Uncertainty and Economic Activity: Evidence from Business Survey Data

Ruediger Bachmann   Steffen Elstner   Eric Sims*

February 8, 2010

Abstract

What is the impact of time-varying uncertainty on aggregate economic activity? Using business survey data from the U.S. and Germany in structural vector autoregressions, we find that positive innovations to business uncertainty lead robustly to very prolonged declines in economic activity. In contrast, their high frequency impact is small. We thus find no evidence of the high-frequency wait-and-see effect – large declines of economic activity on impact and fast rebounds – that the recent literature associates with positive innovations to uncertainty. Rather, positive innovations to business uncertainty have similar effects as negative innovations to business confidence. Once we control for its low frequency impact, we find no statistically or economically significant effect of uncertainty innovations on aggregate activity.

JEL Codes: E30, E32, E37.

Keywords: survey data, uncertainty shocks, confidence shocks, structural VAR.

*Respectively: University of Michigan, CES-IFO and IFO; IFO Institute; University of Notre Dame. We are grateful to seminar participants at NYU Stern and Rochester for their comments. We thank Kai Carstensen and Sigrid Stallhofer from the IFO Institute, Holly Wade from the NFIB as well as Jennifer Peruzzi from Manpower for graciously providing us with their data and introducing us to the institutional backgrounds.

1 Introduction

What is the impact of time-varying uncertainty on aggregate economic activity? A basic and old idea associates innovations to uncertainty with a wait-and-see effect: if firms suddenly find themselves in a more uncertain environment they stop investing and hiring and the economy slips into a recession. This wait-and-see effect has recently attracted attention in the literature: Bloom (2009) and Bloom et al. (2009) use a quantitative RBC model with various adjustment frictions to capital and labor to argue that positive innovations to uncertainty lead to an option value to wait with investing and hiring. This effect manifests itself in very short-run fluctuations, starting with a rapid decline in aggregate activity, then a rebound phase and a prolonged overshoot after approximately four to six quarters. Figure 2 in Section 2 provides a qualitative depiction of these wait-and-see dynamics, while Figure 19 in Appendix A is a replication of the impulse response of an uncertainty shock in actual data, measured by stock market volatility, on industrial production from Bloom (2009). The three phases can be clearly seen in these graphs.

Bachmann and Bayer (2009), using data from a detailed German firm-level panel, argue that the effects in Bloom (2009) and Bloom et al. (2009) are quantitatively small and do not substantially alter unconditional business cycles dynamics. This is confirmed in Chugh (2009), who uses innovations to micro-level uncertainty, which are calibrated to U.S. manufacturing plant data, to explain the dynamics of leverage, but also finds a small business cycle impact of uncertainty shocks. Finally, Gilchrist et al. (2009) argue that in a model with financial frictions increases in uncertainty lead to an increase in the cost of capital through an increase in the firm's bond premium which is followed by a decline in investment activity.[1] These papers employ mostly quantitative models and calibration exercises to study the impact of time-varying uncertainty on aggregate economic activity. What is missing from the literature are more agnostic studies of the economic effects of uncertainty.[2]

In this paper we use monthly data from four different business surveys to investigate the relationship between uncertainty and economic activity with a structural vector autoregressions (SVAR) approach. These business surveys contain qualitative information on actual changes and expectations of the following kind: "business activity has decreased, increased or stayed the same in the last month". Specifically, we use disagreement in business expectations for the

[1]In a related, but slightly different context, Fernandez-Villaverde et al. (2009) argue that innovations to the volatility of interest rates depress economic activity in small open Latin American economies.

[2]The two exceptions we know of are: Alexopolous and Cohen (2009) who uses a narrative approach in a structural vector autoregressions framework (the incidence of the words "uncertainty" and "economy" in New York Times articles) and finds similarly to Bloom (2009) relatively high-frequency decline-rebound-overshoot dynamics; and Popescu and Smets (2009) who show, again with structural vector autoregressions and for German data, that it is innovations to risk aversion rather than innovations to uncertainty per se that explain roughly 10%-15% of output fluctuations.

Third FED District Business Outlook Survey (BOS), the Small Business Economic Trends Survey (SBETS) and the Manpower Employment Outlook Survey (MEOS) to estimate the impact of the uncertainty that actual decision makers at the business level face on aggregate economic activity.[3] The German IFO Business Climate Survey (IFO-BCS) data allow us to go one step further. In particular, we use the confidential micro data of the survey to gauge the accuracy of disagreement in expectations as a measure of the forecast error variance of individual expectations. We find that the two uncertainty measures are positively correlated and that their impact on economic activity is economically very similar and statistically indistinguishable. This justifies our use of survey disagreement as a proxy for uncertainty more broadly.

We argue that these high-frequency business survey data from specific segments of the economy are best suited to measure the direct impact of uncertainty changes on economic decision making. As discussed in the next section, wait-and-see dynamics are very short-run and rely on adjustment frictions, which make high frequency data the best candidate to detect these dynamics. *Business* survey data are advantageous because they are easily available,[4] they capture the subjective element of uncertainty, viz the mind set of actual decision makers (as opposed to outside experts), and they are an easily implementable alternative to panel GARCH techniques, for which the micro data requirements are large. Also, business survey data allow us to compare expectations and realizations of economic variables and thus – as is the case with the IFO-BCS data – utilize two complementary proxies of true ex ante uncertainty: ex ante disagreement and ex post forecast error variance. And, finally, since wait-and-see is a partial equilibrium mechanism, focusing on surveys from specific segments of the economy gives the best hope of eliminating general equilibrium effects and thus identifying the direct impact of uncertainty changes on economic decision making.

We consistently find across all four surveys that in two-variable SVARs innovations to uncertainty have very protracted, if not permanent, negative effects on economic activity in the long-run. The effect on impact, in contrast, is small. This is documented in Figure 1, where we show in the lower panel the aforementioned impulse response from the "general business conditions" uncertainty in the BOS on U.S. manufacturing industrial production. In comparison, the upper panel shows the impulse response from the "general business conditions" confidence in the BOS on U.S. manufacturing industrial production. The former looks almost like a mirror image of the latter. We will show that this finding is robust across surveys and specifica-

[3]Using a dispersion measure of expectations as a measure of uncertainty has a long tradition in the literature (mostly in the context of inflation expectations and inflation uncertainty): see, for instance, Zarnovitz and Lambros (1987), Bomberger (1996), Giordano and Soederlind (2003), Fuss and Vermeulen (2004) for a good literature overview, Bloom et al. (2009) and Popescu and Smets (2009).

[4]They are in particular easily and readily available to those who participated. Given that there is no direct financial incentive to participate, it is reasonable to assume that decision makers at the firms care very much about these survey results and let their decision making be influenced by them.

Figure 1: Confidence and Uncertainty Innovations on Manufacturing Production

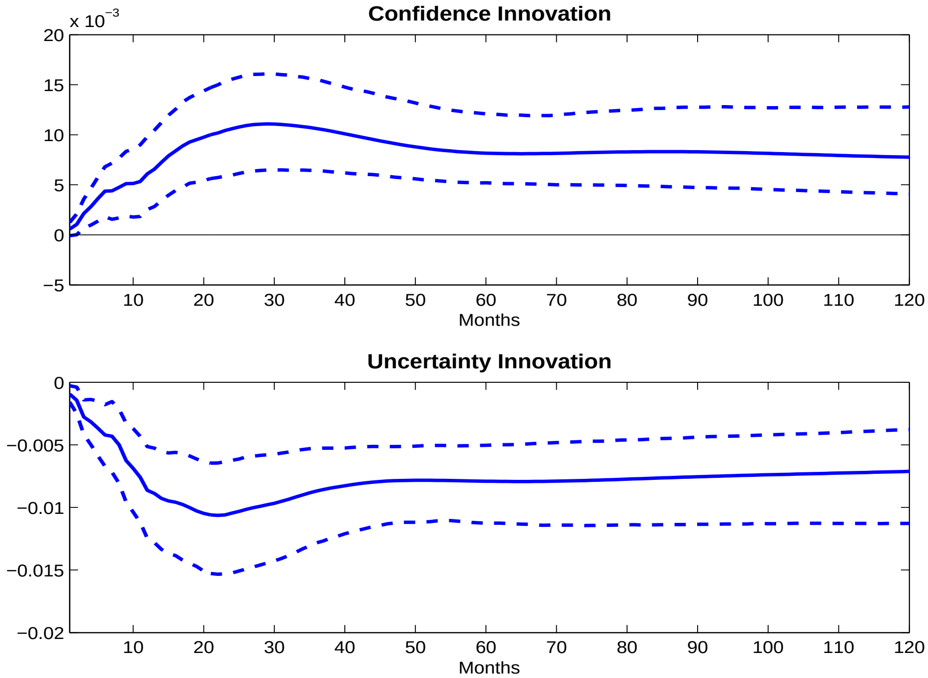

Notes: Both IRFs are based on the "general business conditions" question of the BOS. $Confidence_t \equiv Frac_t(\text{Increase}) - Frac_t(\text{Decrease})$ and $Uncertainty_t \equiv sqrt\Big(Frac_t(\text{Increase}) + Frac_t(\text{Decrease}) - \big(Frac_t(\text{Increase}) - Frac_t(\text{Decrease})\big)^2\Big)$, where $Frac_t(\text{Increase})$ is the fractions of respondents that say that general business activity six months from time t will increase. $Frac_t(\text{Decrease})$ is defined analogously. The upper panel is based on a two-variable SVAR with $Confidence$ ordered first and 12 lags. The lower panel is based on a two-variable SVAR with $Uncertainty$ ordered first and 12 lags. Manufacturing production is the natural logarithm of the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators. All confidence bands are at the 95% significance level using Kilian's (1998) bias-corrected bootstrap.

tions.[5] We then imbed business outlook confidence into the SVARs and orthogonalize business outlook uncertainty with respect to business outlook confidence and find that the impact of business outlook uncertainty on economic activity becomes economically mitigated. In a final step, we use a variation of the technique used by Barsky and Sims (2009b) to identify a shock that maximally explains aggregate activity in the medium and long-run. Specifically, we define a time horizon of five to ten years as the long-run, because this is a time-horizon where the wait-and-see effect advocated in the literature clearly has worn out (see Figure 19 in Appendix A). We find that business uncertainty is generally negatively related to this innovation to the future long-run component of economic activity and that its impact on economic activity

[5]We also show that if anything innovations to uncertainty lead to increases in employment turnover rather than decreases, as implied by wait-and-see.

essentially vanishes, once it is purged of this low-frequency predictive component. We interpret this as "bad news" for the literature that attributes strong high-frequency effects to uncertainty innovations, in particular the wait-and-see mechanism.[6]

The next section provides a sketch of the mechanisms through which uncertainty might affect economic activity, providing a benchmark against which we can compare our empirical results. The third section describes briefly our business survey data. The fourth section outlines the empirical strategy. The fifth section presents the results and illustrates the above claims in detail. A final section concludes. Various details are relegated to appendices.

2 Uncertainty and Activity: Mechanisms and Motivation

In this section we give a brief (and highly stylized) of the wait-and-see mechanism that might give rise to uncertainty-driven fluctuations. In addition to providing a benchmark against which one can compare our empirical results, this exercise will also serve to motivate our use of high frequency, sectoral data in examining the impact of uncertainty on economic activity.

Time-varying uncertainty at the firm level may have important economic consequences when there is a degree of irreversibility to firm actions. For a concrete example, suppose that a firm faces fixed costs to adjusting the size of its labor force and/or physical capital stock. Suppose, at time t, that there is a mean-preserving spread on the distribution of future demand for the firm's product. Put differently, there is an uncertainty shock – in expectation future demand conditions are unchanged, but the variance of future demand is higher. With fixed adjustments costs, higher uncertainty over future demand makes new hiring and investment less attractive. The reason why is intuitive – if a large fixed cost must be paid to adjust the firm's labor or capital, then it makes sense to minimize the number of times this cost must be paid. If the future is very uncertain (in the sense that demand could be either very high or very low relative to the present), then it makes sense to wait until the uncertainty is resolved to undertake new hiring and investment. Why pay a large fixed cost now when a highly uncertain future means that you will likely have to pay the fixed cost again?

As such, an increase in uncertainty at time t makes inaction relatively more attractive. Given a reduction in hiring, employment, and hence output, will fall through exogenous separations. As the future begins to unfold (i.e. as time progresses to $t+1$, $t+2$, etc.), demand conditions are, in expectation, unchanged. As such, there will be pent up demand for labor and capital. More formally, inaction in period t moves firms closer to their lower (S,s) triggers in subsequent

[6] Appendix A provides more details on the empirical findings of Bloom (2009) and argues that the claim "Uncertainty shocks cause recessions" has not been well established.

periods, leading to, in expectation, increased hiring in the future. This increase in hiring will lead to a rebound in economic activity, followed by an eventual return (in expectation) to the starting point. The same story can be told for capital adjustment.

Figure 2: Stylized Wait-And-See Response

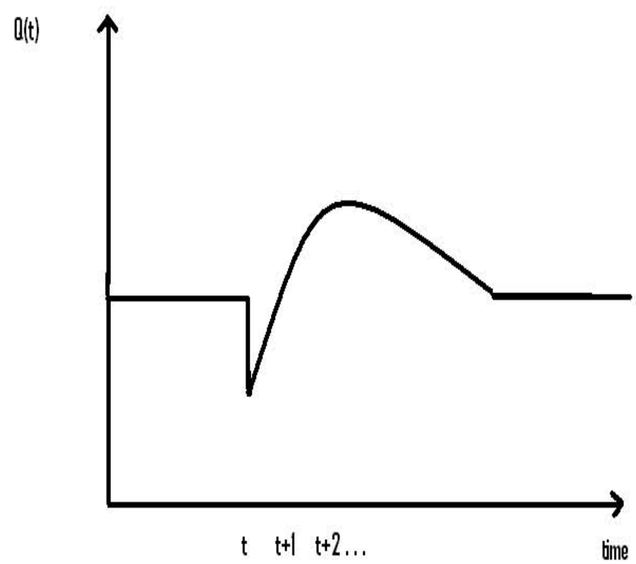

Figure 2 provides a qualitative characterization of an impulse response of firm level production to an increase in uncertainty over future demand conditions. The qualitative predictions of models with fixed costs or irreversibilities as discussed above are clear. The increase in uncertainty causes firms to want to wait and see what demand conditions are in the future, leading to a drop in employment and output on impact. In expectation, firms recover to get back to their steady state employment and output levels in the coming periods, leading to a period of overshoot following the initial decline. This figure provides a benchmark against which we can evaluate our empirical results. Figure 19 in Appendix A is a replication of Figure 2 in Bloom (2009).[7] It shows that U.S. industrial production follows a dynamic path to stock market volatility innovations that is similar to the predicted responses to increased uncertainty in the model verbally sketched here.

This informal sketch of the model above highlights two important points as pertains to our empirical work. First, the economic implications of uncertainty shocks in a wait-and-see model are decidedly high frequency in nature. In the formal model of Bloom (2009), the entire bust-

[7] Please refer to Appendix A for a more thorough documentation and discussion of Bloom's (2009) results.

boom cycle in response to increased uncertainty only takes about a year to play out. As such, any empirical study of uncertainty must make use of high frequency data, which is one of the reasons why we used monthly surveys in this paper. Second, and perhaps even more importantly, general equilibrium forces will mitigate the effects of uncertainty on economic activity. If all firms simultaneously want to shut down hiring, as an example, wages are likely to adjust in equilibrium so that at least some firms do continue hiring. Our focus on sector level data thus gives the wait-and-see mechanism a better chance of shining through than does an analysis on aggregate data, where general equilibrium forces likely play a more important role.[8]

3 Data Description

We use four different business surveys, each with different relative strengths. Specifically, we study uncertainty measures constructed from the Third FED District Business Outlook Survey (BOS), the Small Business Economic Trends Survey (SBETS), the Manpower Employment Outlook Survey (MEOS), and the German IFO Business Climate Survey (IFO-BCS). In the next subsection we briefly describe the main characteristics of each and list the main survey questions we use to measure business uncertainty. We then define the variables used in the empirical analysis. We finally provide a brief analysis of the cyclical properties of our main variables.

3.1 Measuring Business Uncertainty

BOS

The Business Outlook Survey is a monthly survey conducted by the Philadelphia FED since 1968 that is sent to large manufacturing firms in the Third FED District.[9]. According to Trebin (1998), the survey design has essentially been unaltered since its inception. The survey questionnaire is of the "box check" variety. It asks about firms' general business expectations as well as their expectations and actual realizations for various economic decision variables such as shipments, workforce and work hours. Respondents indicate whether the value of each economic indicator has increased, decreased, or stayed the same over the past month. They are also asked about their expectations for each indicator over the next six months. Only those that have 100 or more employees are asked to participate in the survey, and participation is voluntary. The survey is sent to the same individual each month, typically the chief executive, a

[8]Note that Bloom (2009) is also a partial equilibrium analysis.

[9]The Third Federal Reserve District comprises the state of Delaware, the southern half of New Jersey, and the eastern two-thirds of Pennsylvania.

financial officer or other people “in the know”. The group of participating firms is periodically replenished as firms drop out or a need arises to make the panel more representative of the industrial mix of the region.[10] Each month 100-125 firms respond. According to Trebin (1998), occasional telephone interviews are used to verify the accuracy of the survey responses.

The advantages of the BOS are its long time horizon, its focus on one consistent, economically relatively homogenous class of entities – large manufacturing firms –, an unparalleled number of questions that are useful for our research question and the fact that for each question there is a "current change" and an "expectation" version.[11] Its main drawback is the relatively small number of respondents. Nevertheless, given its advantages we use the results for the BOS as our baseline results. We mainly focus on the following two questions:[12]

Q 1 *“General Business Conditions: What is your evaluation of the level of general business activity six months from now vs. [CURRENT MONTH]: decrease, no change, increase?”*

Q 2 *“General Business Conditions: What is your evaluation of the level of general business activity [LAST MONTH] vs. [CURRENT MONTH]: decrease, no change, increase?”*

Both these questions are phrased as asking about general business conditions.[13] All other questions are about the individual firm’s condition, see the appendix.

SBETS

The Small Business Economic Trends Survey is a monthly survey conducted by the National Foundation of Independent Businesses (NFIB) which focuses on small companies across the U.S. and across all sectors. Its monthly part starts in 1986. The survey on a quarterly basis is available since the mid 1970s. As we have discussed above, we prefer the highest possible frequency to give the wait-and-see dynamics the best possible chance to appear in the data. Nevertheless, none of our results depend on that choice of frequency. In terms of participation, the October 2009 issue of the SBETS (see Dunkelberg and Wade, 2009) reports that from January 2004 to December 2006 roughly 500 business owners responded, and that the number has subsequently increased to approximately 750.[14] Almost 25% of respondents are in the retail sector,

[10]See Trebin (1998) for more detailed information.

[11]The only exception is capital expenditures, for which only an "expectation" version exists, which we do not use in this paper.

[12]The other questions we use from the BOS are documented in Appendix B.

[13]Given that the answers to the general business condition question and the shipment question, which is phrased as company-specific, are highly correlated, Trebin (1998) suspects that both are essentially indicators of firm-specific business conditions. Another interpretation is that firm-specific characteristics are of minor importance and aggregate conditions are the only conditions that firms care about. Either way, for our question what matters are firms’ overall expectations and uncertainty.

[14]The participation in the quarterly survey is higher, 1200 on average before January 2007 and 1750 thereafter.

20% in construction and 15% in manufacturing, followed by services, which ranges well above 10%. All other one-digit sectors have a single digit representation fraction. In terms of firm size, the sample contains much smaller enterprises than the BOS, which makes the SBETS and the BOS good complements in our analysis: the modal bin for the number of employees[15] is "three to five", to which over 25% of respondents belong, followed by the "six to nine"-category with roughly 20%. The highest category is "forty or more", which contains just under 10% of firms.[16] The main question we use from the SBETS is similar to the question about general business conditions in the BOS:[17]

Q 3 *"About the* $\boxed{\textit{economy in general,}}$ *do you think that* ***six months from now*** *general business conditions will be better than they are now, about the same, or worse?:* $\boxed{1}$ *Much better,* $\boxed{2}$ *Somewhat better,* $\boxed{3}$ *About the same,* $\boxed{4}$ *Somewhat worse,* $\boxed{5}$ *Much worse,* $\boxed{6}$ *Don't know. "*

Clearly, like the BOS general business conditions question, this is a question intended to ask about aggregate conditions. Notice, that one advantage of this question over its BOS version is that it is slightly more nuanced in that it allows for two "increase"- and two "decrease"-categories. Also, the fact that SBETS focusses in small enterprises makes it particularly attractive to test for wait-and-see dynamics: it is smaller companies where adjustment costs – especially for labor – should be particularly relevant and the interaction with uncertainty most prominent.

MEOS

The Manpower Employment Outlook Survey asks only one question:

Q 4 *"How do you anticipate total employment at your location to change in the three months to the end of next quarter as compared to the current quarter?"*

This is a clear limitation of the MEOS for our purposes, especially since this question is really about workforce expectations rather than aggregate or idiosyncratic general business conditions. The advantage of the MEOS is that it covers the whole U.S. and Puerto Rico and all one-digit sectors in a representative manner. The manufacturing sector is split into a non-durable and durable sector. It is from the latter that we are going to present results. The MEOS is a quarterly survey. It started in 1976. The third quarter 2009 survey covered approximately 28,000 employers in the 200 Metropolitan Statistical Areas in the U.S.

---

[15] This includes full- and part-time employees.

[16] For this and more details, see Dunkelberg and Wade (2009).

[17] The box and the bold font are also used in the original. The other questions that we use from SBETS are documented in Appendix B.

IFO-BCS

The German IFO Business Climate Survey is one of the oldest business confidence surveys available. The IFO-BCS is a monthly survey. We use only the data from the manufacturing sector.[18] Since micro data are only available in a processable form since 1980, we limit our analysis to a time horizon from then until the present.[19] The average number of respondents at the beginning of our sample is approximately 5,000; towards the end the number is about half that at 2,500.[20] Participation in the survey is voluntary and there is some fraction of firms that are only one time participants. However, conditional on staying two months in the survey, many firms continue on and this allows us to construct our measure of ex post forecast error uncertainty. Our final sample comprises then roughly 4,000 respondents at the beginning and 2,000 towards the end. The two main questions we use and that allow us the construction of a qualitative index of ex post forecast error are as follows:[21]

Q 5 *"Expectations for the next three months: our domestic production activities – without taking into account differences in the length of months or seasonal fluctuations – with respect to product XY will: increase, roughly stay the same, decrease."*

Q 6 *"Trends in the last month: our domestic production activities – without taking into account differences in the length of months or seasonal fluctuations – with respect to product XY have: increased, roughly stayed the same, decreased."*

Since this survey is the only one where we have access to the micro data, it provides us with a unique opportunity to compare ex ante uncertainty measures that are derived from cross-sectional disagreement with a qualitative index of ex post forecast error uncertainty, because we can compare firms' qualitative predictions about expected changes with their qualitative answers about realized changes. Furthermore, showing that German economic activity reacts very similarly to changes in uncertainty as its U.S. counterparts, lends additional support to the findings in Bachmann and Bayer (2009) and Popescu and Smets (2009). They are most likely not specific to Germany.

[18] There are also surveys for the construction, retail and wholesale sectors, but they have problems with longitudinal consistency.

[19] From 1991 on, the sample includes East-German firms as well. The raw time series graphs of the survey-based variables we use do not appear to exhibit structural breaks at that time. Nevertheless, we allow for a reunification dummy in all SVARs with German data.

[20] The IFO-BCS is really a survey at the product level, so that these numbers are not exactly numbers of responding firms.

[21] Here we provide a translation, for the German original see Appendix B.

3.2 Variable Definitions

With the exception of the SBETS all answers fall into two categories, $Increase$ and $Decrease$, with one or several neutral categories. We use these two categories to define two forward-looking indices and an index of current activity. We start with the forward-looking indices:

$$Confidence_t \equiv Frac_t(\text{Increase}) - Frac_t(\text{Decrease}).$$

$$Uncertainty_t \equiv sqrt\Big(Frac_t(\text{Increase})+Frac_t(\text{Decrease})-\big(Frac_t(\text{Increase})-Frac_t(\text{Decrease})\big)^2\Big).$$

Notice that $Uncertainty_t$ is the cross-sectional standard deviation of the survey responses, if the $Increase$-category is quantified by $+1$ and the $Decrease$-category by -1 and the residual categories by 0.[22] This is a standard quantification technique for qualitative survey data, under which $Confidence_t$ simply becomes the cross-sectional average of the survey responses. Next, we define a current index of economic activity from the current variables:

$$Activity_t \equiv \sum_{\tau=1}^{t} Frac_\tau(\text{Increase}) - Frac_\tau(\text{Decrease}).$$

Summing up variables that essentially measure changes is intended to capture a qualitative measure of the level of economic activity.

Finally, the availability of micro data in IFO-BCS allows us to compute an index of the ex post forecast error standard deviation, which can be viewed as an alternative measure of uncertainty. To fix ideas, we pretend for now that the production expectation question in IFO-BCS was asked only for the next month instead of the following three months. In that case, when comparing the expectation in month t with the realization in month $t+1$, nine possibilities arise: the company could have predicted an increase in production and realized one, in which case we would count this as zero forecast error. It could have realized a no change, in which case, we would quantify the expectation error as -1 and, finally, it could have realized a decrease, which counts as -2.

Table 1: POSSIBLE EXPECTATION ERRORS - ONE MONTH CASE

	$Increase_{t+1}$	$Unchanged_{t+1}$	$Decrease_{t+1}$
Expected $Increase_t$	0	-1	-2
Expected $Unchanged_t$	+1	0	-1
Expected $Decrease_t$	+2	+1	0

Notes: Rows refer to the qualitative production expectations in IFO-BCS in month t. Columns refer to the qualitative production realizations in IFO-BCS in month $t+1$.

[22] For the SBETS we adapt this quantification as follows: Much better = 2 ,Somewhat better = 1, About the same = 0, Somewhat worse = -1, Much worse = -2, Don't know = 0.

Table 1 summarizes the possible expectation errors. Of course, the production expectation question in IFO-BCS is actually for three months ahead, which we have to take into account in our index of ex post forecast error standard deviation. By way of an example, imagine the case where a firm stated in month t that its production will increase in the next three months. Suppose that in the next three months one observes the following sequence of outcomes: production has increased in $t+1$, remained unchanged in $t+2$ and finally decreased in $t+3$. Due to the qualitative nature of IFO-BCS we have to make some assumptions concerning the definition of the expectation error at the micro level. In our baseline measure, we adopt the following steps. First, we define for every month t a firm-specific activity variable over the next three months $t+3$ by the sum of the $Increase$-instances minus the sum of the $Decrease$-instances over that time period.[23] Denote the variable by $REALIZ_t$. This variable can obviously range from $[-3,3]$. Then the expectation errors are computed as:

Table 2: Possible Expectation Errors - Three Month Case

		$Expectationerror_{t+3}$
Expected $Increase_t$	$REALIZ_t > 0$	0
Expected $Increase_t$	$REALIZ_t \leq 0$	$(REALIZ_t - 1)/3$
Expected $Unchanged_t$	$REALIZ_t > 0$	$REALIZ_t/3$
Expected $Unchanged_t$	$REALIZ_t = 0$	0
Expected $Unchanged_t$	$REALIZ_t < 0$	$REALIZ_t/3$
Expected $Decrease_t$	$REALIZ_t < 0$	0
Expected $Decrease_t$	$REALIZ_t >= 0$	$(REALIZ_t + 1)/3$

Notes: Rows refer to the qualitative production expectations in IFO-BCS in month t.

Notice that the procedure in Table 2 is analogous to the one month case. Dividing by three is simply a normalization. $Expectationerror_{t+3}$ ranges from $[-\frac{4}{3}, \frac{4}{3}]$, where for instance $-\frac{4}{3}$ indicates a really negative forecast error: the company expected production to increase over the next three months, yet every single subsequent month production actually declined.

Averaging and taking the cross-sectional standard deviations at each month t allows us to compute a series of aggregate expectation errors and forecast error standard deviations, at least in a qualitative way. Specifically:

$$Uncertainty_t^{fe} \equiv STD(Expectationerror_{t+3}).$$

[23] We also experiment with a weighted sum approach: we value realizations in $t+1$ one half, realizations in $t+2$ one third and realizations in $t+3$ one sixth. This reflects that naturally when asked in t about the next three months the firm may bias its answer towards the immediate future. None of our results depends on the precise weighting scheme.

Notice the timing in the definition of $Uncertainty_t^{fe}$: the standard deviation of *realized* expectation errors in $t+3$ does not constitute uncertainty per se in $t+3$, but the knowledge of the standard deviation of these forecast errors in t is what makes decision makers uncertain at time t. It should be emphasized that this timing does not require decision makers to know anything about the future, other than that it is more or less uncertain. The advantage of $Uncertainty_t^{fe}$ over $Uncertainty_t$ is that it is based on actual "uncertain-at-time-t" innovations as opposed to potentially heterogeneous, but certain, expectations of the future. However, the raw correlation coefficient between $Uncertainty_t^{fe}$ and $Uncertainty_t$ is reasonably high for monthly data, 0.64, and when we aggregate both series up to the quarterly level to eliminate high-frequency noise the correlation becomes 0.71. This means that $Uncertainty_t$ is a reasonable proxy for true business uncertainty. Most importantly, the impulse responses on economic activity look qualitatively and quantitatively the same and are statistically not distinguishable (see Section 5.4). This lends support to the widespread procedure of proxying uncertainty with survey disagreement.

3.3 Cyclicality of Business Survey Variables

In this section, we report basic cyclical properties of the survey-based variables. All these variables – $Confidence_t$, $Uncertainty_t$, $Activity_t$ and $Uncertainty_t^{fe}$ – were seasonally adjusted with the SAS X12 procedure, which is itself an adaptation of the U.S. Bureau of the Census X-12-ARIMA Seasonal Adjustment method. Table 3 displays the cyclical properties of the various survey-based uncertainty measures we described in Section 3.2. It shows that our survey-based uncertainty measures are countercyclical. This confirms previous findings by Bloom (2009), Bloom et al. (2009), Chugh (2009) and Bachmann and Bayer (2009), from various data sources, such as stock market volatility and balance-sheet-based cross-sectional measures of uncertainty. The relatively low numbers for MEOS can be explained by the fact that this employment-based disagreement measure can only be a proxy for the disagreement measures based on the general business situation, which we have shown are good, but imperfect measures of true underlying uncertainty. And the fact that the MEOS uncertainty measures are not contemporaneously correlated with employment activity is simply a result of the usual lag to shocks and other activity variables that frictional employment adjustment induces: the correlations between the MEOS uncertainty measures and employment reach the order of magnitude of the ones with production at employment lagged one or two quarters. Finally, the discrepancy between the first column and second column for the IFO uncertainty measures is partly, just like for the other surveys, simply the result of an imperfect representation of the entire population by the survey sample. For Germany, however, we find that the industrial production measure exhibits a lot of high frequency noise-like movements, which in part contributes to this low correlation. The correlation becomes more negative when we aggregate up to the quarterly frequency.

Table 3: CYCLICAL PROPERTIES OF $Uncertainty_t$ AND $Uncertainty_t^{fe}$

Uncertainty-Measure	Industrial Production	$Activity_t$
General Situation-$Uncertainty_t^{BOS}$	-0.28	-0.47
Shipments-$Uncertainty_t^{BOS}$	-0.27	-0.29
General Situation-$Uncertainty_t^{SBETS}$	-0.57	-0.71
Durable Manufacturing-$Uncertainty_t^{MEOS}$	-0.22	-0.08
Non-Durable Manufacturing-$Uncertainty_t^{MEOS}$	-0.18	-0.03
Production-$Uncertainty_t^{IFO}$	-0.14	-0.50
Production-$Uncertainty_t^{fe^{IFO}}$	-0.05	-0.45

Notes: This table displays the contemporaneous raw correlations between the survey-based variables in the rows and the appropriate differences of two different activity measures in the columns. Industrial production measures are logged. The General Situation-$Uncertainty_t^{BOS}$ measure, based on Q 1, is paired with the month-over-month difference of the (seasonally adjusted) manufacturing industrial production index from the OECD main economic indicators and the General Situation-$Activity_t^{BOS}$ measure based on Q 2. The Shipments-$Uncertainty_t^{BOS}$ measure, based on Q 7 (see appendix), is paired with the month-over-month difference of the (seasonally adjusted) manufacturing industrial production index from the OECD main economic indicators and the Shipments-$Activity_t^{BOS}$ measure based on Q 10 (see appendix). The General Situation-$Uncertainty_t^{SBETS}$, based on Q 3, is paired with three-months differences in the log of (seasonally adjusted) total industrial production from the FED and the $Activity_t^{SBETS}$ measure based on Q 13 (see appendix). Since the latter asks for a three-months difference, we take three-months differences for industrial production as well. The Durable Manufacturing-$Uncertainty_t^{MEOS}$ measure, based on Q 4, is paired with the quarter-over-quarter difference of the (seasonally adjusted) durable manufacturing industrial production index from the OECD main economic indicators and the log quarter-over-quarter difference of (seasonally adjusted) durable manufacturing employment from the BLS-CES database (the analogous holds for the Non-Durable Manufacturing-$Uncertainty_t^{MEOS}$-measure). The Production-$Uncertainty_t^{IFO}$ measure, based on Q 5, is paired with the month-over-month difference of the (seasonally adjusted) manufacturing industrial production index from the German Federal Statistical Agency and the $Activity_t^{IFO}$-measure based on Q 6. For the definition of Production-$Uncertainty_t^{fe^{IFO}}$, see Section 3.2; it is paired with the same activity measures as the Production-$Uncertainty_t^{IFO}$ measure.

Table 4 displays the cyclical properties of the various survey-based (differenced) activity measures we described in Section 3.2. They are, not surprisingly, procyclical.

Table 4: CYCLICAL PROPERTIES OF $Activity_t$

Activity-Measure	Industrial Production
General Situation-$Activity_t^{BOS}$	0.55
Shipments-$Activity_t^{BOS}$	0.46
Sales-$Activity_t^{SBETS}$	0.37
Production-$Activity_t^{IFO}$	0.28

Notes: This table displays the contemporaneous raw correlations between the differenced survey-based variables in the rows and the appropriate differences of industrial production indices. Industrial production measures are seasonally adjusted and logged. The General Situation-$Activity_t^{BOS}$ measure, based on Q 2, is paired with the month-over-month difference of the manufacturing industrial production index from the OECD main economic indicators. The Shipments-$Activity_t^{BOS}$ measure, based on Q 10 (see appendix), is paired with the month-over-month difference of the manufacturing industrial production index from the OECD main economic indicators. The Sales-$Activity_t^{SBETS}$ measure, based on Q 13 (see appendix), is paired with three-months differences in the log of total industrial production from the FED. The Production-$Activity_t^{IFO}$ measure, based on Q 6, is paired with the month-over-month difference of the manufacturing industrial production index from the German Federal Statistical Agency.

4 Empirical Strategy

We are interested in characterizing, in a relatively atheoretic way, the dynamic responses of economic activity to various measures of business uncertainty. We do so by use of vector autoregressions, similarly to other work by Alexopolous and Cohen (2009), Bloom (2009), and Popescu and Smets (2009). This exercise will allow us to draw inferences about the dynamic consequences of surprise movements in uncertainty, and, in so doing, to say something about the importance of uncertainty shocks in business cycle dynamics as well as to speak to the suitability of existing DSGE models for capturing these features of the data.

Let $\mathbf{x}_t$ be a vector of endogenous variables.[24] We assume that these variables can be characterized by a finite order vector autoregressive process, written in lag operator notation as:

$$\mathbf{A}(\mathbf{L})\mathbf{x}_t = \mathbf{u}_t.$$

$\mathbf{A}(\mathbf{L})$ is a matrix lag polynominal of order p, while $\mathbf{u}_t$ is a vector of innovations of conformable size, with variance-covariance matrix Σ_u. It is necessary to orthogonalize the innovations in

[24] As a baseline, we estimate these VARs with all variables appearing in levels, which is consistent with the recommendations of Sims, Stock, and Watson (1990). Levels estimation will produce consistent estimates of VAR impulse responses. Robustness checks where clearly trending series (such as sectoral production or employment) enter in first differences or as deviations from a trend produce nearly identical results.

some way. We assume that there exists a linear mapping, given by the square matrix $\mathbf{B}$, from "structural" shocks, ε_t, into reduced form innovations, $\mathbf{u}_t$:[25]

$$\mathbf{u}_t = \mathbf{B}\varepsilon_t$$

After normalizing the variance-covariance matrix of structural shocks to an identity matrix, the elements of the "impact matrix" can be found as the solution to: $\mathbf{BB}' = \Sigma_u$. We discuss the restrictions used to identify $\mathbf{B}$ in more depth below.

For a first pass at the data, we estimate VARs with two variables – a measure of sector level uncertainty and a measure of sector level economic activity, where the measure of activity is chosen so as to most closely correspond with the survey questions from the forecast.[26] Two variables in the reduced form system is identical to the benchmark systems estimated in Alexopolous and Cohen (2009). Since we are only interested in characterizing the responses to uncertainty shocks, the inclusion of other variables is unnecessary, though our results are quite robust to larger dimensional systems. After estimating the reduced-form VAR with a suitable number of lags,[27] we order the measure of uncertainty first in a recursive Choleski decomposition to identify $\mathbf{B}$. In the models now coming into popularity (e.g. Bloom, 2009), uncertainty innovations have their largest impact on economic activity on impact and at high frequencies, and so they must be allowed to affect economic activity on impact.

In the two variable VARs, we consistently find that positive innovations to uncertainty are prognostic of long-lasting and, in some cases at least, apparently permanent reductions in various measures of economic activity. This finding suggests that a significant fraction of innovations in uncertainty may simply be correlated with "news", broadly defined, about future long-run economic activity (see Figure 1 in the introduction). We thus consider alternative specifications in which the variable $Confidence_t$ is included in the vector of endogenous variables. Ordering $Confidence_t$ before the uncertainty measure attempts to control for any information about long-run economic activity contained in our uncertainty measures. The impulse responses of activity to uncertainty orthogonal to confidence is then more likely to isolate the high-frequency impacts of time-varying uncertainty.

To deal more explicitly with the long-run component of uncertainty innovations, we adopt a variant of the Barsky and Sims (2009b) strategy for identifying "news shocks". It is important to note that we are not attempting to identify news shocks in the sense formalized by Beaudry and Portier (2006), Barsky and Sims (2009a), Sims (2009) and Kurman and Otrok (2010), which

[25]We put "structural" in quotation marks because we only interpret these innovations as structural insofar as they are mutually uncorrelated.

[26]For example, if the survey question concerns "general conditions" we use a sector level of production as the activity measure, whereas if the survey concerns hiring, we use a measure of employment.

[27]Twelve for monthly data and four for quarterly data. Our basic results remain unaltered if we optimize the lag number with AIC or BIC.

are shocks to future economic state variables orthogonal to the present. Rather, what we do in the context of our VARs is to isolate structural disturbances which have predictive ability for movements in real activity at different frequencies – in particular the "long-run". We identify a shock which best explains movements in real sectoral activity at low frequencies (in particular between five and ten years, which are horizons beyond which uncertainty shocks, if the simple wait-and-see channel is operating, should have much of an effect). We then orthogonalize uncertainty measures with respect to this disturbance, with the hope of extracting large high frequency impacts on economic activity.

More formally, rewrite the VAR as a reduced-form moving average process:

$$\mathbf{x}_t = \mathbf{A}(\mathbf{L})^{-1}\mathbf{u}_t$$

For notational ease, let $\mathbf{C}(\mathbf{L}) = \mathbf{A}(\mathbf{L})^{-1}$, and express $\mathbf{x}_t$ as a structural moving average process for some arbitrary orthogonalizing matrix, $\widetilde{\mathbf{B}}$:

$$\mathbf{x}_t = \mathbf{C}(\mathbf{L})\widetilde{\mathbf{B}}\epsilon_t$$

For this arbitrary orthogonalizing matrix, the fraction of the forecast error variance of variable i (indexing elements of $\mathbf{x}_t$), to structural shock j (indexing elements of ϵ_t), at horizon h can be expressed as:

$$\Omega_{i,j}(h) = \frac{\sum_{\tau=0}^{h} \mathbf{C}_{i,\tau}\widetilde{\mathbf{B}}_j\widetilde{\mathbf{B}}_j'\mathbf{C}_{i,\tau}'}{\sum_{\tau=0}^{h} \mathbf{C}_{i,\tau}\Sigma_u\mathbf{C}_{i,\tau}'}$$

Here $\widetilde{\mathbf{B}}_j$ refers to the j^{th} column of $\widetilde{\mathbf{B}}$, while $\mathbf{C}_{i,\tau}$ denotes the row of reduced-form moving average coefficients of variable i at horizon τ. The denominator is the total forecast error variance of variable i at horizon h, and is independent of the manner in which the reduced form innovations are orthogonalized.

The Barsky and Sims (2009b) identification is essentially a principal components approach to choosing elements of **B**. They identify a news shock as the structural shock which best explains future movements in TFP subject to the restriction that it have no impact on TFP on impact. We take a less restrictive approach and define our long-run component innovation to be a shock which maximally explains the variation in a measure of economic activity at "low" frequencies (equivalently long horizons), to be made more precise below.

In particular, we identify the long-run component by choosing the j^{th} column of the orthogonalizing matrix, which we denote as $\mathbf{B}_j$, to maximize the total explanatory power for a measure of economic activity from horizon $\underline{h}$ to horizon $\overline{h}$:

$$
\begin{aligned}
\mathbf{B}_j &= \arg\max \sum_{h=\underline{h}}^{\overline{h}} \Omega_{i,j}(h) \\
& s.t. \\
\mathbf{B}_j &\in \widetilde{\mathbf{B}}
\end{aligned}
$$

Note that we only identify one column of the orthogonalizing matrix, not the entire matrix. As such, the maximization problem is done subject to the constraint that the solution belong to the space of possible orthogonalizing matrixes, denoted by $\widetilde{\mathbf{B}}$. In practice this is accomplished by first picking an arbitrary orthogonalizing matrix (say a Choleski decomposition), and then post-multiplying it with an orthonormal matrix of conformable size. We then conduct the maximization by searching over the columns of the space of orthonormal matrixes. For more details, see Barsky and Sims (2009b), Sims (2009), or Kurman and Otrok (2010). The identification in the present paper differs from the Barksy and Sims (2009b) identification in that there is no contemporaneous orthogonality restriction between the long-run component and activity and because the maximization problem begins at some horizon other than on impact. In this sense, the identification here is very similar to Uhlig (2004), where he identifies shocks in a VAR context which best explain movements in output at business cycle frequencies (i.e. from horizons $\underline{h} = 6$ to $\overline{h} = 32$ quarters).

Armed with an estimate of $\mathbf{B}_j$, we can then form a time series of structural shocks to the long-run component from the reduced-form innovations, and can trace out the implications of these shocks for a variety of different variables. We robustly find that these shocks are positively correlated with forecast expectations and negatively correlated with uncertainty – across all data sets. We then ask the following question: controlling for the long-run component in the uncertainty measures (identified as above), what effects do movements in uncertainty have for real economic activity at the sectoral level? To accomplish this, we simply regress the reduced-form innovations to uncertainty on the identified long-run component series and take the residual. We then regress a measure of economic activity on that residual, and trace out the dynamic implications and compute the forecast error variance decomposition of movements in activity attributable to uncertainty orthogonalized with respect to the long-run component. This approach – which essentially controls for the low frequency predictability of uncertainty for economic activity – should give the uncertainty series the best chance of leading to large, high frequency fluctuations in output, consistent with the implications of wait-and-see models. As shown in the next section in far greater detail, we robustly find that there is little significant response of activity to uncertainty once we control for its low frequency predictive power.

5 Results

In this section we present and discuss our main empirical results. We robustly find that innovations to business uncertainty are associated with small and slowly-building reductions in measures of economic activity. This finding is robust across all surveys we use. Controlling for lower frequency predictability – either by orthogonalizing uncertainty with respect to confidence or by explicitly identifying the main source of low frequency movements in activity, as discussed in Section 4 – renders the responses of economic activity to uncertainty essentially zero at all horizons. These findings are difficult to reconcile with the wait-and-see channel from uncertainty to aggregate dynamics, as emphasized in the recent literature.

We sequentially analyze survey and economic activity data from a variety of different sources, beginning with the Federal Reserve Bank of Philadelphia Third District Business Outlook Survey, followed by the Small Business Economic Trends survey and the Manpower Employment Outlook. We conclude with an analysis of the German IFO Business Climate Index. In addition to providing verification of our main qualitative findings from data in another country, the IFO index allows us to compare and contrast the differences that arise when using ex ante survey disagreement versus the ex post forecast error standard deviation as measures of business uncertainty. We show that our SVAR results are nearly identical using either measure as a metric for uncertainty.

5.1 Third FED District Business Outlook Survey

We begin and indeed focus our analysis on the Federal Reserve Bank of Philadelphia's Third District Business Outlook Survey. Figure 1 in the introduction shows impulse responses from two variable SVARs with US manufacturing industrial production and either the confidence series or the uncertainty series. The SVARs are estimated with 12 lags (i.e. a year's worth), all variables in levels, and the variables ordered such that innovations to the survey measure influence economic activity on impact but not vice versa.[28] The dashed lines are 95 percent confidence bands from Kilian's (1998) bias-corrected bootstrap.

As noted in the introduction, the impulse response of production to an innovation to uncertainty is slightly negative on impact with effects that build over time. The peak decline is at about 1 percent, occurring about two years after impact, with no tendency to revert. Even at

[28]One might be worried that uncertainty should not affect economic activity on impact because of various information or decision lags. For instance, one might assume that companies know uncertainty only through the published surveys themselves. To test whether this affects our results, we present the analog to the lower panel of Figure 1 with economic activity ordered first in Figure 23 in Appendix C. From this graph it is clear that timing does not drive our results.

horizons outside of standard business cycle frequencies the impulse response is negative at all conventional significance levels. As the upper panel of the figure shows, the response of production to uncertainty is roughly the mirror image of its response to a confidence innovation.

Figure 3: Uncertainty Innovation on the BOS General Activity Index

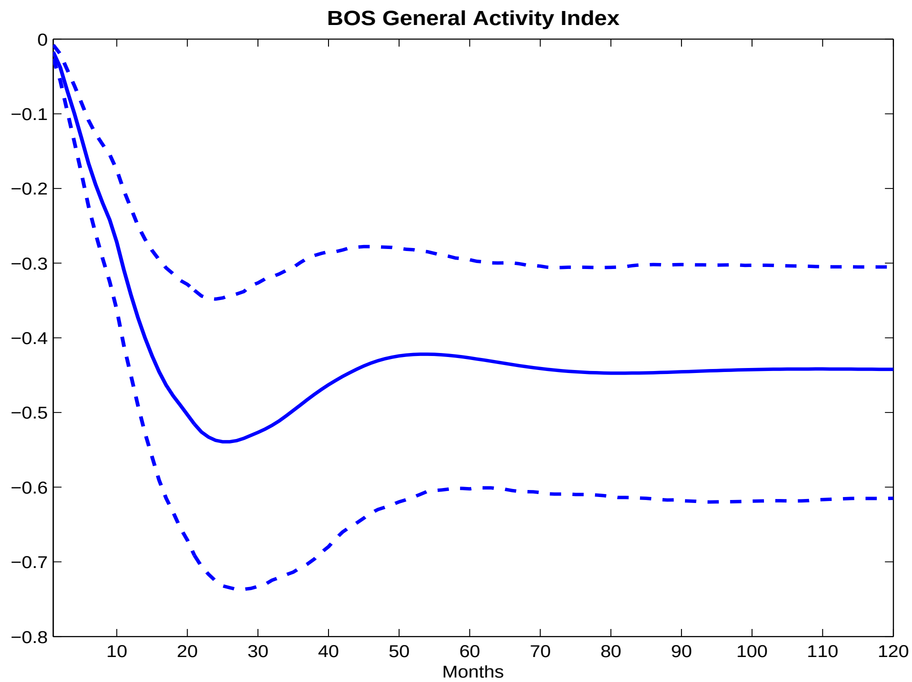

Notes: The IRF is based on a two-variable SVAR with $Uncertainty$ (based on Question 1 of the BOS) ordered first and $Activity$ (based on Question 2 of the BOS) ordered second and 12 lags. All confidence bands are at the 95% significance level using Kilian's (1998) bias-corrected bootstrap.

Figure 3 provides corroborating evidence in the form of a completely different measure of sectoral economic activity. In addition to the forward-looking confidence question, Question 2 of the BOS asks about current conditions relative to the recent past. We estimate a bivariate SVAR with the uncertainty measure and this current activity index, again with uncertainty ordered first. The response is strikingly similar to that using overall manufacturing production.

Figure 4 plots the forecast error variance decomposition of activity due to uncertainty innovations at various horizons. At horizons under one year, innovations to uncertainty account for less than 20 percent of the forecast error variance of manufacturing production. This fraction rises at horizons beyond a couple of years, providing some quantitative support for the notion that uncertainty provides some useful information about changes in activity in the relatively far off future, but has little to say about fluctuations at high frequencies.

Figure 4: Forecast Error Variance Decomposition

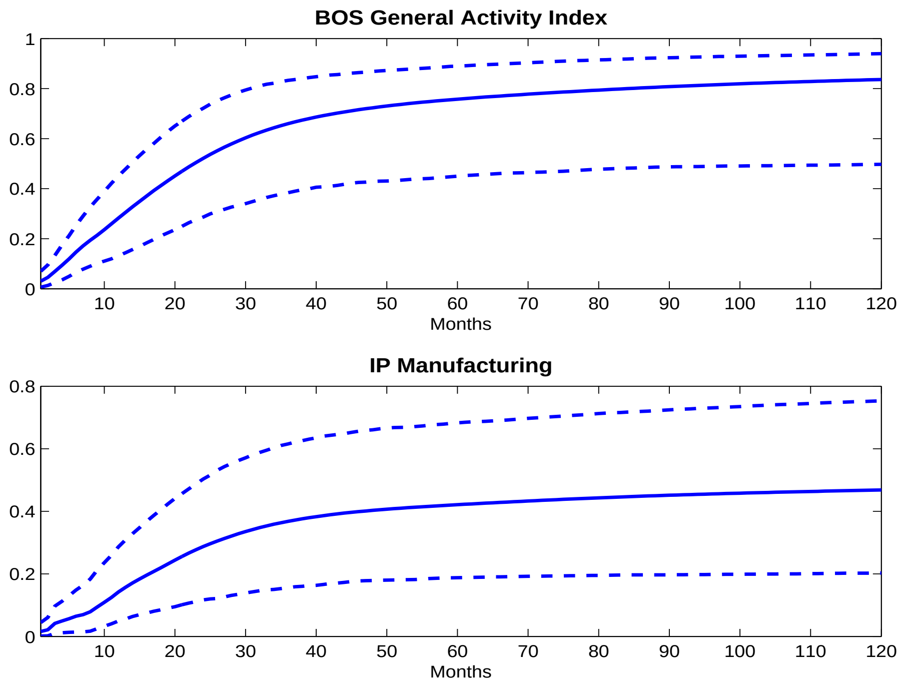

Notes: see notes to Figures 1 and 3.

We provide further robustness checks in Figures 5 and 6, using a variety of different measures of economic activity, both qualitative and quantitative. Figure 5 shows the responses of a variety of different qualitative activity indexes from the BOS survey: shipments, workforce and work hours. The impulse responses are qualitatively (and quantitatively) similar across the three different activity indexes, further reinforcing our main finding that uncertainty innovations are prognostic of long-lasting reductions in economic activity, with little effect at high frequencies.

Figure 6 shows responses of various different measures of employment within the manufacturing sector.[29] Yet again, the qualitative nature of the responses is very similar. Employment moves very little on impact with a tendency to decline further as the horizon grows. There is no evidence consistent with a significant wait-and-see channel.

[29] We thank Mark Bils for the suggestion to distinguish between production and non-production workers, given that they might be subject to different hiring and firing costs.

Figure 5: Uncertainty Innovations on Various BOS Activity Indices

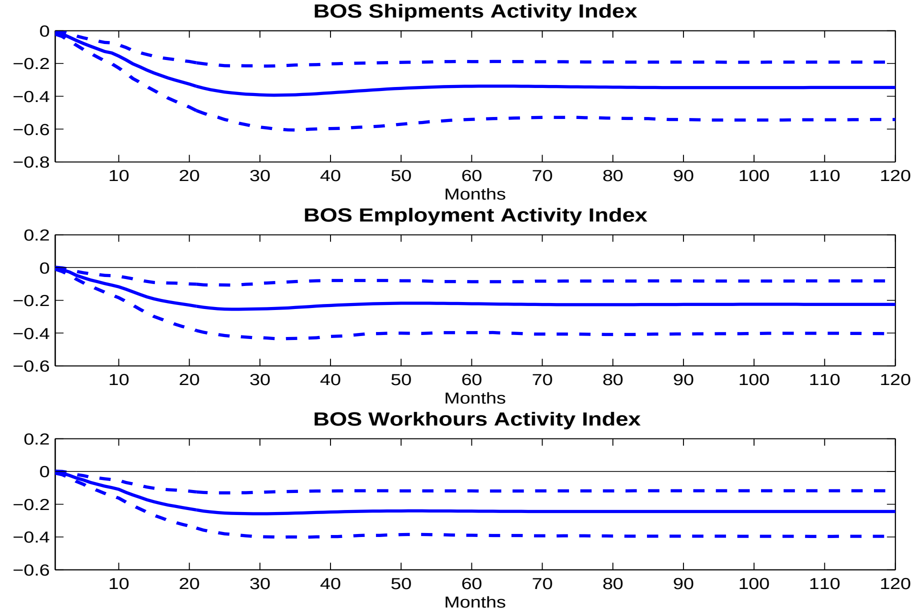

Notes: see notes to Figure 3. The activity indices for the three panels are based on Questions 10, 11 and 12, respectively.

Figure 6: Uncertainty Innovations on Manufacturing Employment

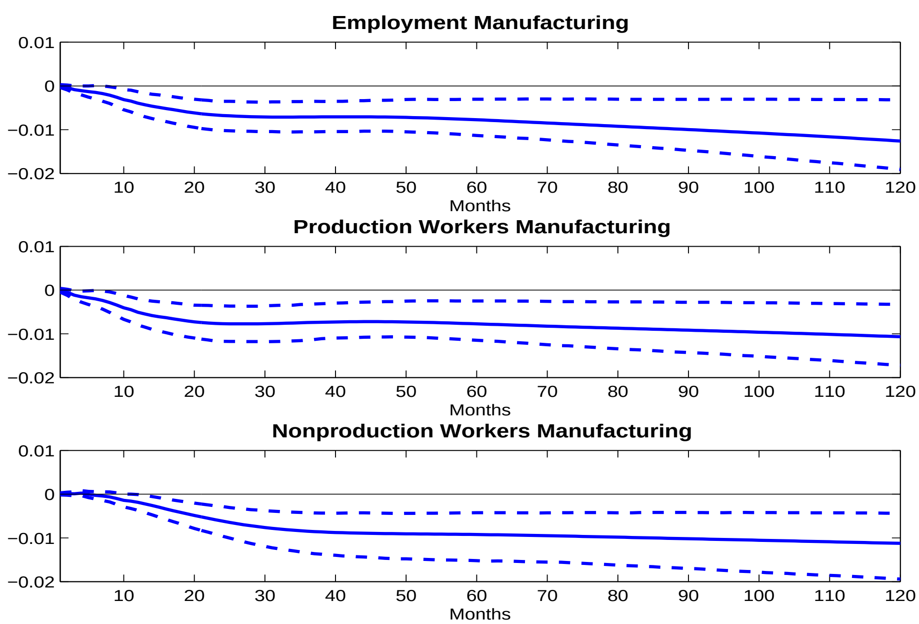

Notes: see notes to Figure 3. The employment measures are seasonally adjusted and logged and come from the BLS-CES data base.

Figure 7: Uncertainty Innovation on Aggregate Investment

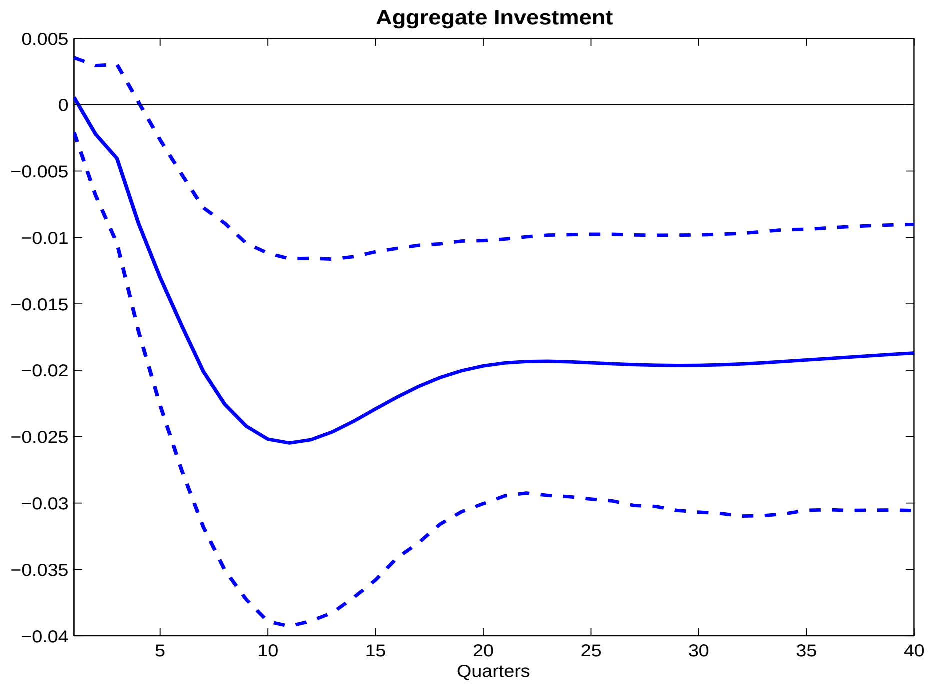

Notes: see notes to Figure 3. Since investment is available only at the quarterly frequency, we aggregate the uncertainty measure up to the quarterly level and estimate with four lags. Investment is seasonally adjusted real private nonresidential fixed investment from the NIPA data. It is logged.

In Figure 7 we show an impulse response from a bivariate SVAR featuring the BOS baseline uncertainty measure and the log of aggregate US investment. The response shown is that of investment to uncertainty, with uncertainty ordered first. Wait-and-see theories of the transmission from uncertainty shocks to business cycles emphasize the partial irreversibility of certain firm actions, in particular investment in new physical capital. As such, if the wait-and-see channel were important, we would observe a large reduction in investment followed by a quick recovery in response to an uncertainty shock, similarly to the output response in Figure 19 in Appendix A. The response of aggregate investment to the uncertainty innovation looks nothing like those theories would predict. Rather, consistent with our other results, investment moves little on impact (indeed the point estimate is positive), followed by a period of sustained reductions, with no obvious tendency for reversion, even at very long horizons.

Another direct prediction of the wait-and-see theory is that job turnover – defined as the sum of job creation and job destruction – should decline following an increase in uncertainty (wait and do nothing). Yet again, our qualitative survey data are inconsistent with this prediction. Figure 8 shows the response of the extensive margin of job turnover[30] to an innovation

[30]Given the (S,s) rationalization for wait-and-see, the extensive margin of job turnover is the correct measure.

in uncertainty. The point estimate on and near impact is positive and insignificant from zero, turning more significant at horizons well beyond one year.

Figure 8: Uncertainty Innovation on BOS Job Turnover Index

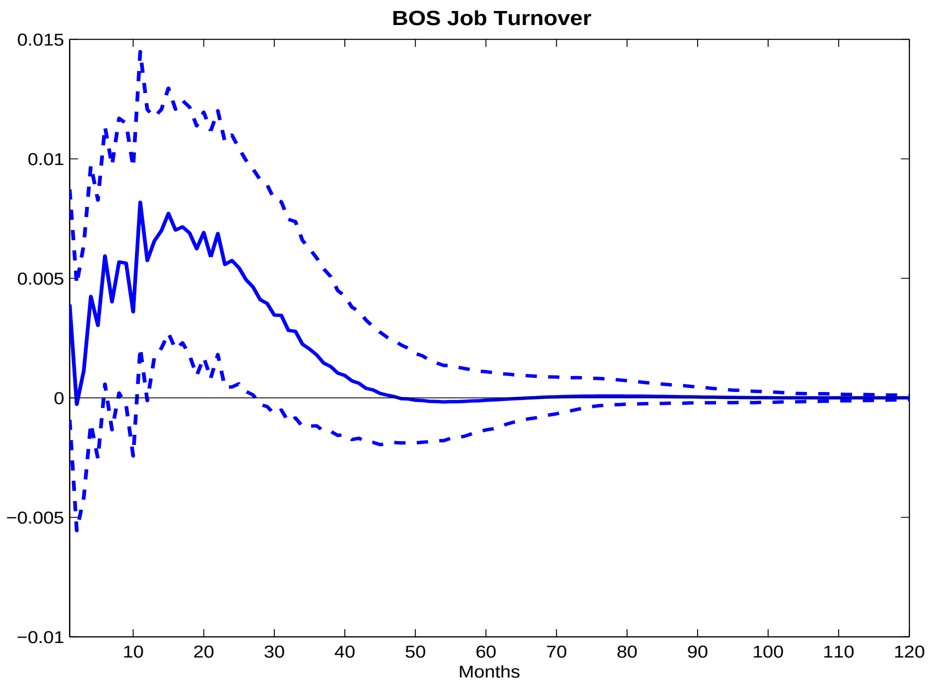

Notes: see notes to Figure 3. The turnover variable is based on Question 11 and is defined as: $Turnover_t \equiv Frac_t(\text{Increase}) + Frac_t(\text{Decrease})$.

Finally, we consider alterations to our baseline measure of uncertainty.[31] Bloom's (2009) preferred uncertainty measure is a $(0,1)-$indicator, taking on values of one for periods in which stock market volatility is abnormally high and zero for periods in which it is not. While his results are roughly invariant using either this discrete measure of uncertainty or the actual continuous measure of volatility, there are reasons to consider the discrete measure as well. First, the month-to-month movements in measured uncertainty may partially reflect noise, making the detection of important changes in uncertainty more difficult. Secondly, models with non-convexities typically predict that it is large increases in uncertainty that matter, not the frequent increases and decreases observed in most months.

To consider this possibility, we construct a $(0,1)-$indicator similarly to Bloom (2009), with the indicator taking on the value of one if our measure of uncertainty is more than one standard

[31] Figures 24 and 25 in Appendix C show more of that. Figure 24 uses uncertainty measures based on the forward-looking questions on shipments, the workforce and work hours. Figure 25 uses information-theoretic entropy instead of the standard deviation as the measure of uncertainty. This has been advocated by Rich and Tracy (2006). For an exact definition, see Appendix C.

deviation above its mean. There are 60 such observations, or about 12% of the sample. We use a cutoff of one standard deviation above the mean to determine the indicator variable because of the left-skewness of the distribution of uncertainty in these data. For example, there are only nine observations more than 1.5 standard deviations above the mean out of 488 observations in the sample. At any rate, using higher cutoffs, or using a cutoff in which uncertainty is not high relative to the entire sample but instead relative to the recent past, produces very similar results. Using indicator variables in a VAR analysis is similar to the "event study" identifications not only in Bloom (2009) but also, for example, in Ramey and Shapiro (1998).

Figure 9: Uncertainty Innovation (Indicator Variable) on Manufacturing Production

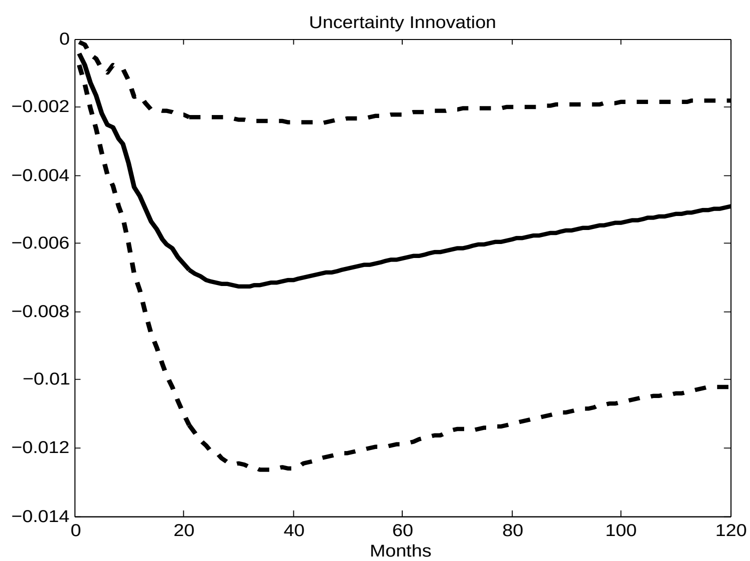

Notes: see notes to Figures 1 and 3. The uncertainty variable here is an indicator variable that takes on a value of one, if the measure of uncertainty, which is based on Question 1, is one standard deviation above its mean, and zero otherwise.

Figure 9 shows the impulse responses of manufacturing production to this indicator measure of uncertainty in a two variable SVAR, with uncertainty ordered first. It is very similar to the benchmark responses – the high frequency effects on production are small followed by sustained periods of low growth. There is no evidence of the kind of "bust-boom" path as implied by the wait-and-see mechanism.

There are two main results from our analysis thus far – one negative, and one positive. The negative result is that there is little evidence supporting the wait-and-see mechanism. The positive result, on the other hand, is that innovations to uncertainty appear to contain significant predictive information for the future path of economic activity. To explore these conclusions further, as well as to give uncertainty a better chance of leading to important wait-and-see-type dynamics, we include a qualitative measure of confidence in the SVAR, orthogonalizing uncertainty with respect to confidence. As noted previously, doing so should, to the extent to which confidence is informative about the long-run, control for the long-run predictive component of uncertainty, thereby allowing any existing wait-and-see channel to shine through.

Figure 10: Uncertainty Innovations Orthogonalized to Confidence Innovations

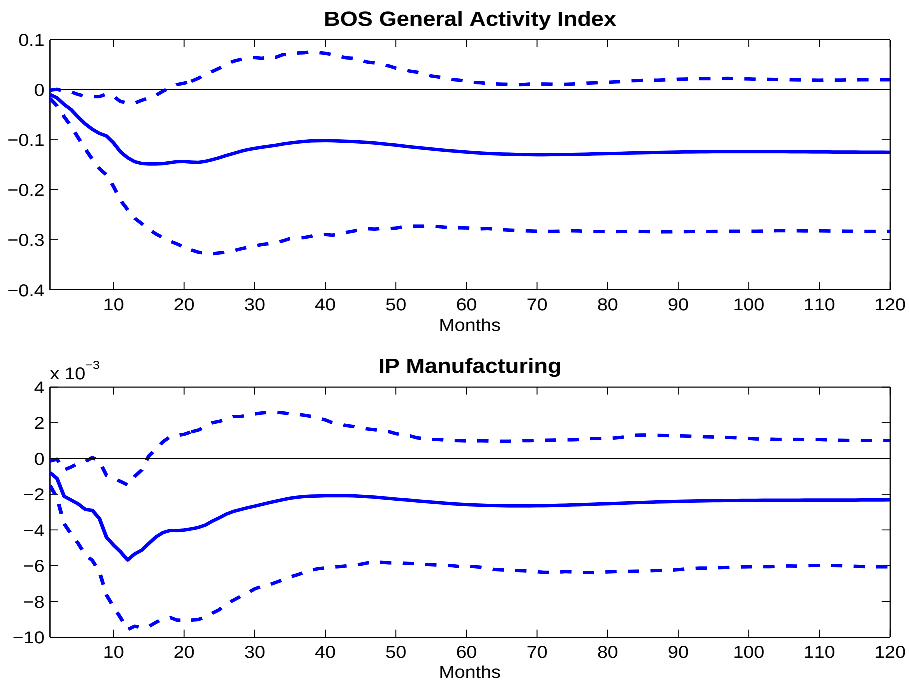

Notes: see notes to Figures 1 and 3. $Confidence_t$ is ordered first, then $Uncertainty_t$, then the activity variable.

Figure 10 depicts impulse responses of two different measures of activity – the general BOS activity index and manufacturing production – to an uncertainty innovation orthogonalized with respect to confidence. Orthogonalizing with respect to the confidence series lowers the quantitative magnitude of the responses of activity to uncertainty, but it appears to do little to change the qualitative nature of the responses. Compare the responses shown in Figure 10 to the comparable ones in Figures 1 and 3. The response of activity, however measured, to uncertainty is small and insignificant on impact, followed by further reductions, and then some evidence of reversion at longer horizons. Thus, orthogonalizing with respect to the confidence series does not point to an important wait-and-see effect. The impact of uncertainty on economic activity does become less statistically significant once we control for information about future long-run economic activity contained in the confidence series.[32]

[32] Figure 26 in Appendix C shows the analogous picture for labor productivity in manufacturing.

Figure 11: Uncertainty Innovations Orthogonalized to the Long Run Component - Manufacturing Production

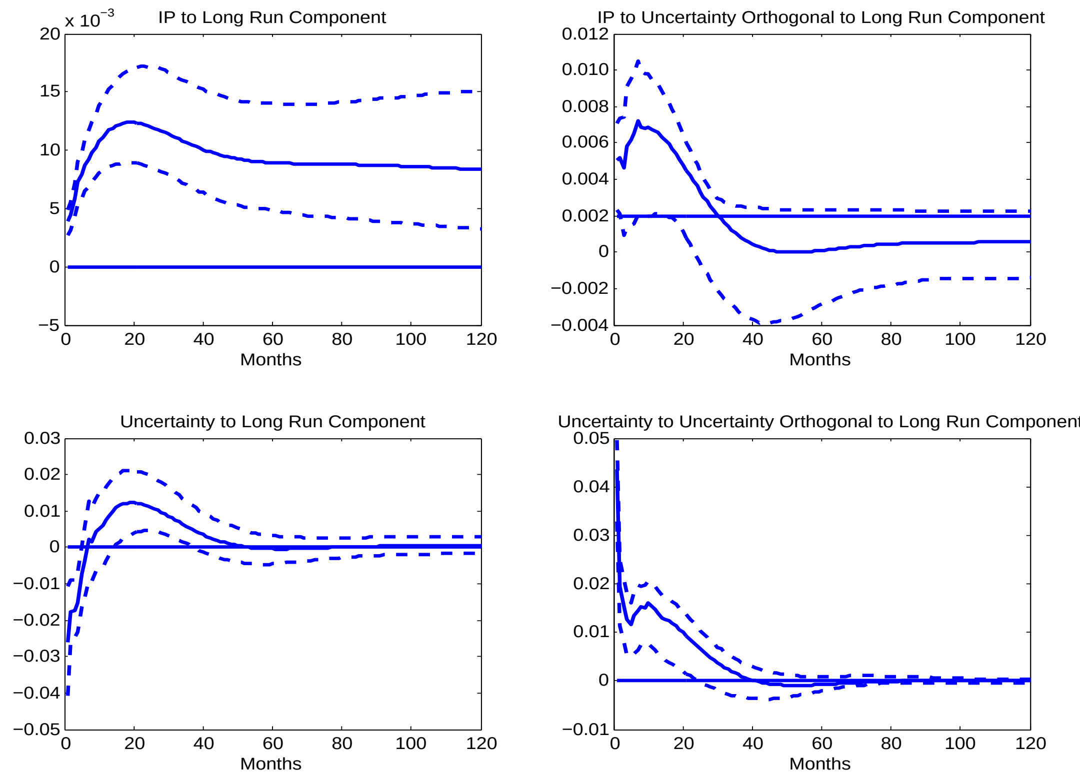

Notes: see notes to Figure 1. The reduced form VARs are estimated with three variables – confidence, uncertainty, and economic activity. Confidence and uncertainty are based on Question 1. The activity series is the natural logarithm of the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators. Identification is explained in detail in Section 4.

To address these issues more formally, we adopt a variant of the Barsky and Sims (2009b) identification strategy to fully "purge" uncertainty innovations of their long run predictive component, thereby hoping to isolate the wait-and-see channel. As described more completely in Section 4, we identify a shock which accounts for the maximum forecast error variance of economic activity at horizons from five to ten years. We then orthogonalize the uncertainty innovation with respect to this shock in order to see what is left. The reduced form VARs are estimated with three variables – confidence, uncertainty, and economic activity. We show responses of both uncertainty and activity to the long-run component shock, as well as the responses to the uncertainty innovations orthogonalized with respect to the long-run component. The results are depicted graphically in Figure 11, which uses manufacturing production as the activity variable.[33]

[33]Figure 11 in Appendix C shows the same exercise when the general BOS activity measure, based on Question 2, is used.

The results are similar regardless of the activity measure. The innovation to the future long-run component of economic activity leads to a small and growing response of activity, with little or no tendency to revert to zero even at very long horizons.[34] Consistent with our earlier findings, the long-run component is significantly negatively correlated with innovations to uncertainty; the unconditional correlation between the two is -0.51. Perhaps most interestingly, the response of activity to uncertainty orthogonal to the future long-run component of economic activity is statistically insignificant from zero at all horizons. Further, the point estimates on impact and at high frequencies are positive, as opposed to the negative responses that would be predicted by the wait-and-see channel. In short, once we control for the long horizon predictive component of uncertainty for economic activity, there is almost no effect of uncertainty on economic activity and no evidence in support of the wait and see channel.

5.2 Small Business Economic Trends

In this subsection we present results from the Small Business Economic Trends Survey. The results are similar to those for the BOS. Since the SBETS complements the BOS insofar as it covers small companies, this lends additional support to our findings.

Figure 12: Uncertainty Innovations on SBETS Sales Activity Index

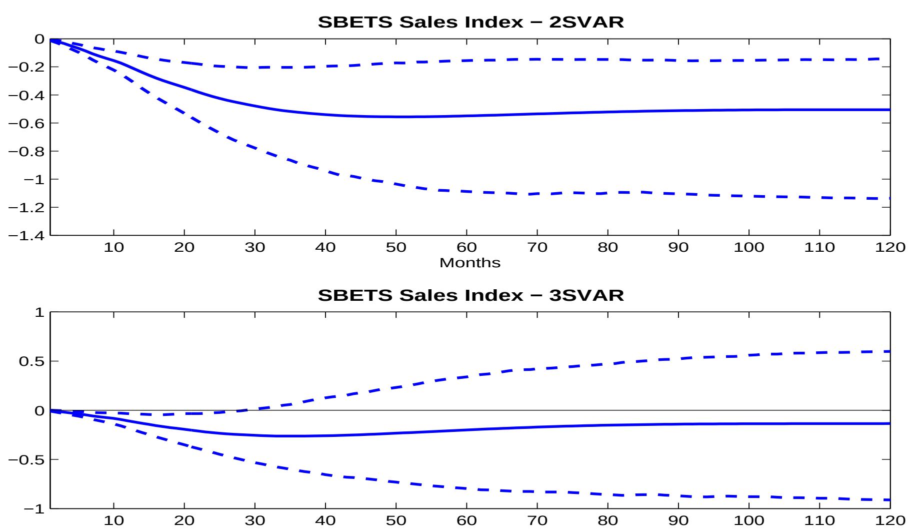

Notes: Confidence and uncertainty are based on Question 3. The activity variable is based on Question 13. The upper panel is based on a two-variable SVAR with uncertainty ordered first and then activity. The lower panel contains confidence ordered first on top of the variables in the upper panel.

[34]Note that, unlike in the strict "news shock" identifications of Sims (2009), Barsky and Sims (2009b), economic theory makes no predictions as pertains the sign of the impact effects of news on activity here. This is because we do not impose that the "news shock" be orthogonal to the present, but rather just predictive for the long-run future.

The upper panel of Figure 12 shows the response of the SBETS sales index to an innovation to uncertainty, ordered first in a bivariate recursive SVAR. Similarly to the previous section, there is little or no high frequency impact followed by a period of sustained negative growth.The lower panel of the figure shows the same response to an uncertainty innovation orthogonalized with respect to the corresponding confidence series. As in the BOS, the response is quantitatively smaller – particularly at lower frequencies – and less statistically significant.

Figure 13: Uncertainty Innovation on SBETS Job Turnover Index

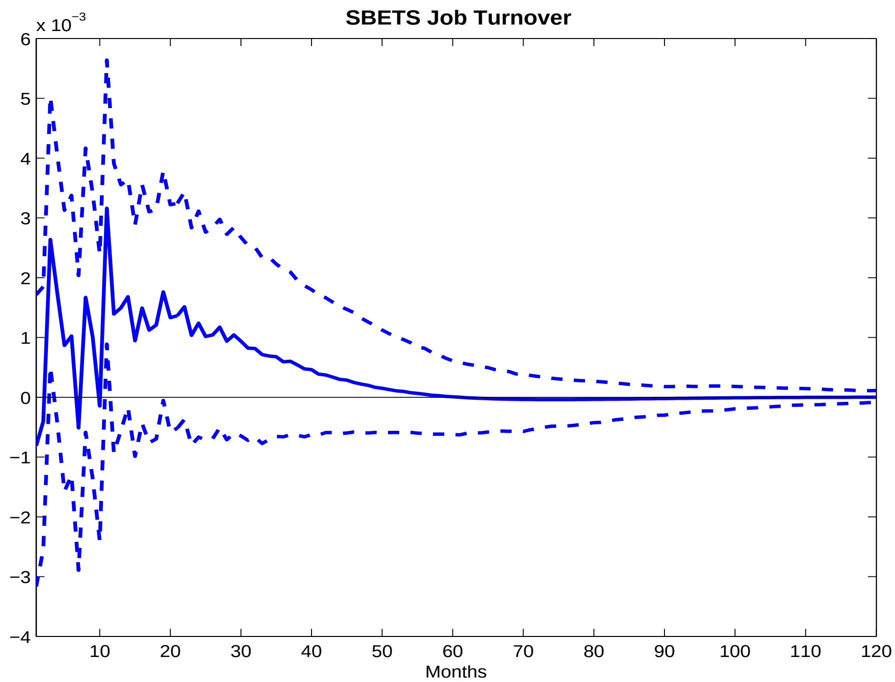

Notes: see notes to Figure 12. The IRF is based on a two-variable SVAR with uncertainty ordered first and then job turnover. Job turnover is based on Question 14.

Figure 13 is similar to Figure 8 from the previous subsection. It shows the impulse response of the job turnover measure to an innovation to uncertainty. As before, to the extent to which job turnover reacts to dispersion at all, it rises (at least the point estimate), which is inconsistent with wait-and-see theories of uncertainty shocks.

Finally, Figure 14 shows responses of activity and uncertainty using the more formal long-run component identification. As with the Third Fed District data, the impulse response of activity to dispersion orthogonalized with respect to this long-run component is, for all intents and purposes, zero at all horizons.

Figure 14: Uncertainty Innovations on SBETS Sales Activity Index Orthogonalized to the Long Run Component

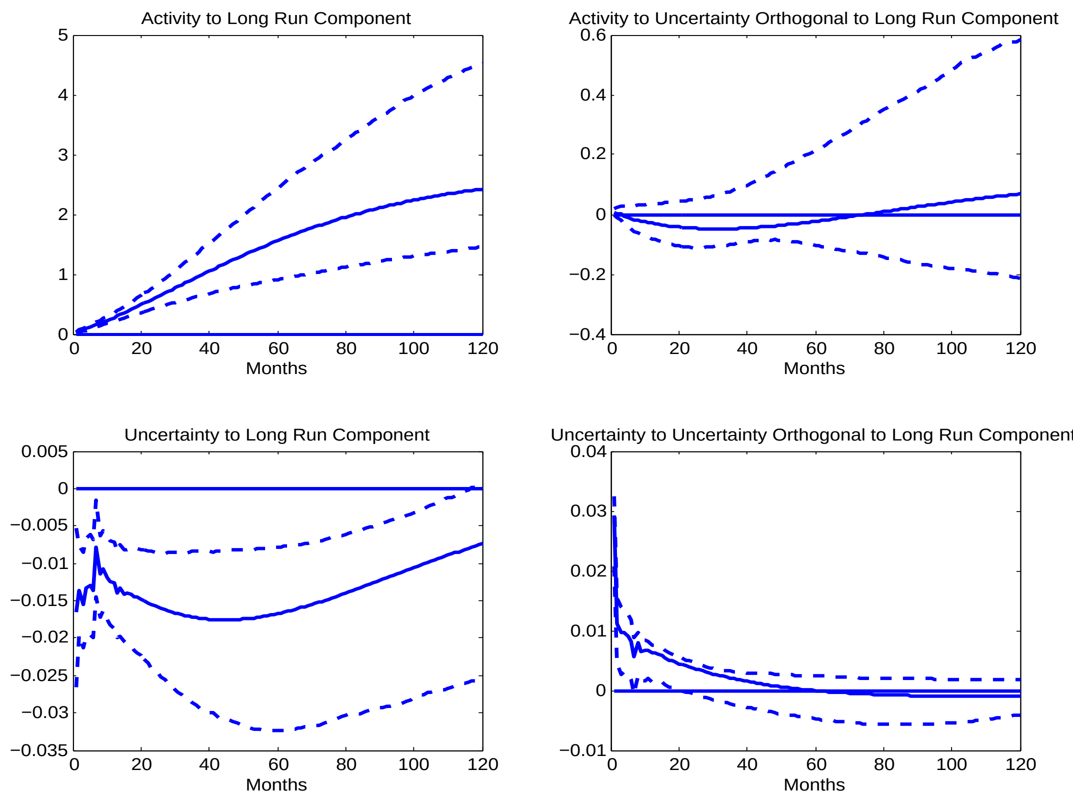

Notes: see notes to Figure 12. The the unconditional correlation between the long-run component and uncertainty innovations is -0.48.

5.3 Manpower Employment Outlook

Here we present results from the Manpower Employment Survey. As noted in Section 3.1, this survey asks specifically about hiring trends and not economic activity or production per se, and so the measure of economic activity we use in the SVARs of this section is employment. For the sake of brevity we limit our analysis to the durable goods manufacturing sector; the non-durable goods sector behaves similarly.

Figure 15 shows impulse responses to uncertainty innovations from two and three variable SVARs. As before, in the two variable SVAR uncertainty is ordered first in a Choleski decomposition. In the three variable SVAR, uncertainty is orthogonalized with respect to the confidence variable. Somewhat surprisingly, the impulse response of sectoral employment to uncertainty is roughly the same regardless of orthogonalizing with respect to confidence – little or no impact effect followed by sustained reductions that persist for quite some time. While these responses are not at all indicative of an important wait-and-see channel, they do suggest, at least for this sector and this survey, that uncertainty innovations convey long-run information above and beyond that which is available in the confidence series.

Figure 15: Uncertainty Innovation on Durable Manufacturing Employment

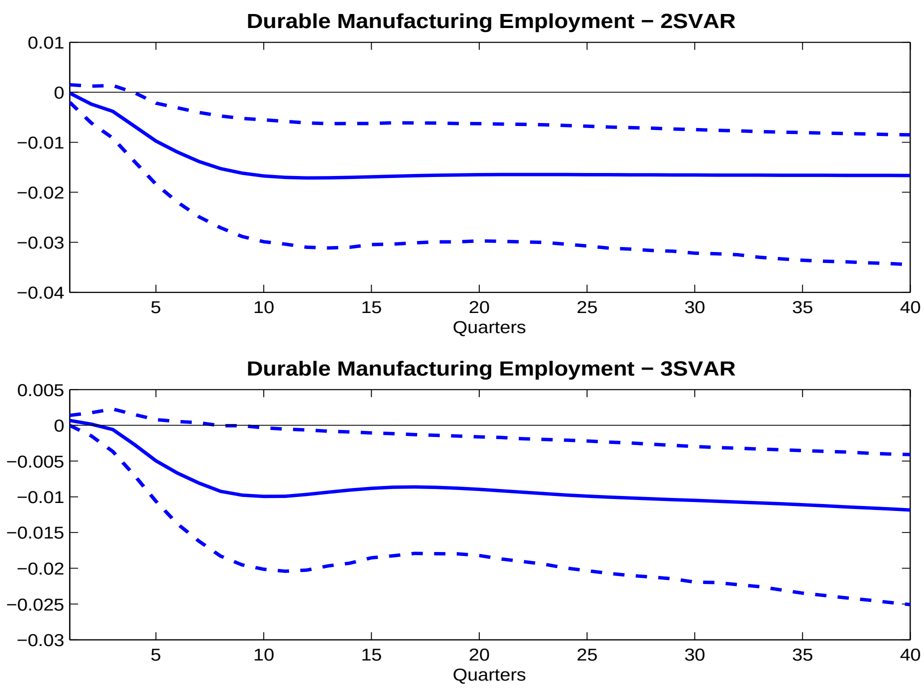

Notes: Confidence and uncertainty are based on Question 4. The activity variable is the seasonally adjusted and logged durable manufacturing employment series from the BLS-CES data base. The upper panel is based on a two-variable SVAR with uncertainty ordered first and then activity. The lower panel contains confidence ordered first on top of the variables in the upper panel.

Figure 16 shows the responses for the Manpower data for the modified Barsky and Sims (2009b) identification. The identified innovation to the future long-run level of durable goods manufacturing employment leads to little impact effect followed by sustained and apparently permanent growth in employment. As would be expected, given earlier results, uncertainty falls sharply on impact in response to favorable news about future employment. The correlation coefficient between the two is -0.41. The upper right panel shows the response of employment to uncertainty purged of its long run predictive component. Here the impact effect is mildly statistically significant and positive. This mild positive increase is quite transitory and not economically large. At any rate, it is inconsistent with the predictions of wait-and-see.

Figure 16: Uncertainty Innovation on Durable Manufacturing Employment Orthogonalized to the Long Run Component

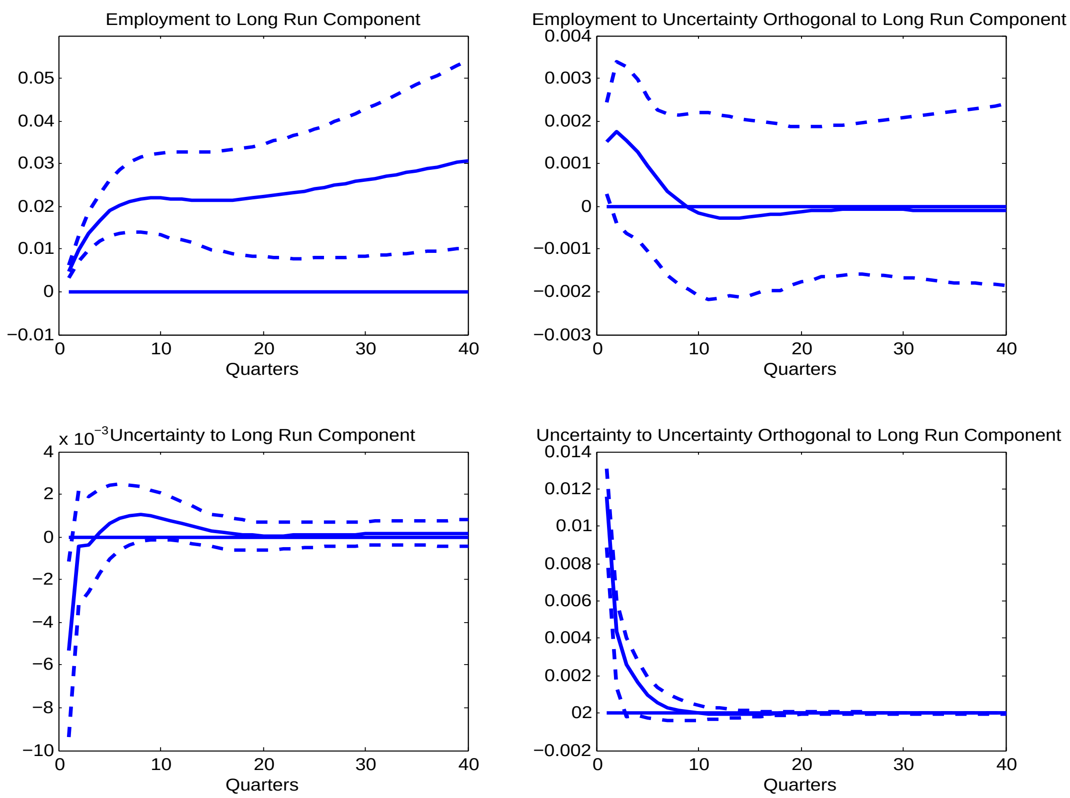

Notes: see notes to Figure 15.

5.4 IFO Business Climate Index

Finally, we present results from the German IFO Business Climate Index. The main features of the results using uncertainty measures are quite similar to our earlier findings from a variety of U.S. data. The main advantage of these German data is that we also have access to micro data, which allows us to compute a measure of uncertainty based on the ex post forecast error standard deviation and compare the results with the uncertainty measure based on ex ante disagreement. We show that the results are very similar. This serves two purposes. First, it provides independent corroboration of our results with U.S. data. Second, it provides a test for our use of a disagreement measure as a measure of actual uncertainty, when micro data are unavailable.

Figure 17 shows responses, for two variable SVARs, to the two types of uncertainty innovations, which are both based on Questions 5 and 6. The SVARs here include a dummy variable from 1991 on to account for structural breaks associated with the German reunification, though our results are quite insensitive to alternative ways of dealing with that event. We use an index of industrial production (upper panel) and the IFO production activity index (lower panel) as activity variables. There are two important results: First, we see that using either measure of activity the responses of activity to the two different measures of uncertainty are quite similar to each other, in fact statistically indistinguishable. This clearly lends support to our proxying of uncertainty with a disagreement measure, when micro data are unavailable. Second, when the IFO production activity index is used the results are also similar to what was reported in the various U.S. data. The impact effects on activity are small, with the peak negative response occurring roughly two years subsequent to the shock. When industrial production for the manufacturing sector is used as the measure of sectoral activity, the responses are somewhat different from before in that there is more reversion towards zero in the response of activity at longer horizons. This is particularly true for the upper confidence bands which almost look like the wait-and-see impulse responses in Figure 19 in Appendix A. Nevertheless, the point estimates are at least inconsistent with a strong rebound phase.

Figure 17: Uncertainty Innovations on Manufacturing Production and the IFO General Activity Index

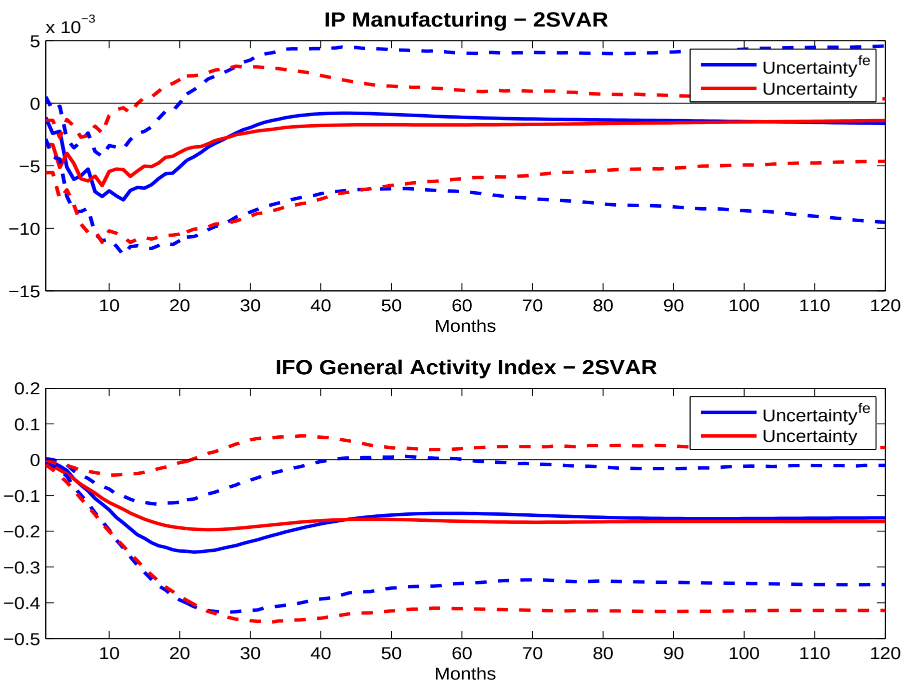

Notes: Uncertainty is based on Questions 5 and 6. The activity variable in the upper panel is the seasonally adjusted and logged manufacturing industrial production series from the German Federal Statistical Agency. The activity variable in the upper panel is based on Question 6. In both SVARs uncertainty is ordered first. We include a dummy variable from 1991 to account for the German reunification.

Figure 18, however, shows that even the case of IFO data with industrial production for the manufacturing sector is ultimately not indicative of high-frequency wait-and-see dynamics.[35] If they were, the modified Barsky and Sims (2009b) methodology, which controls for a long-run component only, would not be able to eliminate the effect of uncertainty. Figure18 shows the opposite. As before, activity responds to neither measure of uncertainty, after orthogonalization with respect to the long-run component of activity. Again, both measures of uncertainty behave similarly in the Barsky and Sims (2009b) identification method. And finally, the German results under the Barsky and Sims (2009b) identification are again very similar to the ones in the U.S., which suggests that the findings in Bachmann and Bayer (2009) and Popescu and Smets (2009) are not specific to Germany.

Figure 18: Uncertainty Innovations on Manufacturing Production Orthogonalized to the Long Run Component

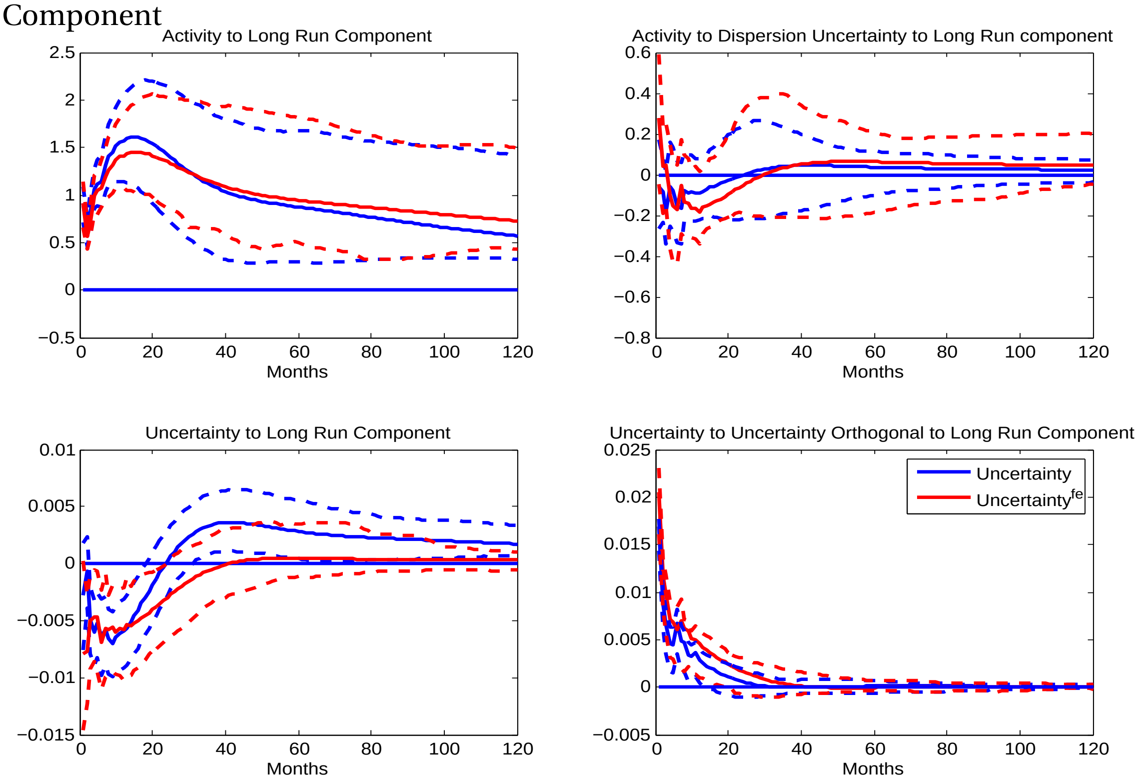

Notes: see notes to Figure 17. The confidence measure is based on Question 5. The unconditional correlation between the long-run component and innovations to uncertainty is -0.16 for $Uncertainty_t$ and -0.35 for $Uncertainty_t^{fe}$.

[35]See Figure 28 in Appendix C for the case with the IFO production activity index.

6 Final Remarks

We use two measures of business uncertainty from high-frequency, sectoral business surveys to investigate the impact of business uncertainty on economic activity in an agnostic structural vector autoregressions framework. Specifically, we construct ex ante disagreement measures from three complementary U.S. and one German business surveys. For the latter, we also use the survey micro data to compute an index of ex post forecast error variance. We find that both measures of uncertainty are positively correlated and exhibit the same impact on economic activity. In a series of SVAR specifications, we show that innovations to business uncertainty have a protracted effect on economic activity that vanishes once we control for measures of business confidence or, using a variant of the identification technique in Barsky and Sims (2009b), directly for the long-run component of economic activity. These results are inconsistent with the high-frequency wait-and-see dynamics that the literature has recently advocated. Rather, given our findings that uncertainty innovations are robustly negatively correlated with the long-run component of economic activity, uncertainty shocks appear to have similar effects as bad news shocks.

We should be clear that our analysis shows that innovations to uncertainty have important economic effects, but they appear inconsistent with a simple wait-and-see mechanism. We think that our SVAR analysis provides important restrictions that the growing literature with business cycle models and uncertainty shocks at least in their partial equilibrium specifications have to satisfy. For instance, a wait-and-see effect to R&D investment or a wait-and-see effect amplified by embodied technology in a vintage capital model are potential structural interpretations of our findings.

References

[1] Alexopoulos, M. and J. Cohen (2009). "Uncertain *Times*, Uncertain Measures", mimeo.

[2] Bachmann, R. and C. Bayer (2009). "Firm-specific Productivity Risk over the Business Cycle: Facts and Aggregate Implications", mimeo.

[3] Barsky, R. and E. Sims (2009a). "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence", NBER-WP 15049.

[4] Barsky, R. and E. Sims (2009b). "News Shocks", NBER-WP 15312.

[5] Beaudry, P. and F. Portier (2006). "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, **96**, 1293–1307.

[6] Bloom, N. (2009). "The Impact of Uncertainty Shocks", *Econometrica*, 77/3, 623–685.

[7] Bloom, N., M. Floetotto and N. Jaimovich (2009). "Really Uncertain Business Cycles", mimeo.

[8] Bloom, N., Bond, S. and Van Reenen, J.(2007),"Uncertainty and Investment Dynamics", *Review of Economic Studies*, 74, 391–415.

[9] Bomberger, W., (1996),"Disagreement as a Measure of Uncertainty", *Journal of Money, Credit and Banking*, 28/3, 391–415.

[10] Christiano, L., R. Motto and M. Rostagno (2007). "Financial Factors in Business Cycles", mimeo.

[11] Chugh, S., (2009). "Firm Risk and Leverage-Based Business Cycles", mimeo.

[12] Dunkelberg, W. and H. Wade (2009). "NFIB Small Business Economic Trends", October 2009.

[13] Gilchrist, S., J. Sim and E. Zakrajsek (2009). "Uncertainty, Credit Spreads and Aggregate Investment", mimeo.

[14] Giordano, P., and Soederlind, P.(2003),"Inflation Forecast Uncertainty", *European Economic Review*, 47, 1037–1059.

[15] Fernandez-Villaverde, J., P. Guerron-Quintana, J. Rubio-Ramirez and M. Uribe (2009). "Risk Matters: The Real Effects of Volatility Shocks", mimeo.

[16] Fuss, C., and P. Vermeulen (2004). "Firm's Investment Decisions in Response to Demand and Price Uncertainty", ECB WP 347.

[17] Kilian, L. (1998). "Small Sample Confidence Intervals for Impulse Response Functions" *Review of Economics and Statistics*, 80, 218–230.

[18] Kurman, A., and C. Otrok (2010). "News Shocks and the Slope of the Term Structure of Interest Rates", mimeo.

[19] Popescu, A. and F. Smets (2009). "Uncertainty, Risk-Taking and the Business Cycle" , mimeo.

[20] Ramey, V. and M. Shapiro (1998). "Costly Capital Reallocation and the Effects of Government Spending" *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–194.

[21] Rich, R. and J. Tracy (2006). "The Relationship between Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts", Federal Reserve Bank of New York Staff Report No. 23.

[22] Sims, C., J. Stock and M. Watson (1990). "Inference in Linear Time Series Models with Some Unit Roots", *Econometrica*, 58, 113–144.

[23] Sims, E. (2009). "Expectations Driven Business Cycles: An Empirical Evaluation", mimeo.

[24] Trebing, M. (1998)."What's Happening in Manufacturing: "Survey Says..."", *Federal Reserve Bank of Philadelphia Business Review*, September/October 1998, 15–29.

[25] Uhlig, H. (2004). "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association*, 2, 361–371.

[26] Zarnovitz, V., and L. Lambros (1987),"Consensus and Uncertainty in Economic Prediction", *The Journal of Political Economy*, 95/3, pp. 591–621.

A Appendix A - Wait-and-See

Figure 19: Replication of the Evidence for Wait-and-See in Bloom (2009)

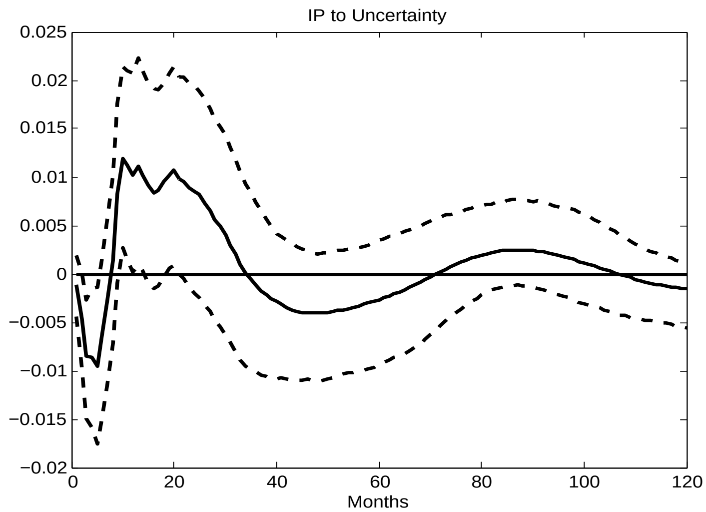

Notes: This IRF is a replication of Figure 2 in Bloom (2009). It shows the response of U.S. industrial production with respect to a stock market volatility shock, simply adapted to our time horizon of 10 years. The variables in the estimation order are log(S&P500 stock market index), a stock-market volatility indicator, Federal Funds Rate, log(average hourly earnings), log(consumer price index), hours, log(employment), and log(industrial production). All variables are Hodrick-Prescott (HP) detrended (λ = 129,600). The main stock-market volatility indicator is constructed to take a value 1 for a month with particularly high volatility, see Figure 20 below.

It is worth noting that Figure 19 shows that the wait-and-see dynamics in Bloom (2009) are barely statistically significant. The next three figures cast some doubt on the claim in Bloom (2009) that uncertainty shocks are a major cause of recessions. Figure 20 indicates the months of high volatility that Bloom (2009) uses to construct his volatility index. Two findings are important: First, 8 out of 15 high volatility dates are outside of NBER recessions. Second, only in the 1975, 1980 and 1991 recessions do we find that volatility was high at the beginning of a recession, in no case was volatility high prior to a recession. In 1982 and 2001 the high volatility events took place towards the end of the recession, and similarly for 1975 we have a higher volatility event towards the end (the Franklin National financial crisis) than at the beginning (OPEC I, Arab-Israeli War). Finally, the NBER beginning of the most recent recession is the fourth quarter of 2007, whereas stock market volatility spiked in late 2008. This seems to indicate that recessions cause stock market volatility, but not the other way around. As a simple test of whether stock market volatility is really an independent structural shock, we repeat in Figures 21 and 22 the analysis in Figure 19, which is in turn identical to the exercise in Bloom (2009), adding an NBER recession dummy to the SVAR. This allows us to compute impulse responses of industrial production to an innovation to stock market volatility, conditional on being in or out of an NBER recession. As can be seen, the IRF is quite different between the two cases. Outside of recessions, stock market volatility has no impact on aggregate activity, in fact the point estimate of the impact is positive. [36] In a multi-shock world, where structural uncertainty shocks sometimes cause recessions and sometimes do not (because other shocks prevent the recession from occurring), we would nevertheless expect the *conditional* response of activity to uncertainty to be the same. This is clearly not the case, casting doubt on the claim that stock market volatility innovations constitute structural (uncertainty) shocks.

Figure 20: The Bloom (2009) Dates

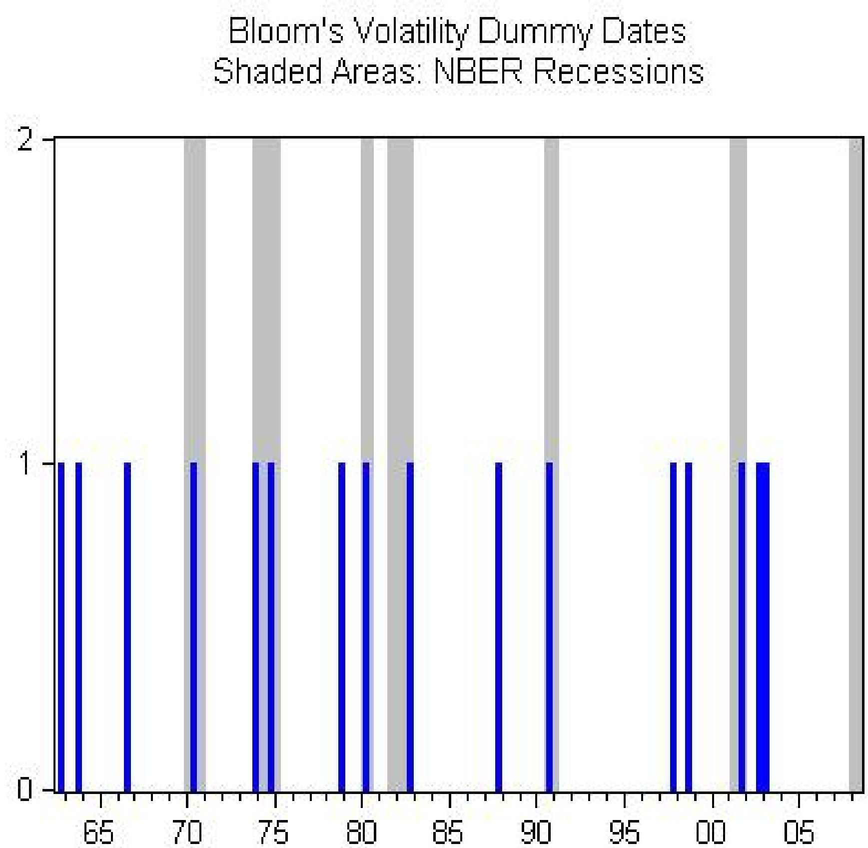

Notes: This graph is a summary of Figure 1 in Bloom (2009). It shows the high stock market volatility dates.

[36]We find similar IRFs when we use employment as opposed to industrial production as the activity variable.

Figure 21: The Bloom (2009) SVAR Conditional on an NBER-Recession

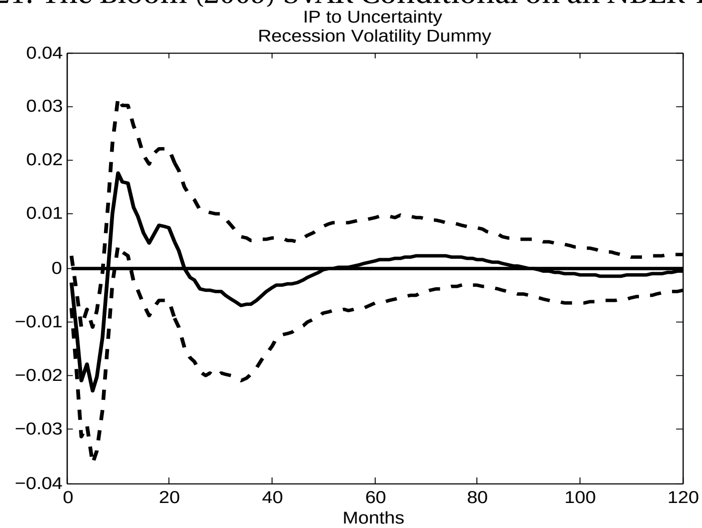

Notes: see notes to Figure 19. This is a replication of Figure 19 with the inclusion of an additional dummy for an NBER recession. The IRF is a response of industrial production to an innovation to stock market volatility, conditional on being in an NBER recession.

Figure 22: The Bloom (2009) SVAR Conditional on not Being in an NBER-Recession

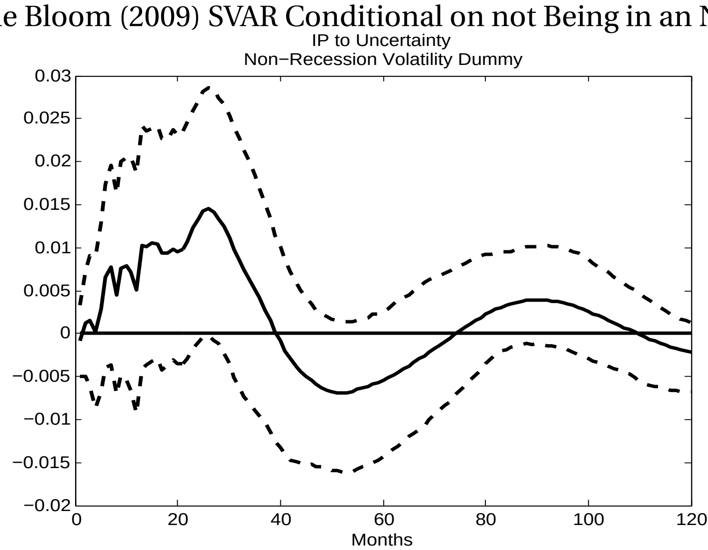

Notes: see notes to Figure 19. This is a replication of Figure 19 with the inclusion of an additional dummy for an NBER recession. The IRF is a response of industrial production to an innovation to stock market volatility, conditional on NOT being in an NBER recession.

B Appendix B - Additional Questions and Variables

Additional BOS Questions and Variables

Q 7 *"Company Business Indicators: Shipments six months from now vs. [CURRENT MONTH]: decrease, no change, increase?"*

Q 8 *"Company Business Indicators: Number of Employees six months from now vs. [CURRENT MONTH]: decrease, no change, increase?"*

Q 9 *"Company Business Indicators: Average Employee Workweek six months from now vs. [CURRENT MONTH]: decrease, no change, increase?"*

Q 10 *"Company Business Indicators: Shipments [LAST MONTH] vs. [CURRENT MONTH]: decrease, no change, increase?"*

Q 11 *"Company Business Indicators: Number of Employees [LAST MONTH] vs. [CURRENT MONTH]: decrease, no change, increase?"*

Q 12 *"Company Business Indicators: Average Employee Workweek [LAST MONTH] vs. [CURRENT MONTH]: decrease, no change, increase?"*

For the question on actual employment changes we also construct a turnover index, defined as:

$Turnover_t \equiv Frac_t(\text{Increase}) + Frac_t(\text{Decrease})$.

Additional SBETS Questions and Variables

Q 13 *"During the **last calendar quarter**, was your dollar sales volume higher, lower, or about the same as it was for the quarter before?* $\boxed{1}$ *Much higher* $\boxed{2}$ *Higher* $\boxed{3}$ *About the same,* $\boxed{4}$ *Lower* $\boxed{5}$ *Much lower. "*

Q 14 *"During the **last three months**, did the **total** number of employees in your firm increase, decrease or stay about the same?* $\boxed{1}$ *Increased* $\boxed{2}$ *Decreased* $\boxed{3}$ *Stayed the same. "*

As with BOS we construct a turnover index for employment from the actual employment change question.

Original IFO-BCS Questions and Variables

Q 15 *“Erwartungen für die nächsten 3 Monate: Unsere inländische Produktionstätigkeit – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY wird voraussichtlich: steigen, etwa gleich bleiben, abnehmen.”*

Q 16 *“Tendenzen im vorangegangenen Monat: Unsere inländische Produktionstätigkeit – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY ist: gestiegen, etwa gleich geblieben, gesunken.”*

C Appendix C - Additional Graphs

Figure 23: BOS - Uncertainty Innovations on Manufacturing Production - Reverse Ordering

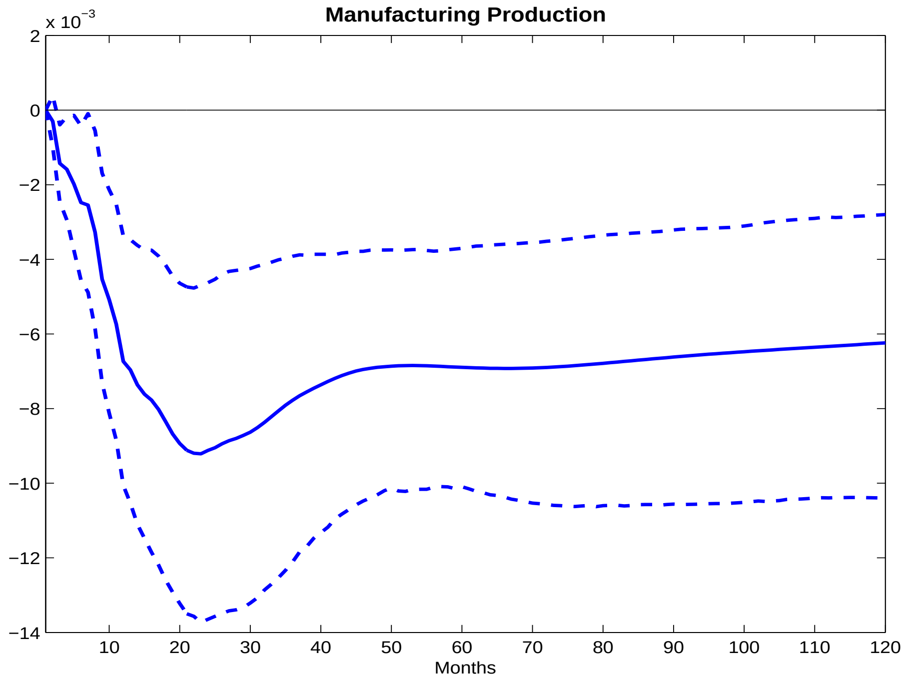

Notes: The IRF is based on a two-variable SVAR with $Uncertainty$ (based on Question 1 of the BOS) ordered second and 12 lags. Manufacturing production is the natural logarithm of the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators. All confidence bands are at the 95% significance level using Kilian's (1998) bias-corrected bootstrap.

Figure 24: Uncertainty Innovations on Various BOS Activity Indices

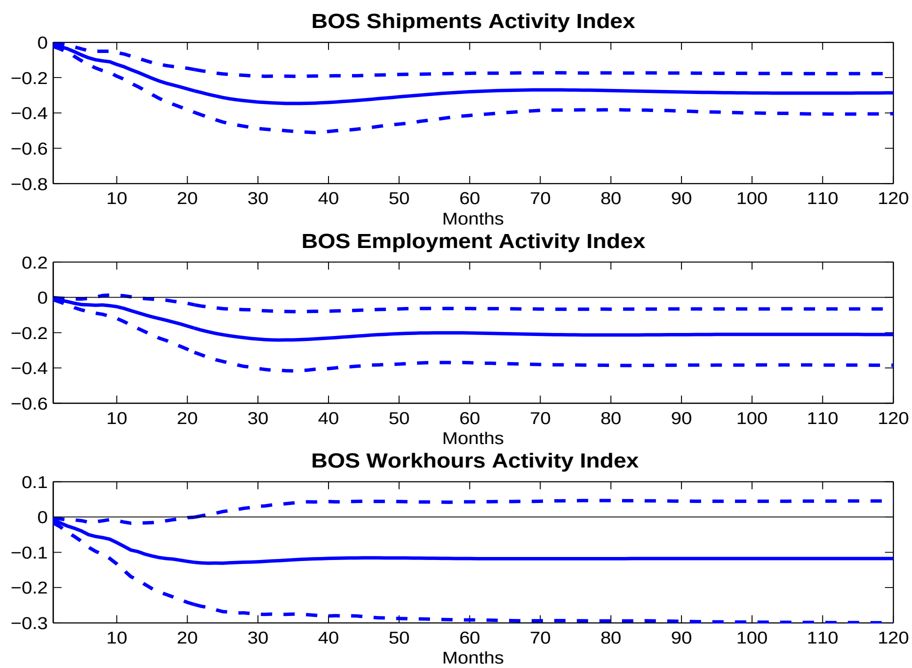

Notes: see notes to Figure 23. The uncertainty variables for the three panels are based on Questions 7, 8 and 9, respectively. The activity indices for the three panels are based on Questions 10, 11 and 12, respectively. Uncertainty is ordered first.

Figure 25: BOS - Uncertainty Innovations on Manufacturing Production - Entropy

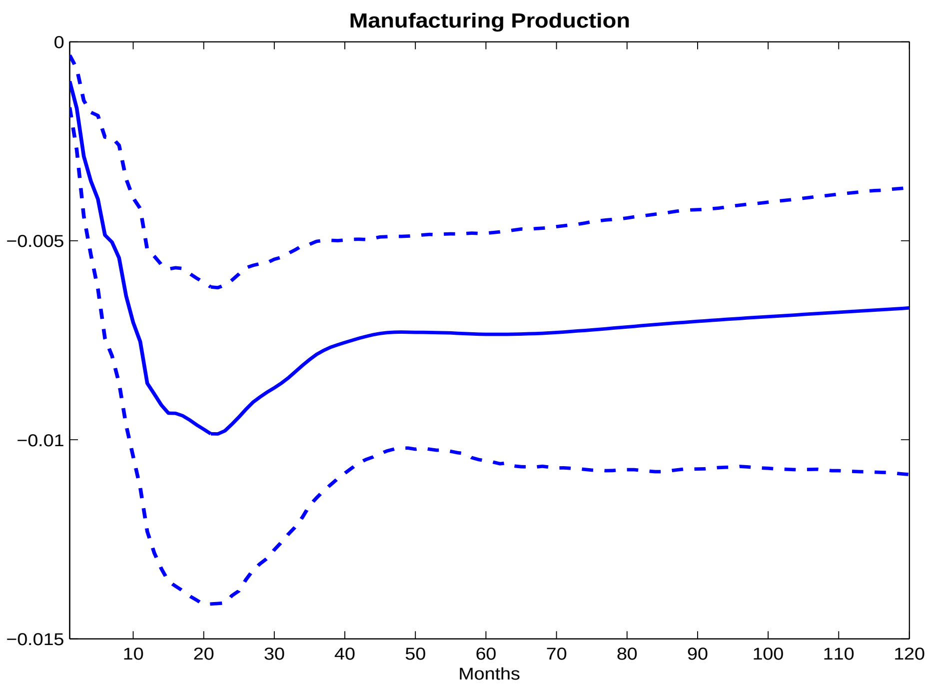

Notes: see notes to Figure 23. Uncertainty is ordered first. It is measured as $Uncertainty_t^{Entrop} \equiv Frac_t(\text{Increase})\log(1/Frac_t(\text{Increase})) + Frac_t(\text{Decrease})\log(1/Frac_t(\text{Decrease})) + Frac_t(\text{Neutral})\log(1/Frac_t(\text{Neutral}))$.

Figure 26: BOS - Uncertainty Innovations on Manufacturing Labor Productivity

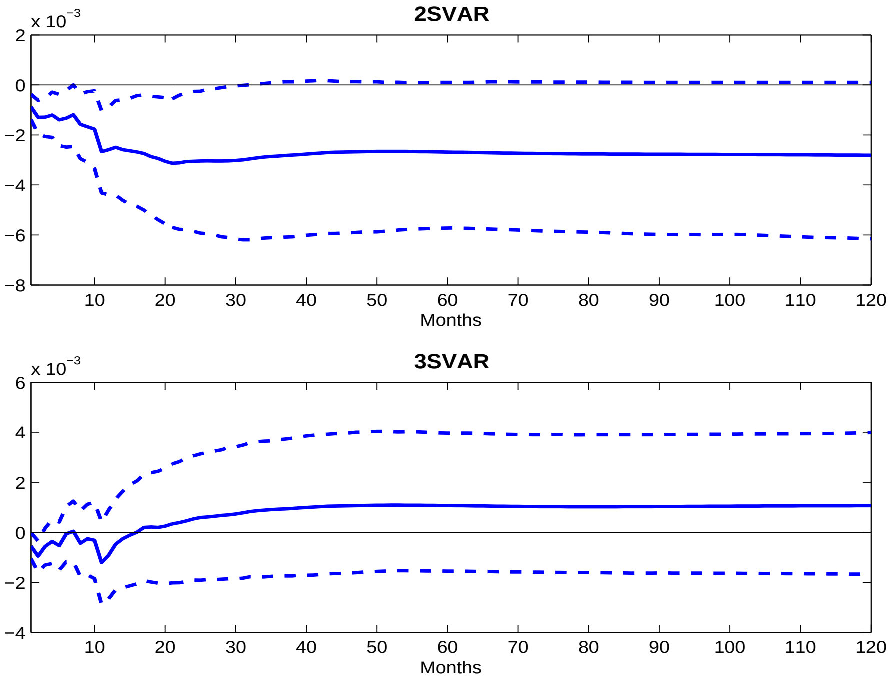

Notes: see notes to Figure 23. In the upper panel uncertainty is ordered first. In the lower panel confidence is ordered before uncertainty and then labor productivity. Labor productivity is the log-difference between the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators and the (seasonally adjusted) monthly manufacturing total hours series, which is itself based on the manufacturing employment and weekly hours series from the BLS-CES data base.

Figure 27: BOS - Uncertainty Innovations Orthogonalized to the Long Run Component - General Activity

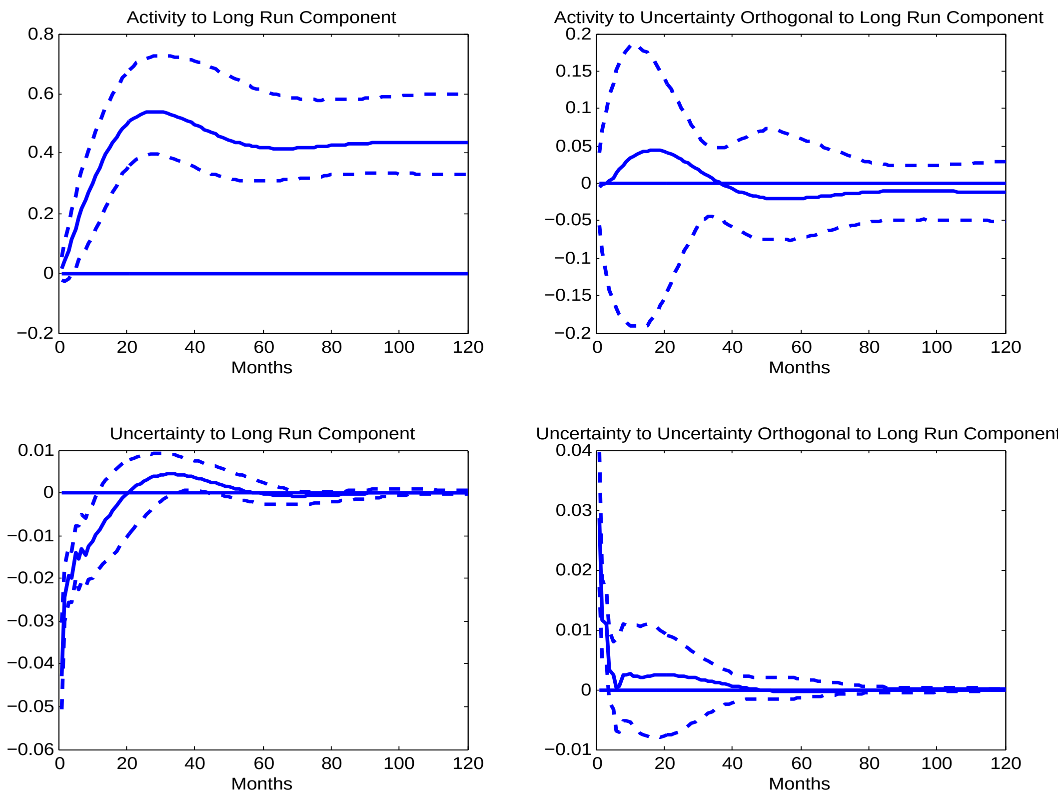

Notes: see notes to Figure 11. The reduced form VARs are estimated with three variables – confidence, uncertainty, and economic activity. Confidence and uncertainty are based on Question 1. The activity series is based on Question 2. The unconditional correlation between the long-run component and innovations to uncertainty is -0.83.

Figure 28: Uncertainty Innovations on IFO General Activity Index Orthogonalized to the Long Run Component

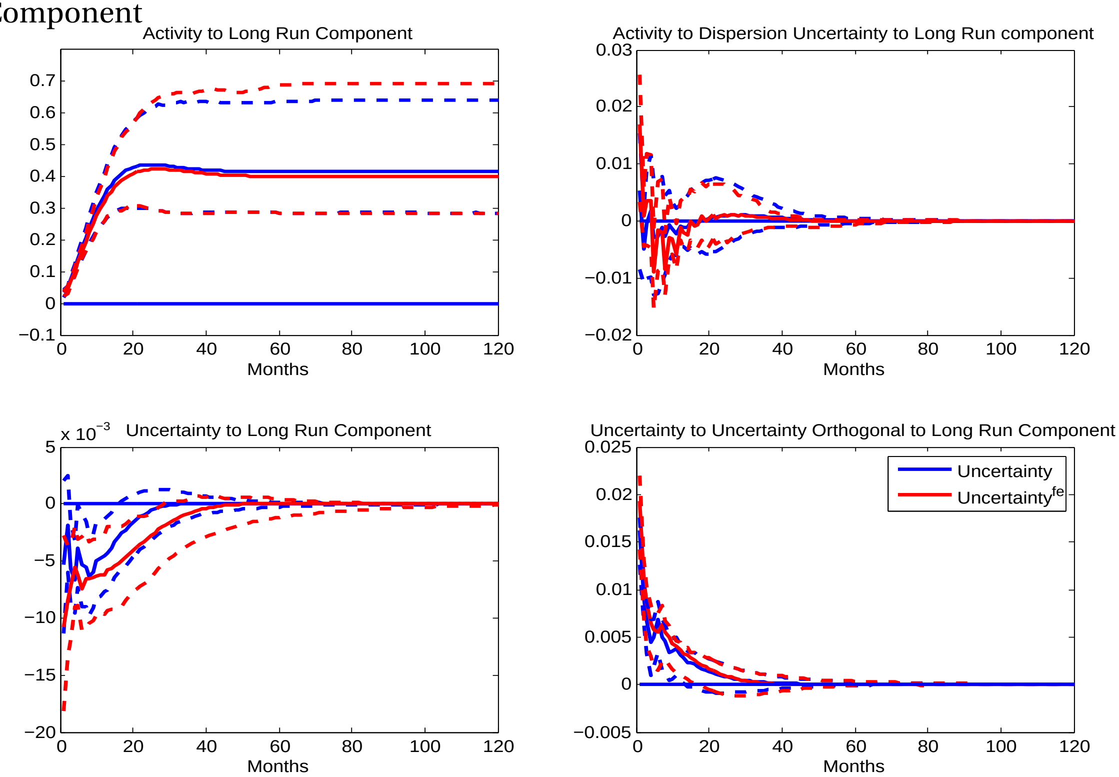

Notes: see notes to Figure 18. The unconditional correlation between the long-run component and innovations to uncertainty is -0.30 for $Uncertainty_t$ and -0.48 for $Uncertainty_t^{fe}$.