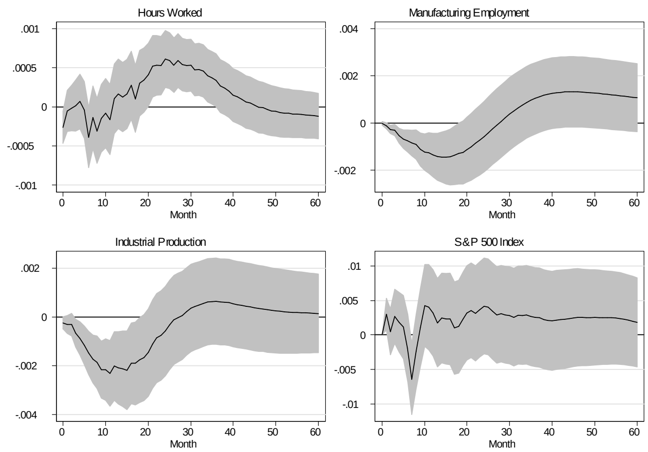

Figure $C13$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with BUS uncertainty.

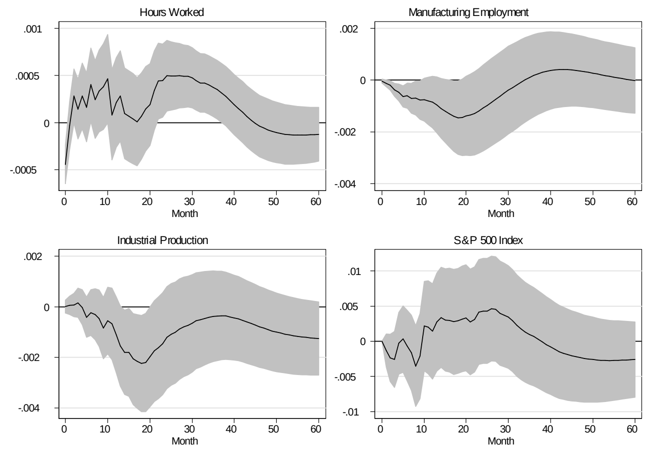

Figure $C14$ – Impulse response of hours worked, manufacturing employment, industrial production, and stock index from estimation of VAR with CMP uncertainty.

D Other results

As for the financial sector, we do similar exercises for other hub sectors. In particular, we include both the uncertainty measured using the forecast errors based on the data for each hub industry - wholesale, iron, management, and the overall uncertainty *excluding* these industries, which we term non-wholesale*, non-iron, and non-management* uncertainty. The VAR system for each hub sector has the same variables as the benchmark VAR except that, (i) sector specific uncertainty replaces the benchmark overall uncertainty, and (ii) the non-sector uncertainty measure is added and ordered *after* the sector uncertainty. Figure $D1, D2$,and $D3$ show the responses of industrial production and stock market to wholesale, iron, management uncertainty, and non-wholesale, non-iron, non-management uncertainty, respectively. First, none of these hub sector uncertainty significantly affects the stock market; and the uncertainty based on the data excluding these hub sectors does have significant impact on the stock market. This suggests that the uncertainty other than that associated with these hub sectors is the one that generate stock market fluctuations, and we have seen that the financial uncertainty (and IT uncertainty) are two candidates. Second, with the presence of the other uncertainty measure, wholesale uncertainty and iron uncertainty only have slightly significant effects on industrial production on impact, and last for less than 12 months. However, the magnitude and persistence are much weaker than that generated by non-wholesale and non-iron uncertainty shocks.

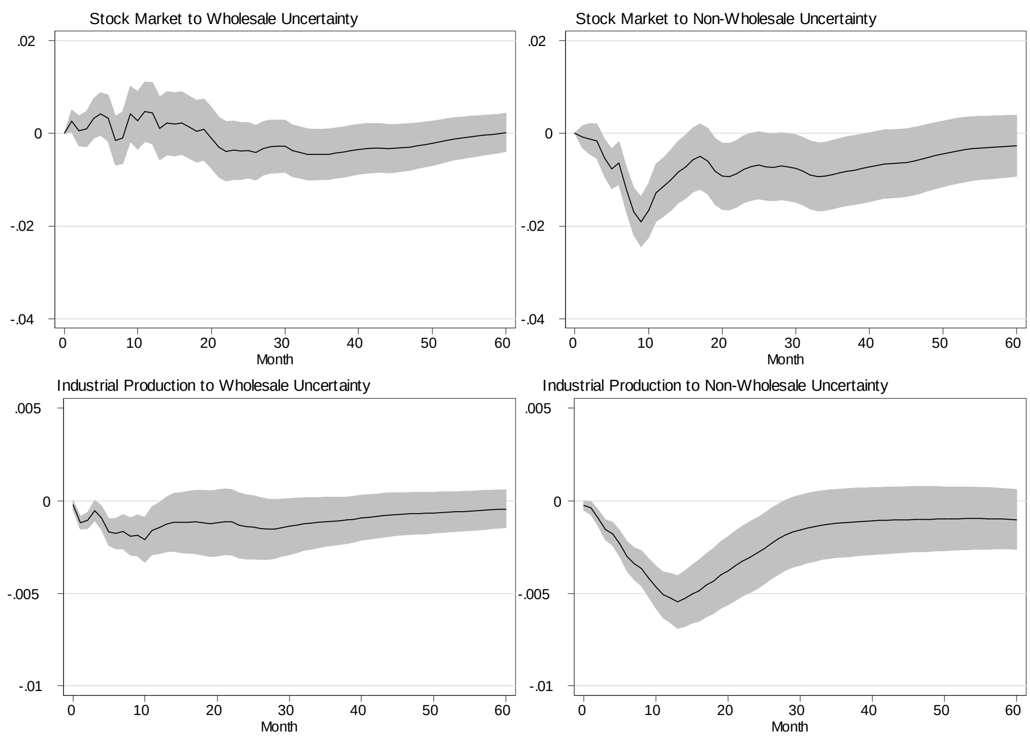

Figure $D1$ – Impulse response of stock market and industrial production from estimation of VAR with wholesale uncertainty and non-wholesale uncertainty. The VAR includes both forms of uncertainty, with wholesale uncertainty ordered first.

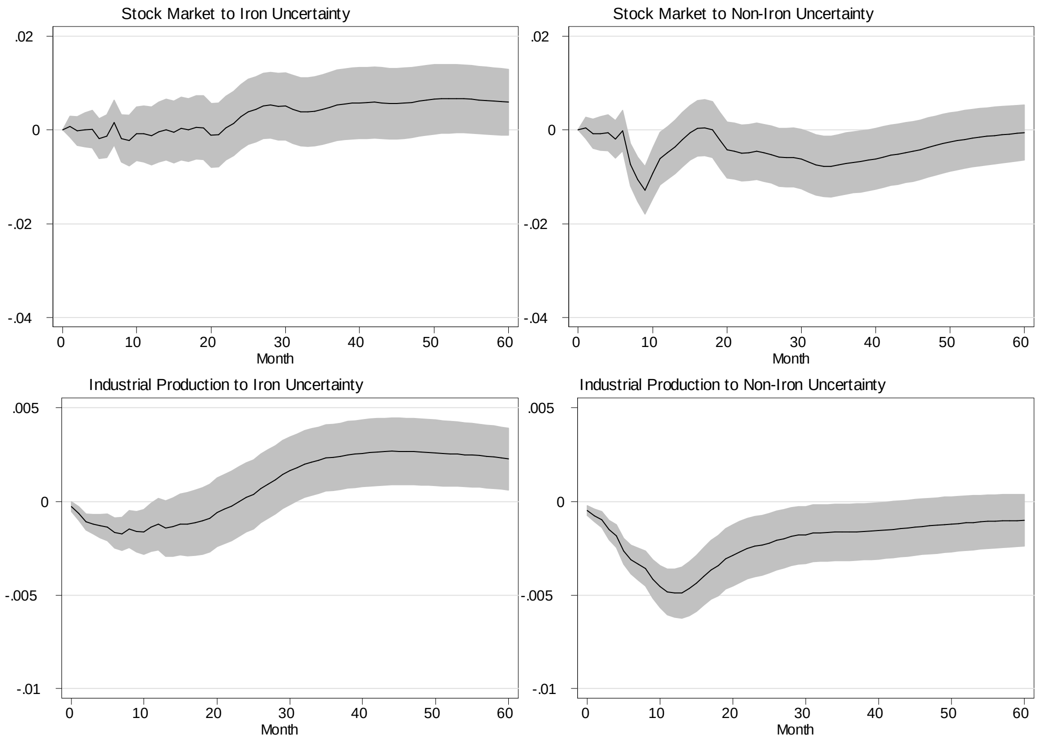

Figure $D2$ – Impulse response of stock market and industrial production from estimation of VAR with iron uncertainty and non-iron uncertainty. The VAR includes both forms of uncertainty, with iron uncertainty ordered first.

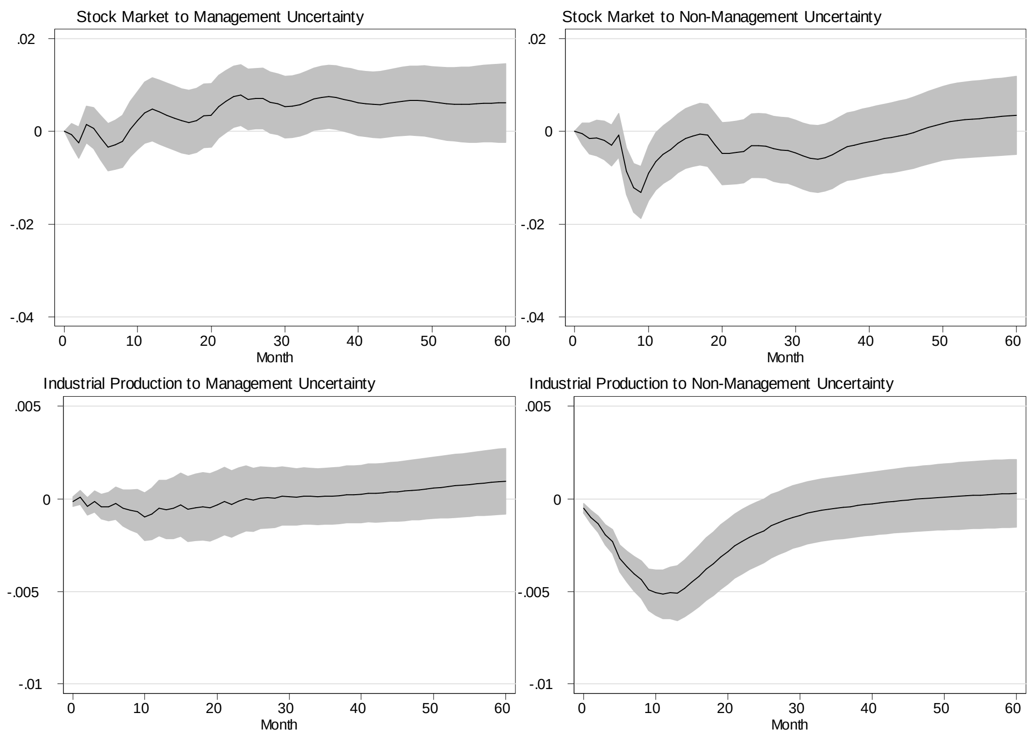

Figure $D3$ – Impulse response of stock market and industrial production from estimation of VAR with management uncertainty and non-management uncertainty. The VAR includes both forms of uncertainty, with management uncertainty ordered first.

To check whether the previous results depend on the order of the uncertainty measures in the VAR system, we redo the exercises by putting the non-sector specific uncertainty as the second variable in the VAR, and the corresponding sector specific uncertainty as the third, while keep all the other variables the same ordering. The results are shown in the following figures. The dynamics suggest that the ordering does not matter for our conclusion: wholesale and iron uncertainty still has non persistent and slightly significant impacts on the output, but no impact on the stocks market; management uncertainty has no impact on either the stock market or the output.

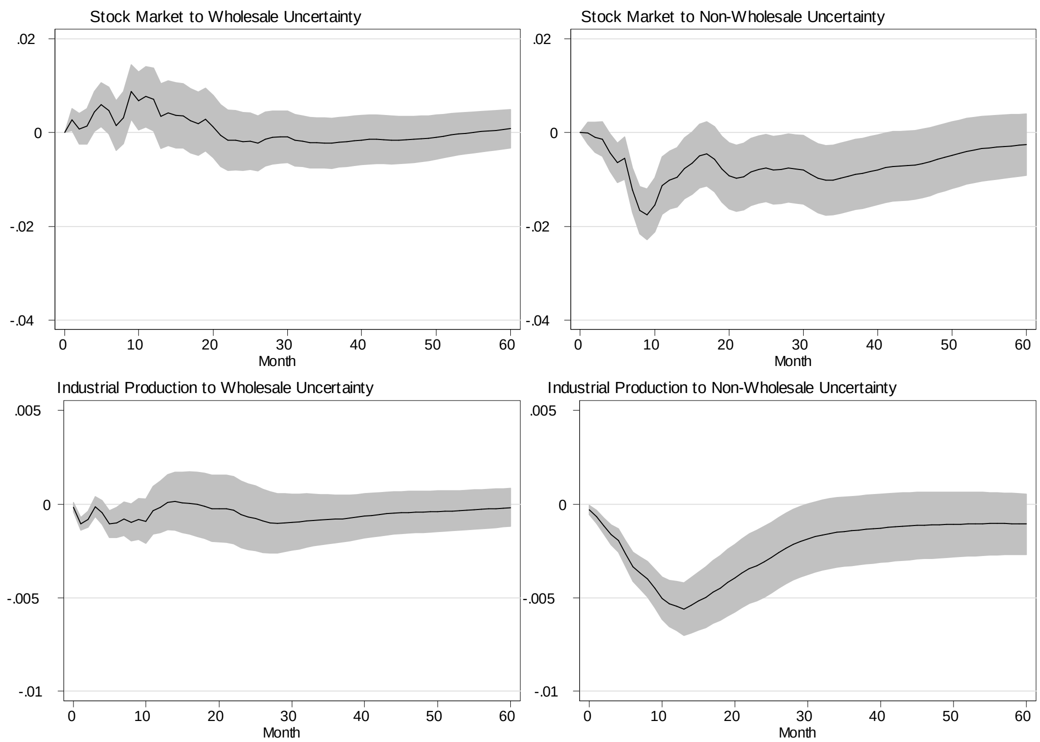

Figure $D4$ – Impulse response of stock market and industrial production from estimation of VAR with wholesale uncertainty and non-wholesale uncertainty. The VAR includes both forms of uncertainty, with non-wholesale uncertainty ordered before wholesale uncertainty.

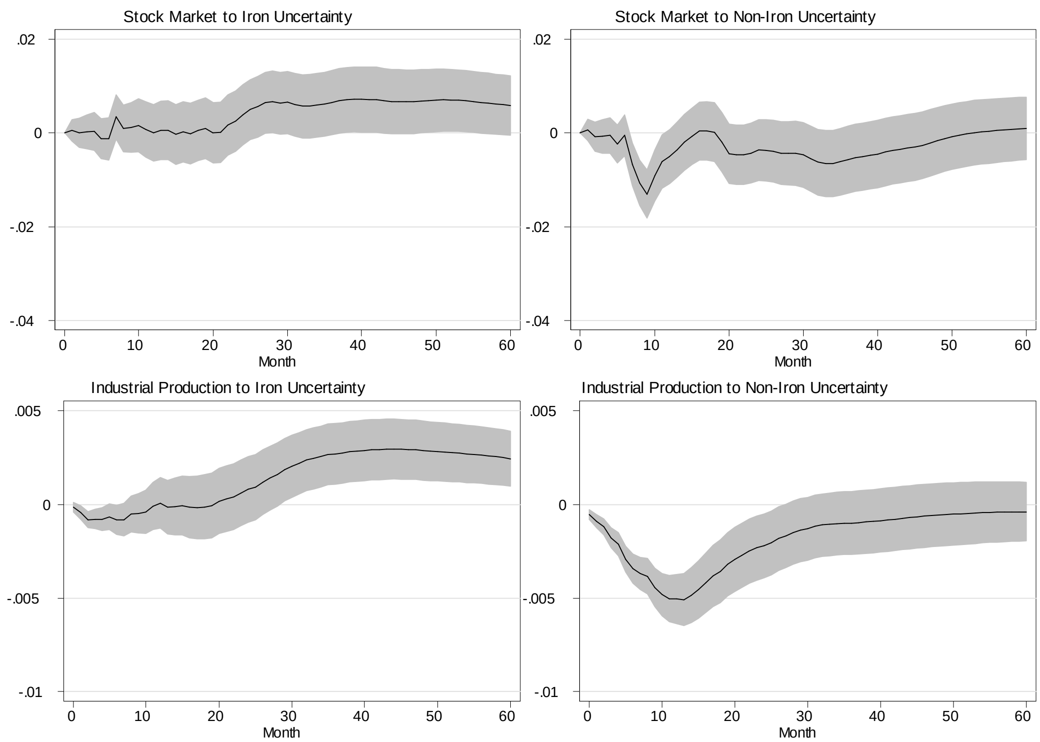

Figure $D5$ – Impulse response of stock market and industrial production from estimation of VAR with iron uncertainty and non-iron uncertainty. The VAR includes both forms of uncertainty, with non-iron uncertainty ordered before iron uncertainty.

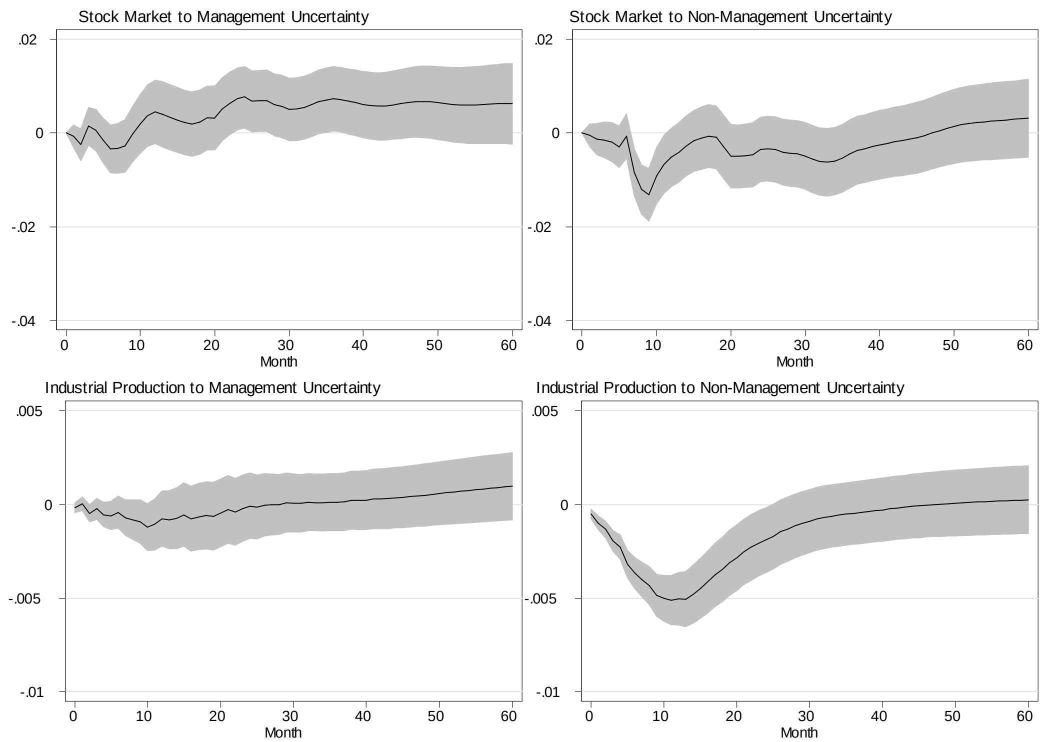

Figure $D6$ – Impulse response of stock market and industrial production from estimation of VAR with management uncertainty and non-management uncertainty. The VAR includes both forms of uncertainty, with non-management uncertainty ordered first.

E Alternative approaches to measuring uncertainty

As the baseline, we measure uncertainty using the absolute value of the median forecast errors for firms in the dataset for each month. However, we also explored two alternative approaches.

One approach was to use the median of the absolute value of the forecast errors – in other words, to compute absolute values first before computing the median.

When we did this, we found that the series displayed clear monthly effects, with "uncertainty" being systematically lower in October-January than in other months. This suggests that the series is capturing something other than uncertainty, e.g. some seasonal effects perhaps related to holiday sales. In contrast, the benchmark series do not display these effects. When removed these effects (by using dummies for each month of the year and removing the series predicted by those dummies) we found a series that resembled our baseline series in that it identifies spikes in uncertainty at similar events. The correlation between the two series is 0.27 and highly significant. See Figure $E1$.

Another approach was to use means instead of medians. The problem was that there were some outliers orders of magnitude larger than most of the others, leading to a series that was not useful. On the other hand, removing outliers conceptually is problematic because large errors are precisely what we use to identify periods of uncertainty. In any case, we computed a mean-based measure dropping the top and bottom 5 percent of mean forecast error observations. The measure again displayed some monthly effects. After removing them again we ended up with a series that resembles the baseline measure, this time with a correlation of 0.90.

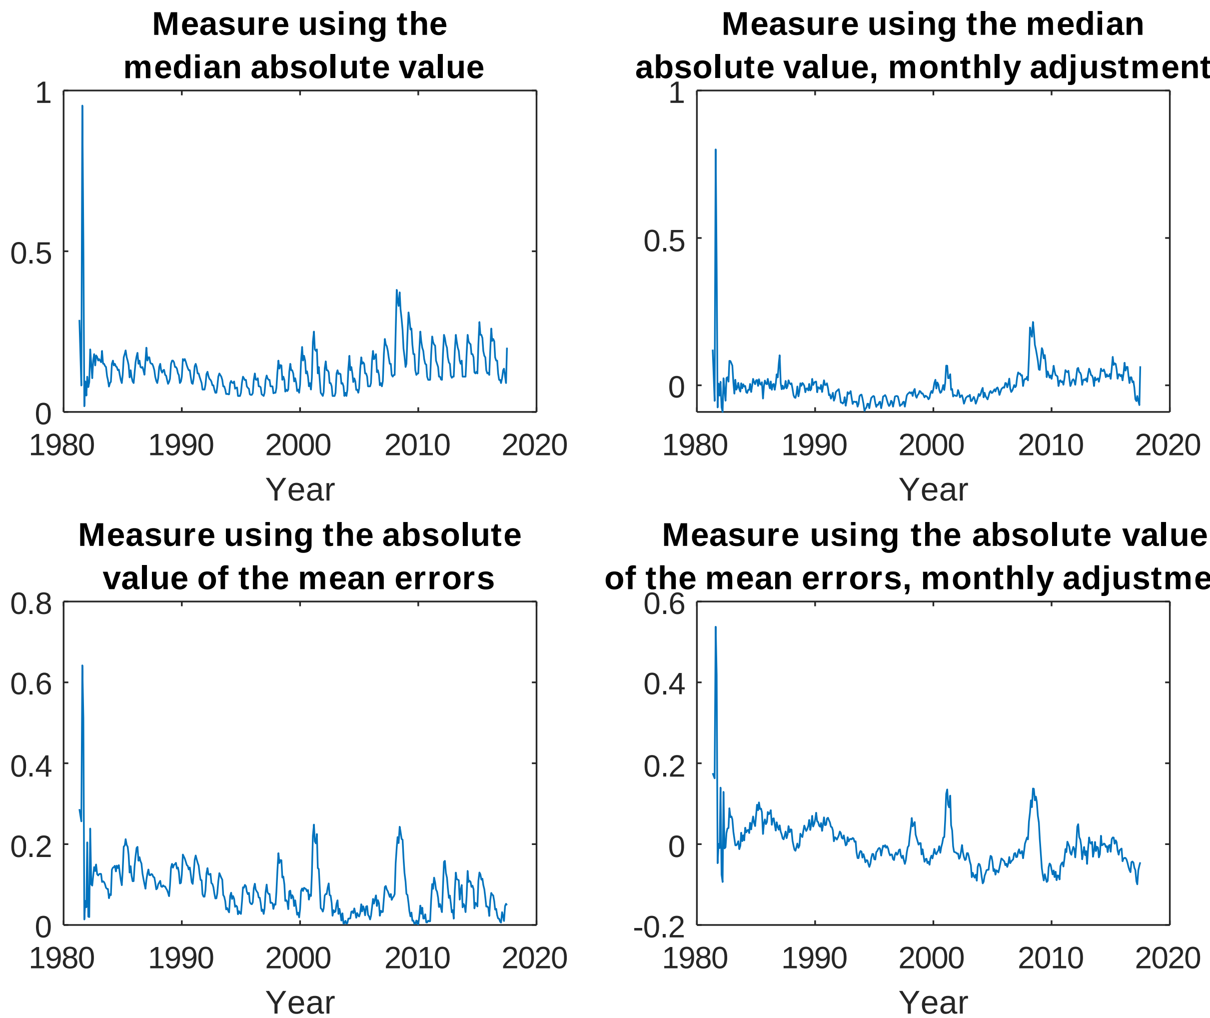

Figure $E1$ – Time series of measures computed using different procedures. See text for definitions

The impact of shocks to these alternative uncertainty measures is also similar with that based on the benchmark measure. Figure $E2$ and $E3$ display the impulse responses of macroeconomic variables to these shocks. The responses are calculated from the same VAR system as the baseline, expect that the uncertainty index is instead the median of absolute value of forecast errors and the absolute value of mean forecast errors, respectively. Increases in both uncertainty shocks lead to negative impact on real variables, stock market, and easing monetary policy, as that demonstrated by Figure 2 for the baseline uncertainty measure.

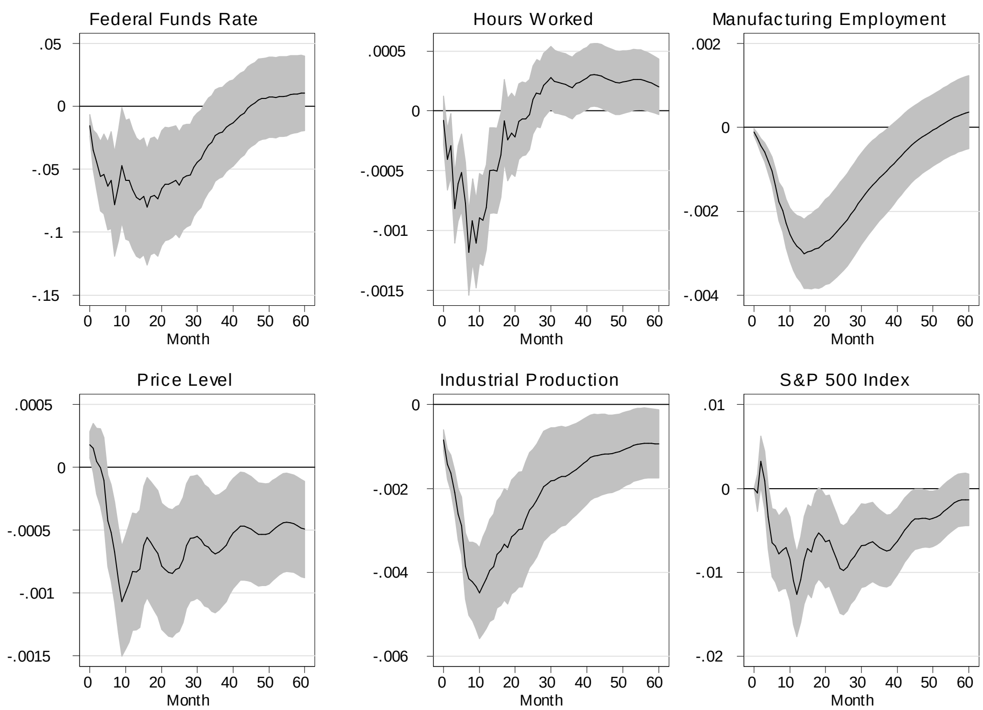

Figure $E2$ – Impulse response of federal funds rate, price level, hours worked, industrial production, manufacturing employment, and stock index from estimation of VAR, with uncertainty measured using the median of absolute forecast error. The shock is one standard deviation, and the error bands are +/- one standard error confidence bands. The sample period is 1981m9 to 2016m12.

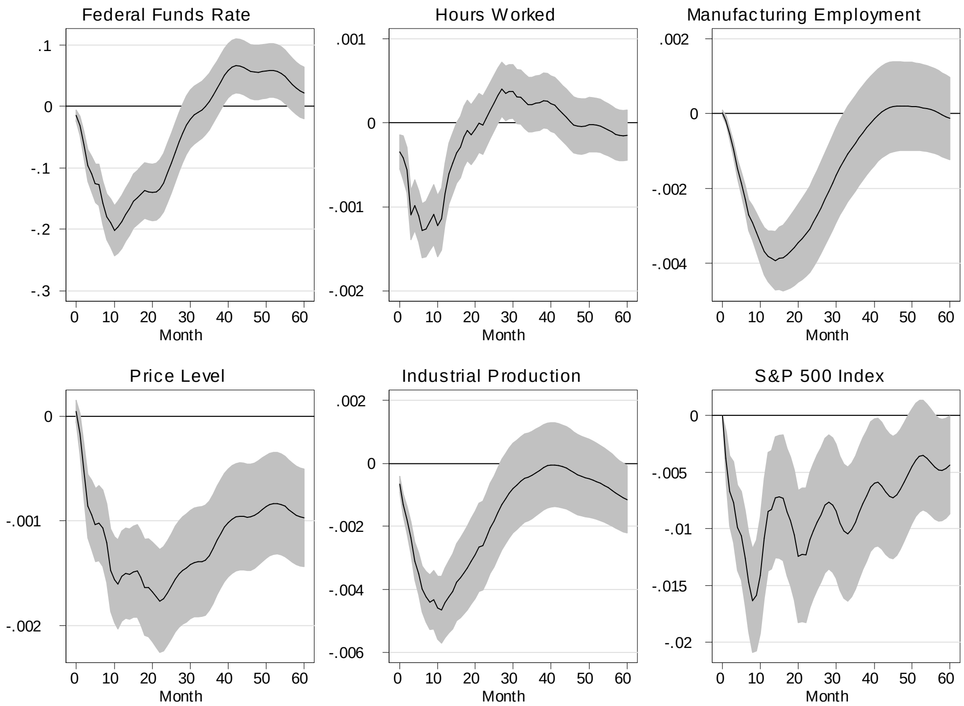

Figure $E3$ – Impulse response of federal funds rate, price level, hours worked, industrial production, manufacturing employment, and stock index from estimation of VAR, with uncertainty measured using the absolute mean forecast error. The shock is one standard deviation, and the error bands are +/- one standard error confidence bands. The sample period is 1981m9 to 2016m12.