

Uncertainty and Change: Survey Evidence of Firms' Subjective Beliefs

Rüdiger Bachmann Kai Carstensen Stefan Lautenbacher
Martin Schneider*

July 5, 2018

Abstract

This paper provides survey evidence on firms' subjective uncertainty about future sales growth from a new representative panel data set of the German manufacturing sector. The main finding is that *uncertainty reflects change*: firms report more subjective uncertainty after either high or low growth realizations. In the cross section of firms, subjective uncertainty differs from statistical measures of uncertainty such as volatility: fast-growing and large firms report lower subjective uncertainty than fast-shrinking and small firms, respectively, even if they face shocks of similar size. In contrast, the substantial time variation in firm-specific subjective uncertainty resembles that in conditional volatility: both measures of uncertainty are mildly persistent and rise more when growth is temporarily low.

Keywords: expectation formation, firms, measurement, subjective uncertainty, survey data

JEL codes: C83, D22, E23

*Respectively: University of Notre Dame, CEPR, CESifo and ifo, e-mail: rbachman@nd.edu (corresponding author); University of Kiel, CESifo and ifo, e-mail: carstensen@stat-econ.uni-kiel.de; ifo and LMU Munich, e-mail: lautenbacher@ifo.de; Stanford University, CEPR and NBER, e-mail: schneidr@stanford.edu. We thank Pete Klenow, Giuseppe Moscarini, Matthew Shapiro, Stephen Terry, Michael Woodford as well as conference participants at the 2017 SED annual meeting in Edinburgh, the 2018 ASSA meetings, the 2018 Cowles Summer Conference on Macroeconomics, the 2018 CESifo Venice Summer Institute, and seminar participants in Kiel, Munich, and Nuremberg for helpful comments, Simeon Häfele for computational assistance, and the team of ifo's Economics and Business Data Center for excellent support and provision of the data. Financial support by the Thyssen Foundation is gratefully acknowledged.

1 Introduction

Businesses deal with uncertainty every day. Any intertemporal decision, be it investment, hiring, or pricing, must be made without knowledge of future shocks to cost and demand. While many macroeconomic models simplify by assuming that firms' decisions depend only on forecasts (that is, expectations of shocks), a recent literature has argued that uncertainty matters for firm actions, see for example the survey by [Bloom \(2014\)](#). A common theme is that uncertainty leads to more conservative hiring or investment decisions. This mechanism is attractive for modeling recessions as times of high uncertainty coupled with low economic activity.

Unfortunately, to date there is little direct evidence on the uncertainty perceived by decision makers within businesses. Empirical work with micro data on firms tends to rely on proxy measures such as stock return volatility or the dispersion of analyst forecasts. As a result, we do not really know which firms are uncertain or how firms' perceived uncertainty evolves with experience. Nevertheless, answers to both questions are crucial to evaluating the new models of uncertainty and its role for the business cycle. In particular, it is still common practice to assume rational expectations: subjective uncertainty is equated with uncertainty measured by an econometrician. If this assumption is not satisfied in the data, we need new models that explicitly incorporate subjective expectation formation.

This paper provides evidence on firms' perceived uncertainty about future sales growth. We exploit a new survey module in the ifo Business Tendency Survey to develop a representative panel data set of firms' subjective beliefs, characteristics and performance for the German manufacturing sector. We then document how subjective uncertainty varies in the time series as well as in the cross section of firms. Our main result is that *uncertainty reflects change*: both at high frequencies and over the medium term, realized change, either good or bad, makes firms more uncertain. In the cross section of firms, subjective uncertainty from our survey differs significantly from volatility measured by an econometrician. At the same time, there is substantial time series variation in firm-specific uncertainty, which is remarkably similar to time variation in conditional volatility experienced by firms.

The new survey module asks firms for a forecast of one quarter-ahead sales growth together with two numbers for their best and worst case scenarios. We then define the *span* between the best and worst case one quarter ahead sales growth scenarios as our quantitative measure of subjective uncertainty. Our data set is based on 14 consecutive quarterly survey waves from 2013 to 2016, with more than 300 firms per wave. It is special in at least three ways. First, the panel dimension allows us to study firms' responses to past growth experience, as well as to their own past forecast errors. Second, attachment to the ifo survey has led to participation by large firms: the 75th percentile of the size distribution is at around 250 employees. Finally, it covers a rather quiet phase of the German business cycle, so that essentially all variation is firm-specific.

Our finding that uncertainty reflects change can be broken down into a series of stylized facts. In the cross section of firms, we show that span comoves with two measures of ongoing change in a firm's environment. First, subjective uncertainty is higher for firms that face more turbulence, as captured by their in-sample volatility. This result is perhaps not surprising; it is consistent with any mechanism that relates realized volatility to subjective uncertainty. Second, subjective uncertainty is higher for firms on strong good or bad growth trends, as compared to firms that grow at moderate rates. This result is

consistent with models of learning in which fast expansion or shrinkage leads firms to a less familiar, and hence more uncertain, state of business. In contrast, it is not consistent with simple models of firm dynamics in which every firm knows its trend growth.

Our third stylized fact is that firms become more uncertain after temporary change. Indeed, the time series relationship between span and the most recently observed sales growth rate is V-shaped: firms become more uncertain after experiencing either very good or very bad sales growth. This result stands in contrast to the negative comovement between growth and uncertainty emphasized in the literature. A key distinction here is that most firm level variation in our sample is idiosyncratic: our results suggest that individual firms' uncertainty is shaped by its individual performance, and increases when an unfamiliar event occurs. This pattern is again consistent with models of learning.

Fourth, we show that the relationship between uncertainty and change is asymmetric: bad change entails more uncertainty than good change. This pattern appears both in the cross section and the time series. On the one hand, firms on a bad growth trend exhibit more average uncertainty than firms on a good growth trend. On the other hand, the V-shaped relationship between span and past sales growth is also asymmetric: while a one percentage point lower negative growth rate is followed by 30 basis points wider span between firms' best and worst case scenarios, a one percentage point higher positive growth rate widens span by only 17 basis points.

How does firms' subjective uncertainty compare to statistical measures of uncertainty such as the volatility of shocks measured by an econometrician? In the cross section of firms, we find two important differences. First, the asymmetric relationship between trend growth and uncertainty is a purely subjective phenomenon. Indeed, statistical measures of firm level uncertainty such as the volatility of growth and firms' average absolute forecast error are also higher for firms on good or bad growth trends. In contrast to the behavior of subjective uncertainty, however, the effect is symmetric: statistical uncertainty is equally high for fast-growing as for fast-shrinking firms. While firms in unfamiliar territory thus indeed experience larger shocks, their perception is less acute if the territory is more favorable.

A second key difference between subjective and statistical uncertainty is that large firms perceive lower subjective uncertainty, holding fixed the volatility of shocks as well as trend growth rates. In particular, the average span reported by very small firms is 5 percentage points wider than that for large firms, for the same trend and turbulence. At the same time, statistical measures of uncertainty do not depend on size: both the volatility of growth as well as the mean absolute forecast error are entirely accounted for by trend and turbulence. This independent effect of size on the perception of uncertainty suggests that simple models of firms with homogenous beliefs and identical objective functions cannot adequately capture firms' behavior under uncertainty.

In the time series dimension, we document large time variation in firm-specific subjective uncertainty. Indeed, the average firm level standard deviation of span is 5.9%. Time varying subjective uncertainty is thus a volatile potential source of changes in firm behavior. In fact, it is almost as volatile as the usual driver of firm behavior in economic models, namely changes in conditional expectations: the average firm level standard deviation of expected growth is 7.1%. Most of the variation in subjective uncertainty is firm-specific: time and sector fixed effects explain only a negligible share.

We also show that firm-specific time variation in uncertainty is remarkably similar to the conditional volatility of shocks experienced by firms: both are mildly persistent, increase with bad past growth and increase somewhat less with good past growth. Indeed,

the panel dimension of our data set allows us to estimate dynamic panel regressions for our subjective uncertainty measure span as well as panel GARCH models for firms' subjective forecast errors and forecast errors from a benchmark statistical model. Measures of uncertainty from all three models share the above properties. They also suggest that firms perceive higher uncertainty not only after large forecast errors, but also after transitory but predictable changes in growth.

To shed further light on mechanisms linking uncertainty and growth, we show that the time series relationship between uncertainty and previous quarter growth does not simply reflect firms' actions *over the previous quarter*. In addition to our uncertainty module, the ifo survey elicits information on firms' capacity utilization, inventory, production and pricing. Controlling for any of those variables does not change our main result. This finding calls into question the role of exogenous "uncertainty shocks" as drivers of firm behavior. It points instead to the role of sales growth as a signal about firms' environment that leads firms to update their views.

To elaborate, compare two mechanisms that could generate correlation between growth in quarter t and uncertainty in the following quarter. First, it could be that persistent increases in uncertainty lead firms to alter their production or inventory policy and thereby affect sales growth. Persistence then implies that higher uncertainty is recorded at the beginning of the following quarter. Under this mechanism, firms learn nothing from observing sales growth – it simply reflects their actions. Alternatively, suppose that sales figures sometimes provide a signal about the environment that increases uncertainty. This does not require any effect of the uncertainty on firms' actions. If the first mechanism were important, including firm actions directly in our regressions should at least partly drive out growth itself. The fact that it does not favors the learning mechanism.

Our paper is related to several strands of literature. First, we contribute to the growing body of work on quantitative survey measures of uncertainty. [Juster \(1966\)](#)'s work on household uncertainty about purchasing decisions first emphasized the relevance of quantitative (as opposed to qualitative categorical) answers to uncertainty-related questions. For households, there is now a fairly large set of surveys that elicit quantitative answers. Examples include studies of the US Health and Retirement Study ([Juster and Suzman, 1995](#); [Hurd and McGarry, 2002](#)), the Bank of Italy's Survey of Household Income and Wealth ([Guiso et al., 1992, 2002](#)), the Survey of Economic Expectations ([Dominitz and Manski, 1997](#)), the Michigan Survey of Consumers ([Dominitz and Manski, 2004](#)) as well as the New York Fed's very recent Survey of Consumer Expectations ([Armantier et al., 2015](#)).

In contrast, only a handful of papers use quantitative survey measures to study uncertainty perceived by firms. [Guiso and Parigi \(1999\)](#) pioneered this line of research using data from the Bank of Italy (see also [Bontempi et al., 2010](#)). Another source of quantitative survey answers are the studies of executives' stock return expectations by [Ben-David et al. \(2013\)](#) and [Gennaioli et al. \(2016\)](#). The Italian studies also point to important cross sectional variation in subjective uncertainty. Our study is unique in its focus on uncertainty and change, both by characterizing the cross section in terms of trend and turbulence and through its focus on time series dynamics in uncertainty and firms' asymmetric updating.

Our findings are relevant for the active literature on time-varying uncertainty in macroeconomics. An empirical literature ([Leahy and Whited, 1996](#); [Guiso and Parigi, 1999](#); [Bachmann et al., 2013](#)) argues that heightened economic uncertainty depresses economic activity. Following [Bloom \(2009\)](#), many studies have explored mechanisms

by which exogenous shocks to uncertainty change firm actions. Examples include wait-and-see effects in the face of higher uncertainty, financial frictions that lead to higher cost of finance, when uncertainty is high (Arellano et al., 2016; Gilchrist et al., 2014), and precautionary saving effects (Fernández-Villaverde et al., 2015; Basu and Bundick, 2017). In these studies, causality usually runs from uncertainty to growth. More recently, several papers have considered feedback effects from growth to uncertainty (Bachmann and Moscarini, 2012; Ilut et al., 2017; Berger and Vavra, 2017; Baley and Blanco, 2017; Ilut and Valchev, 2017). Our results suggest that these feedback effects are particularly important for understanding the comovement of growth and subjective uncertainty.

Finally, our paper is related to a large body of work on deviations from rational expectations. The analysis of miscalibration and overconfidence goes back at least to Oskamp (1965). The earlier literature is summarized by Koriat and Fischhoff (1980); see Morgan (2017) for a more recent treatment. A typical finding is that survey respondents are miscalibrated in the sense that their subjective confidence intervals are narrower than confidence intervals constructed by a statistician from the data. We find a similar result: many growth rate realizations are outside of best and worst case bounds provided by firms. Our results on updating relates to recent work to test Bayesian updating from experimental data (Augenblick and Rabin, 2017).

The paper is structured as follows. Section 2 explains our new survey questions and properties of the data. Section 3 introduces the raw relationship between uncertainty and change and presents a simple organizing framework. Section 4 studies uncertainty and change in the cross section, while Section 5 looks at time series variation after controlling for fixed firm characteristics. Section 6 compares the dynamic properties of subjective uncertainty and statistical measures of uncertainty.

2 Data

The ifo Business Tendency Survey, run by the Munich-based ifo Institute, is a long running survey of German businesses. It provides input for a leading indicator of the German business cycle, the ifo Business Climate Index.¹ We work with the *manufacturing sector supplement* to the main survey which is available for a sample representative of the German manufacturing sector. A firm in the survey is either a stand-alone firm or a division of a large conglomerate. For simplicity, we refer to “firms” throughout this paper. A meta-study performed by ifo (Abberger et al., 2011) suggests that survey questions are filled out by knowledgeable personnel, typically senior management for small and medium-sized firms and executives in controlling for larger corporations.

Since 2013, the manufacturing sector supplement contains an online module of questions on subjective uncertainty. Our current sample comes from 14 survey waves spanning 2013:Q2 through 2016:Q3. The survey is administered at the beginning of every quarter. Participation in the new module is stable at between 300 and 400 firms per wave. The selection of participants is similar to that of the manufacturing supplement as a whole. Questions in the manufacturing supplement that ask about realized outcome (such as

¹The ifo Business Tendency Survey is part of the EU-harmonized business surveys commissioned by the Directorate General for Economic and Financial Affairs of the European Commission. Aggregate survey results for Germany are presented at www.ifo.de/w/3fvxPxj2P, the harmonized European results, including the European Economic Sentiment Indicator, can be found here: https://ec.europa.eu/info/business-economy-euro/indicators-statistics/economic-databases/business-and-consumer-surveys_en

production) explicitly ask firms to ignore seasonal fluctuations. Consistent with this, we observe only negligible seasonal effects in our data; we thus treat the variables below as seasonally adjusted at the individual firm level.

2.1 Eliciting subjective uncertainty

At the beginning of every quarter, firms are asked a three-part question. As an example, Figure 1 displays a sample questionnaire for April 2014 in the original German. The data for this paper come from responses to the following questions. In English, firms are asked:

1. *By how much (in percentage terms) have your sales changed in the first quarter of 2014?*
2. *By how much will your sales change in the second quarter of 2014?*
 - a. in the best possible case (in percent):
in the worst possible case (in percent):
 - b. *Taking into account all contingencies and risks, I expect for the second quarter of 2014 an overall change (in percent) of:*

The questionnaire form contains four boxes for respondents to provide their four numerical answers. Next to every box, there is a reminder to provide positive or negative integers. The default option is to skip the question by checking “don’t know” (“weiss nicht” in German) behind the box, as shown in the empty form in the figure. Once a respondent enters a number, the “don’t know” option becomes unchecked. Finally, underneath both questions 1 and 2, firms are invited to provide free text comments (“Anmerkungen”).

Figure 1: Original survey questionnaire in German

April 2014

Hinweis zu diesen Zusatzfragen:

Das Wirtschaft zu 50% aus Psychologie besteht, wusste schon Ludwig Erhard. Ein wichtiges Element sind dabei Erwartungen über eine unsichere Zukunft, mit der Sie als Unternehmer tagtäglich umgehen müssen. Das haben die Wirtschaftswissenschaften zu lange vernachlässigt. Diese Erwartungen und diese Unsicherheit zu messen und zu evaluieren, ist das Ziel der folgenden Fragen. Mit Ihren Antworten helfen Sie uns sehr.

Für Rückfragen steht Ihnen Frau Wieland zur Verfügung; Tel. 089-9224-1247 - E-Mail: wieland@ifo.de

Die folgenden Fragen beziehen sich auf Änderungen gegenüber dem Vorquartal.

1. Um wieviel Prozent hat sich der Umsatz in Ihrem Bereich im ersten Quartal 2014 verändert?

Veränderung um: % (bitte ganze, positive oder negative Zahlen eingeben) weiß nicht

Anmerkungen:

2. Um wieviel Prozent wird sich der Umsatz in Ihrem Bereich im zweiten Quartal 2014 verändern?

a) Im bestmöglichen Fall: % (bitte ganze, positive oder negative Zahlen eingeben) weiß nicht

Im schlechtestmöglichen Fall: % (bitte ganze, positive oder negative Zahlen eingeben) weiß nicht

b) Unter Berücksichtigung aller Chancen und Risiken erwarte ich im zweiten Quartal 2014 alles in allem eine Veränderung um: % (bitte ganze, positive oder negative Zahlen eingeben) weiß nicht

Anmerkungen:

Note: Original questionnaire from ifo’s online module on subjective uncertainty in German; snapshot from April 2014.

To clarify the timing, consider a firm responding in early April 2014, that is, at the beginning of 2014:Q2. Question 1 asks for the change in sales between 2013:Q4 and 2014:Q1. This is the most recent sales growth realization that the firm has experienced. Question 2 then asks for the firm’s outlook over the current quarter 2014:Q2, as compared to the last quarter 2014:Q1. This is the next growth rate realization that the firm expects.

Our preferred quantitative measure of subjective uncertainty is the *span* between the best and worst case scenarios for sales growth that firms provide in response to question 2.a. Asking for best and worst case scenarios fits well with the prevalence of scenario analysis in German business practice.² For firms that routinely compute adverse and favorable scenarios as part of their planning process, filling out the survey does not impose an additional forecasting task and is likely to generate more thought out answers.

A popular alternative approach to measuring subjective uncertainty starts from a distributional assumption on respondents’ subjective beliefs. Indeed, consider a class of distributions such that its parameters can be identified from the support and a measure of central tendency such as the mean or mode. An example is the triangular distribution: its density looks like a triangle parameterized by upper and lower limits as well as a mode. By inferring upper and lower limit from the bounds in question 2.a and the mode from the forecast in question 2.b, we can construct a subjective distribution for each survey response. We can then compute the variance or interquartile range of that distribution as our measure of subjective uncertainty.

In our data, the alternative approach delivers measures that are almost perfectly correlated with span. This is true for a variety of distributions imposed and statistics used to capture uncertainty. Intuitively, it makes sense that for any class of distributions where the above inference “works”, uncertainty moves closely with the distance between the maximal and minimal support point. One might still expect imperfect correlation between span and the alternative measure if the location of the mean within the support varies a lot across responses. Since this is not the case in our data, many uncertainty measures behave similarly, see Appendix C.

Working directly with span has the advantage that it is a meaningful statistic whether or not the respondents think in terms of probabilities. On the one hand, a Bayesian respondent who forms probabilistic beliefs can provide the support of his subjective distribution. On the other hand, consider a respondent who views uncertainty as Knightian, that is, he cannot assign unique probabilities to every scenario; instead he plans to be robust to multiple possible probability distributions. While a Knightian decision maker cannot accurately express his subjective uncertainty by a probabilistic measure such as, say, variance, he can readily provide best and worst cases.

2.2 Other variables and sample selection

We use responses to the uncertainty question to define a number of other variables. A firm’s *forecast error* is the difference between its actual sales growth in the current quarter and its expected growth rate at the beginning of that quarter, that is, its answer to part 2.b. At the beginning of every quarter, firms cannot perfectly predict the flow of sales over the entire quarter; the forecast errors captures the mistakes they make.

We further define two variables that measure the medium term dynamics in a firm’s environment, based on its realized growth rates (that is, answers to part 1). We refer to a

²Mietzner (2009) provides an overview of the literature on strategic planning in German firms. In many industries, the majority of firms engage in some sort of scenario analysis.

firm’s sample average sales growth as its *trend*. Moreover, the *turbulence* experienced by a firm is measured by the sample standard deviation of its growth rates. While turbulence is a measure of uncertainty, we emphasize that it differs from span for two reasons: First, it is purely statistical as it is based on realized growth rates. Second, it is an average over three years, whereas span measures conditional uncertainty one quarter ahead.

We essentially clean the data in two steps, see also Appendix A for more details. We first exclude outliers by dropping responses with a lower bound below -100% or an upper bound above 100% in parts 1 and 2, with the exception that we set the upper bound to 300% for the best case growth rate in part 2.a.³ We thereby lose 36 firm quarters. In a second step, we check whether respondents consistently order the numbers in question 2 as worst case < forecast < best case. We drop 234 firm quarters for which the forecast is above or below both bounds since it is unclear what the respondent intended. There are also 80 firm quarters for which the forecast is in between the bounds but the worst case is higher than the best case. We attribute this case to a simple mistake when filling out the questionnaire and invert the order of the two bounds.

The baseline sample for our analysis comprises 400 firms for which we have at least 5 observations. After cleaning, we have 2762 firm quarter observations from 14 quarters. For many of our results below, we check for robustness with the full sample of all pooled firm quarters. For some exercises, we need consecutive forecast errors: in this case we work with a restricted sample of 1,329 firm quarters of which 380 firm quarters are used as lagged “pre-sample” observations so that an effective sample of 949 firm quarters remains. The results are similar across all three samples.

We measure firm size by the number of employees. Importantly, our data set contains a substantial number of large firms: the 75th percentile is at about 250 employees. The median firm employs 100 workers while the 25th percentile is at 40. We also know the firms’ sector at the two digit manufacturing level. We construct sector dummies by forming 14 super-sectors for which we have a sufficient number of observations. For further details, see Table 13 in Appendix B.

In addition to responses to the new uncertainty question, we use other responses from the Ifo Business Tendency Survey. Two key continuous indicators of the strength of a firm’s business are capacity utilization (in percent) and the state of business (coded between 0 and 100). There are also categorical questions on the state of orders, inventory, production, prices, demand and credit appraisal. Those variables allow us to construct a rich set of control variables that capture a firm’s information set at the time when it makes its forecast and responds about its subjective uncertainty.

2.3 Summary statistics

Table 1 reports summary statistics for the answers to the new uncertainty question. Panel A pools all firm quarter observations and reports mean, volatility and key quantiles for this pooled sample. The numbers here reflect variation both in the time series and in the cross section of firms. For Panel B, we compute, for each individual firm, the time series mean and standard deviations. The panel reports mean, standard deviation and quantiles of the cross sectional distributions of firm level statistics.

³This choice is motivated by text comments that accompany large (two-digit) growth rates. Comments often suggest that large numbers are deliberate responses rather than mistakes: firms give explanations such as “Many projects were moved into this quarter” and “Invoice of a major project”.

Table 1: Summary statistics of survey answers, pooled and by firm

Variable	N	Mean	Std. Dev.	P10	P25	P50	P75	P90
<i>PANEL A: Pooled data</i>								
Sales growth rate in the previous quarter	2762	1.71	14.69	-15	-5	2	10	15
Absolute value of sales growth rate in the previous quarter	2762	9.96	10.92	1	3	7	12	20
Expected sales growth rate for the current quarter	2710	2.22	10.63	-10	0	2	5	10
Worst case sales growth rate for the current quarter	2762	-4.75	11.82	-20	-10	-2	0	5
Best case sales growth rate for the current quarter	2762	7.36	12	0	2	5	10	20
Span between worst and best case forecast	2762	12.11	9.89	3	5	10	15	25
Absolute forecast error	1664	8.69	10.95	0	2	5	10	20
<i>PANEL B: Data by firm</i>								
Mean by Firm: Sales growth rate in the previous quarter	400	1.74	7.87	-6.67	-1.92	1.69	5.3	10.29
Mean by Firm: Expected sales growth rate for the current quarter	399	2.61	6.57	-3.75	-.6	2.25	5	9
Mean by Firm: Worst case sales growth rate for the current quarter	400	-4.48	7.4	-13	-8.31	-3.35	-.15	2.72
Mean by Firm: Best case sales growth rate for the current quarter	400	7.87	8.02	.81	3.41	6.67	10.61	16.67
Mean by Firm: Span between worst and best case forecast	400	12.34	7.35	5.09	7.15	10.57	15.55	22.33
Mean by Firm: Absolute forecast error	389	9.44	9.55	2.25	4	7	11.67	17.57
SD by Firm: Sales growth rate in the previous quarter	397	11.42	9.22	3.44	5.48	8.59	13.7	23.49
SD by Firm: Expected sales growth rate for the current quarter	396	7.36	7.1	1.6	2.89	5.24	9.72	14.92
SD by Firm: Worst case sales growth rate for the current quarter	397	8.09	7.13	2.12	3.18	6.28	10.8	16.33
SD by Firm: Best case sales growth rate for the current quarter	397	8.13	7.78	2.16	3.43	5.87	10.31	15.52
SD by Firm: Span between worst and best case forecast	397	5.85	5.05	2.04	2.87	4.67	7.5	10.61
SD by Firm: Absolute forecast error	338	6.44	6.28	1.41	2.51	4.33	8.08	13.63

German manufacturing firms experience substantial variation in sales growth that is hard to predict. Realized firm sales growth is about 2% per quarter on average, but varies widely: Panel A shows a standard deviation close to 15 percentage points and an IQ range from -5% to almost 10% . Relative to this variation, the distribution of forecasts is compressed, with an IQ range from zero to 5% . The variance of forecasts is about half that of the realizations. Both moments indicate that predicting sales growth is difficult. At the same time, forecasts display little bias on average: the average forecast is only slightly higher than the average realization.

Firms' forecasts take into account the dynamics of realized growth in a nontrivial way. To illustrate, we compare the variation in forecasts to two extreme time series models: a random walk and an iid process. The difference between a firm's forecast and its last realization of growth – that is, the forecast under a random walk model – has a standard deviation of 17.6 percentage points, larger than that of the forecast itself. At the same time, the difference between a firm's forecast and its firm level mean growth rate – that is, an estimate of an iid forecast – has a standard deviation of 11.1 percentage points, again showing substantial deviation from the simple forecasting model.

Best and worst case scenarios and the magnitude of subjective uncertainty

Firms' best and worst case scenarios bracket forecasts almost symmetrically. The average worst and best case bounds are -4.8% and 7.4% , respectively. The midpoint between these bounds is 1.3% and hence less than one percentage point below the average forecast of 2.2% . Individual bounds have slightly higher standard deviations and wider IQ ranges than forecasts. A key difference between the variables is that the distribution of the lower (upper) bound is negatively (positively) skewed.

Our measure of subjective uncertainty is similar in magnitude to firm level volatility. Indeed, the mean span in the pooled sample is 12.1 percentage points. Moreover, Panel B shows the cross sectional mean of firms' time series standard deviation of growth rates at 11.4% . At the same time, the mean absolute value of firms' forecast errors is 9.4% . The difference between best and worst case scenarios reported by the average firm is thus also similar in magnitude to the typical forecast error experienced by a firm.

Subjective uncertainty in the cross section and time series

Subjective uncertainty differs substantially across firm quarters, with a standard deviation of close to 10% and an IQ range from 5% to 15% . The distribution of firm averages in Panel B indicates a sizeable cross sectional component to this variation: the cross sectional standard deviation of firms' mean span is 7.4% . The magnitude of this cross sectional variation in subjective uncertainty resembles that of cross sectional variation in actual volatility: the cross sectional standard deviation of firms' time series standard deviation of growth is 9.2% .

Our data also show substantial time variation in subjective uncertainty. The sample standard deviation of span for the average firm in Panel B is 5.9 percentage points and hence more than half of the standard deviation of span in the pooled sample. Time series variation in subjective uncertainty is also large compared to other changes in firms' beliefs. For example, the cross sectional mean of firms' time series standard deviation of forecasts is 7.2% and numbers for best and worst scenarios are only slightly higher. A firm's span usually changes together with its forecast: only 13% of changes in span are not accompanied by a forecast revision.

What do *changes* in uncertainty look like? On average, they consist of moves in both the best and worst case scenarios. In particular, for all instances where a firm increases its span from one quarter to the next, the mean change in the worst case scenario is -4.7% whereas the mean change in the best case scenario is $+2.6\%$. In other words, the average increase in uncertainty thus consists of an outward expansion of span that is slightly asymmetric. The average decrease in span is a symmetric downward compression: conditional on a decrease in span, the worst case increases by 3.8% and the best case decreases by 3.4% .

The variation of subjective uncertainty in both time series and cross section is overwhelmingly firm specific. Indeed, the R-squared of a regression of span on time fixed effects is 0.006, on time and industry fixed effects 0.030, and on time-industry fixed effects it is merely 0.084. This fact does not imply that we cannot uncover strong patterns in the variation of span, as we will see below. It simply means that the cross sectional patterns are not driven by the industry, but rather by differences in firm perceptions within industries. Similarly, the time series patterns are driven by individual firm experiences as opposed to, say, the state of the business cycle.

Change in firms' environment

Firms not only differ in subjective uncertainty, but also in the extent to which their environment is subject to continual change. Figure 2 provides a scatter plot of trend and turbulence, defined above as the firm level mean and standard deviation of sales growth, respectively. Every dot represents a firm, and the color of the dot indicates firm size, as measured by the number of employees. Size increases from light blue to pink according to the color bar provided on the right hand side of the figure.

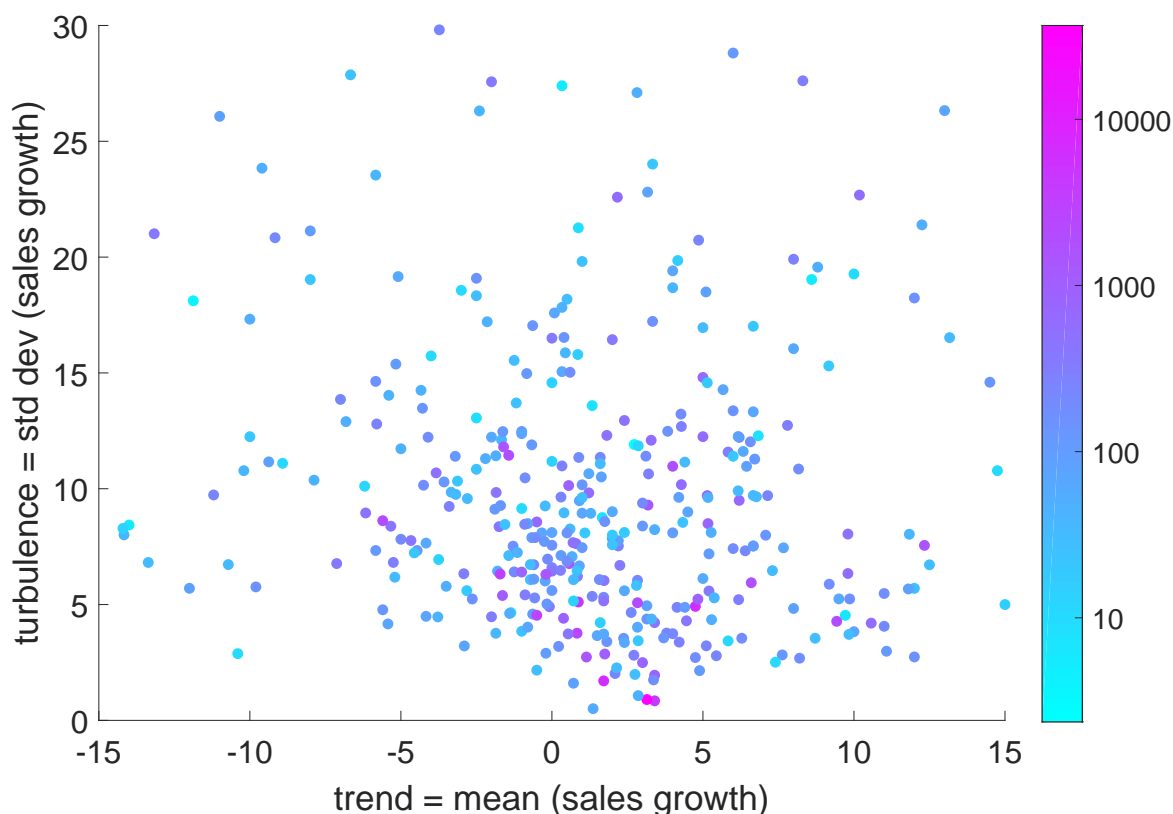
The main takeaway from Figure 2 is that while trend and turbulence vary substantially – as Table 1 also shows –, they are not particularly correlated. Firms that grow or shrink along strong trends need not typically experience large shocks and vice versa. Moreover, the correlation of either environment variable with size is also quite weak. While the very largest firms (identified by bright pink dots) do tend to cluster at the bottom center where trend is small and turbulence is low, we observe firms of all sizes spread out over the plane.

Interpretation of scenarios

The summary statistics reported above suggest that firms' reported best and worst case scenarios are typically not meant to be support bounds. The most direct way to see this is to count how often a realization falls outside the forecast interval between a firm's worst and best case scenario. Table 2 reports this share for different categories of firms. For the overall sample, the first row shows that close to one half of the realizations fall outside the bounds, with about equal shares above the upper and below the lower bound. Sorting firms into size, trend and turbulence bins shows that this patterns is broadly similar across categories.

The finding that the range of scenarios contemplated by firms is narrow compared to the distribution of shocks experienced by firms is familiar from other studies of firms' subjective beliefs. In particular, it also emerges from studies that explicitly elicit probabilistic beliefs. A possible interpretation advanced in the literature is that firms are overconfident, that is, they underestimate the risks they face. While we cannot rule out this possibility, it strikes us as implausible that firms would consistently report support

Figure 2: Trend and turbulence



Note: Every dot represents a firm identified by its trend (firm level mean sales growth) and turbulence (standard deviation of mean sales growth). Color indicates number of employees according to color bar on right hand side.

bounds for quarterly sales growth that are so narrow.⁴

Our preferred interpretation of the data is motivated by the practice of scenario analysis: firms report best and worst case scenarios that are relatively likely, as opposed to, say, tail events or support bounds. From this perspective, it makes sense that average span is of the same order of magnitude as average firm level volatility. However, since we cannot be sure exactly how firms select scenarios, we refrain from judgements about overconfidence. Instead, we focus in this paper on *variation* in subjective uncertainty in cross section and time series, as well as how that variation compares to variation in conditional volatility.

⁴It is important in this context that our observed realized growth rate is already seasonally adjusted by the firm. The presence of seasonals in actual sales growth can therefore not explain the result.

Table 2: Realizations versus the forecast intervals of worst and best case scenarios

	Outside	Below	Above
All observations	0.45	0.20	0.25
<i>By firm size</i>			
Tiny	0.48	0.24	0.24
Small	0.45	0.22	0.23
Medium	0.45	0.21	0.24
Large	0.45	0.18	0.27
<i>By sales growth trend</i>			
'Bad'	0.49	0.29	0.20
'Normal'	0.43	0.20	0.23
'Good'	0.47	0.12	0.35
<i>By turbulence of sales growth rates</i>			
Low	0.33	0.11	0.22
Medium low	0.48	0.23	0.24
Medium high	0.46	0.21	0.25
High	0.55	0.27	0.28

Notes: The numbers in the table represent the shares of realized sales growth rates, for different groups of firms, that fell outside, below, or above the forecast interval of worst and best case sales growth rates.

3 Uncertainty and change

In this section, we first use a simple scatter plot to introduce the main theme of the paper: uncertainty reflects change. We then lay out a simple organizing framework that relates a firm's subjective belief about their sales growth to the distribution of growth measured by an econometrician. In later sections we use the framework to guide our detailed discussion of uncertainty and change in both cross section and time series.


3.1 Uncertainty and past growth: an asymmetric V

Figure 3 displays a scatter plot of responses, with span at the beginning of quarter t along the vertical axis and sales growth realized between quarters $t - 2$ and $t - 1$ along the horizontal axis. Since firms provide integer responses, dots lie on a grid. The vertical gray lines indicate the interdecile range which reaches from -15% to $+15\%$ as reported in Table 1.

Firms that have experienced larger changes are more uncertain. In particular, the relationship between subjective uncertainty and past sales growth looks like the letter V with a minimum at zero. This is illustrated in the figure by two lines. The solid line is a nonparametric regression line. The dashed line is from a simple piecewise linear regression with a breakpoint at zero. The two lines are very similar, and they virtually coincide in the relevant range where most observations are located.

Firms perceive higher uncertainty after negative change than after positive change. Indeed, the slope of the left hand branch of the letter V is about twice as large in absolute value as the slope of the right hand branch. After a one percentage point lower negative sales growth, next quarter's span is half a percentage point wider. In contrast, after one

Figure 3: Uncertainty and past sales growth



Note: Every dot represents a firm-time observation. The solid line is the prediction from a non-parametric kernel regression. The dashed line depicts the predicted values from a piecewise linear regression of subjective uncertainty on past sales growth, with a break at zero. The thin vertical lines mark the interdecile range that extends from -15% to 15% , see Tab. 1.

percentage point higher positive sales growth, span is wider by slightly more than one quarter of a percentage point wider. The regression coefficients are reported in column (2) of Table 4, discussed further below.

The V-shaped regression line relates uncertainty to change; it stands in contrast to the simple linear relationship between uncertainty and growth often emphasized in the literature. At the same time, asymmetry implies that uncertainty and growth are in fact negatively correlated. Indeed, a linear regression returns a small but significantly negative coefficient of -0.06 , shown below in column (1) of Table 4. However, ignoring the V shape drastically lowers explanatory power of past sales growth rate from 19% for a piecewise linear regression to 1% for the simple linear regression.

3.2 Organizing framework

To organize facts about uncertainty and change, we assume that firms have probabilistic beliefs. As will become clear, this assumption is not essential for the points we make here, but it allows us to express those points in simple familiar notation.

Let g_{t+1}^i denote firm i 's sales growth from quarter t to quarter $t+1$, that is, the growth rate that firm i forms beliefs about when it answers our survey questions at the beginning of quarter $t+1$. Firm i 's information set at that point in time includes g_t^i , the last observed growth rate from quarter $t-1$ to t . It may also include other signals that

represent news in quarter t , which we collect in a vector z_t^i . We then use the vector s_t^i to represent all information from past growth rates or other signals that is relevant for forecasting the future dynamics of growth.

We represent firm i 's belief about its sales growth by the state space system

$$g_{t+1}^i = f(s_t^i, x^i) + \sigma(s_t^i, x^i) \varepsilon_{t+1}^i \quad (1)$$

$$s_t^i = S(s_{t-1}^i, g_t^i, z_t^i; x^i) \quad (2)$$

where x^i is a vector of fixed firm characteristics and ε_t^i is an error that has mean zero and variance one under the firms' subjective belief. The observation equation allows firm i 's forecast $f(s_t^i, x^i)$ to depend on the state as well as its fixed characteristics. The state is updated every period to incorporate new information in g_t^i and z_t^i according to the function S .

When firm i answers our survey questions at the beginning of quarter $t + 1$, it provides its forecast $f(s_t^i, x^i)$ as well as best and worst case scenarios. We also observe the subsequent realization g_{t+1}^i and hence the firm's subjective forecast error. We further identify span, the difference between firm i 's best and worst case scenarios, with firm i 's subjective conditional volatility $\sigma(s_t^i, x^i)$. This connection is exact if firm i reports quantiles as scenarios and appropriate distributional assumptions are in place. More generally, we expect firm i 's answer to the survey question to reflect some measure of dispersion in its forecast error, so σ serves as a concrete standin.

Examples

The state space system (1)-(2) nests many models used to describe firms' subjective uncertainty in economic models. As a simple example, consider the case of iid growth together with an orthogonal uncertainty shock:

$$g_{t+1}^i = f + \sigma_t^i \varepsilon_{t+1}^i \quad (3)$$

$$\sigma_t^i = S(\sigma_{t-1}^i, z_t^i) \quad (4)$$

Here the only relevant state is stochastic volatility σ_t^i . Rational expectations models with uncertainty shocks often assume that σ_t^i is correlated across firms and high in recessions, which helps generate observed dispersion in firm growth rates in bad times.

The system (1)-(2) also nests many popular learning rules. Examples include Bayesian models where firms track some latent state such as a regime, or constant gains learning where firms recursively estimate parameters of the one step ahead predictive distribution while downweighting past observations. The common denominator of all these setups is that the state vector contains statistics of the empirical distribution that are relevant for predicting the future dynamics of growth. A natural property in many settings is that high growth g_t^i increases the forecast f and that a large absolute value of the forecast error increases subjective uncertainty σ .

Comparing beliefs and the true data generating process

We would like to distinguish firms' subjective uncertainty from actual volatility, as reflected in the size of innovations measured by an econometrician. We thus consider a change of measure from the firm's belief to the "econometrician's belief", that is, the probability measure that characterizes the true data generating process. We assume that under the econometrician's belief the distribution of growth rates has the alternative

representation

$$g_{t+1}^i = f(s_t^i, x^i) + b^i(s_t^i, x^i) + \hat{\sigma}(s_t, x^i) \hat{\varepsilon}_{t+1}^i \quad (5)$$

$$s_t^i = \hat{S}(s_{t-1}^i, g_t^i, z_t^i) \quad (6)$$

where again the error has mean zero and variance one, now under the econometrician’s belief.

The new observation equation allows for two key differences between firms’ belief and the true data generating process. First, firms might have biased forecasts, represented by the function b . Second, the size of the typical innovation $\hat{\sigma}$ might be different from firms’ subjective uncertainty captured by σ . Both differences may vary either in the cross section with firms’ fixed characteristics x^i or over time with the information set captured by s_t^i . In the special case of rational expectations, there is no bias ($b = 0$) and subjective uncertainty mirrors actual volatility, that is, $\sigma = \hat{\sigma}$.

4 Uncertainty and change in the cross section

In this section, we ask what type of firms perceive more subjective uncertainty *on average*. We thus compute, for each firm, its average span, that is, the time series mean of all observations of span for the firm. We then regress average span on a number of firm level characteristics. In terms of the framework of Section 3.2, we thus characterize the dependence of subjective uncertainty σ on fixed characteristics x^i , assuming that time averaging removes the effects of information s_t^i .

Our selection of firm characteristics x_i is guided by the facts on firm growth rates documented in Section 2. In particular, we consider trend and turbulence, defined above as the firm level mean and volatility of sales growth, respectively. It makes sense that firms in a more turbulent environment – that is, firms that experience larger shocks – are always more uncertain. We would thus expect a V-shaped relationship between growth and span from the cross section alone. To capture turbulence, we use dummies for the quartiles of the distribution of firm level standard deviations in Table 1, with the lowest quartile as the baseline.

The second measure of firms’ environment introduced above is trend, that is, the firm level average growth rate. It is possible that firms on a good or bad trend are more uncertain as they are moving into unfamiliar territory. To investigate this possibility, we again form four quartiles. Since the middle two quartiles turn out to be very similar, we introduce dummies only for a low trend (bottom 25%) as well as a high trend (top 25%). In addition to trend and turbulence, we also divide firms up by size: here we use average employment over our sample to divide firms into four categories. We follow the German Statistical Office in their definition of tiny, small, medium sized, and large firms; lower bounds for the latter three groups are at 10, 50, and 250 employees, respectively. We work with three dummies, with tiny firms as the baseline.

4.1 Subjective uncertainty, size, trend and turbulence

Table 3 presents regression results. The first three columns ask how much can be explained by each fixed characteristic – size, trend and turbulence – separately. All three characteristics show a statistically and economically strong association with span. Col-

umn (1) says that smaller firms perceive less uncertainty. Average span in the entire population of firms is about 12 percentage points, and it falls monotonically from 18 percentage points for very small firms (the omitted category) to 10 percentage points for large firms.

Columns (2) and (3) show that cross sectional variation in trend and turbulence – each by itself – is enough to induce a V shaped relationship between growth and uncertainty, as observed in the pooled scatter plot. On the one hand, trend and span are directly related by an asymmetric V: quickly shrinking or growing firms report higher average spans than firms with normal growth, by 2 and 6 percentage points, respectively. On the other hand, more turbulent firms also report monotonically higher spans. Since more turbulent firms’ growth rate realizations fall more into the tails, this effect also generates a V pattern.

Each of the three firm characteristics has independent effects on the average subjective uncertainty of firms. This is established in column (4) where we consider all three in the same regression. The positive turbulence gradient is qualitatively and quantitatively unchanged compared to the results in column (3). For trend, the interaction with other characteristics is more subtle. In particular, once size and turbulence are controlled for, growing firms no longer perceive higher uncertainty. At the same time, the negative branch of the V remains large and statistically significant.

While trend and turbulence are correlated with size, controlling for them does not remove an independent role for size in explaining subjective uncertainty. Indeed, comparing columns (1) and (4), the negative size gradient is quantitatively reduced, but remains in place qualitatively. Column (5) further shows that our three firm characteristics are not simply reflective of industry characteristics: including industry dummies neither changes significantly the R-squared of compared to column (4) nor the coefficient estimates.

Table 3: Regressions of firm-level average subjective uncertainty, firm-level volatility, firm-level average absolute forecast error and firm-level average forecast error on firm characteristic dummies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	POLS aver. span	POLS aver. span	POLS aver. span	POLS aver. span	POLS aver. span	POLS volatility	POLS aver. abs. FE	POLS aver. FE
Dummy small firms	-3.267* (1.923)			-1.610 (1.655)	-1.936 (1.729)	0.232 (1.504)	1.258 (2.256)	4.156 (2.803)
Dummy medium sized firms	-6.402*** (1.794)			-3.711** (1.550)	-4.308*** (1.623)	0.484 (1.427)	-0.165 (2.066)	1.557 (2.508)
Dummy large firms	-8.834*** (1.827)			-5.051*** (1.603)	-5.705*** (1.716)	0.490 (1.399)	-0.570 (2.055)	3.035 (2.485)
Dummy 'bad' sales growth trend		5.940*** (0.936)		3.233*** (0.821)	3.209*** (0.850)	1.727** (0.732)	2.713*** (0.939)	-5.310*** (1.232)
Dummy 'good' sales growth trend		2.177*** (0.800)		0.444 (0.730)	0.190 (0.739)	1.593** (0.636)	2.493** (1.126)	5.359*** (1.364)
Dummy medium low turbulence			2.287*** (0.623)	1.731*** (0.613)	1.751*** (0.664)	3.261*** (0.205)	2.853*** (0.517)	-0.0466 (0.737)
Dummy medium high turbulence			6.028*** (0.725)	5.052*** (0.701)	4.985*** (0.723)	6.819*** (0.250)	5.244*** (0.563)	0.0492 (0.892)
Dummy high turbulence			11.28*** (0.892)	9.625*** (0.865)	9.216*** (0.898)	19.78*** (1.005)	13.31*** (1.393)	0.0579 (1.640)
Intercept	18.16*** (1.731)	10.32*** (0.456)	7.456*** (0.366)	10.61*** (1.567)	10.96*** (1.790)	2.895** (1.393)	2.721 (2.081)	-2.744 (2.498)
Sector dummies					YES			
No. of observations	400	400	400	400	400	400	388	388
No. of firms	400	400	400	400	400	400	388	388
No. of parameters (excl. intercept)	3	2	3	8	21	8	8	8
R-squared	0.10	0.11	0.34	0.41	0.43	0.68	0.35	0.14

Notes: aver. span denotes the time-series average of firm-level span, volatility is the time-series standard deviation of firm-level sales growth, aver. abs. FE denotes the time-series average of the firm-level absolute forecast error, and aver. FE denotes the time-series average of the firm-level forecast error. Sales growth rates are in the interval (-100,100). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm; * p < 0.10, ** p < 0.05, *** p < 0.01. Size dummies are defined based on the firm-level average number of employees.

4.2 Subjective uncertainty, volatility and bias in the cross section

How does firms' perceived uncertainty relate to the size of the fluctuations they experience? Column (6) of Table 3 reports a regression of firms' average volatility on fixed characteristics. Since span reflects perceived uncertainty *conditional* on information at the beginning of the quarter, column (7) adds an analogous regression for the average absolute value of the firm's subjective forecast error, a measure of the size of innovations experienced by the firm. The results in the two columns are quite similar.

Along all three cross sectional dimensions we consider, subjective uncertainty is significantly different from the size of shocks experienced by the typical firm. First, for both volatility and the absolute forecast error, there is no independent effect of size once we control for trend and turbulence. It is true that unconditionally larger firms experience smaller shocks. However, this relationship is entirely explained by their trend and turbulence. We conclude that the additional effect of size on span is a subjective phenomenon: large firms' perceive lower uncertainty even if they face the same size shocks as smaller firms.

A second special feature of subjective uncertainty is its asymmetric dependence on trend. While it is true that growing and shrinking firms also experience larger shocks, this effect is symmetric. For the same shocks, shrinking firms thus perceive higher subjective uncertainty. A final property concerns the turbulence gradient. While more turbulent firms – who experience larger shocks by construction – also perceive higher uncertainty, they are relatively much less uncertain than low turbulence firms.

As discussed in Section 3.2, forecast errors experienced by firms may in part reflect systematic bias in firms' forecasts. Column (8) of Table 3 shows a regression of the mean forecast error on characteristics. For the size and turbulence categories, the coefficients on the dummies are not statistically significant. Consistent with this result, group means (not reported) are zero if firms are sorted into size or turbulence categories.

At the same time, there is strong evidence that firms on trends make biased forecasts. In particular, growing firms make large positive forecast errors, defined above as realized growth less forecast. In other words, growing firms are regularly positively surprised; their forecasts are biased towards zero. Analogously, shrinking firms make large negative forecast errors: again the forecast is biased towards zero – firms do not sufficiently anticipate the trend they are on.

5 Uncertainty and change over time

We have seen in the previous section that the V shaped relationship between growth and uncertainty in Figure 3 in part reflects fixed differences between firms. We now turn to time series variation: we ask how much of a V remains once we control for fixed characteristics. In terms of the organizing framework of Section 3.2, we ask whether variation of span σ with firms' information s_t^i also contributes to the V shape, via the correlation of s_t^i with past growth.

Formally, let $\Delta y_{i,t-1}$ denote firm i 's sales growth rate between quarters $t-2$ and $t-1$. All our specifications take the form

$$span_{it} = \beta^- \Delta y_{i,t-1}^- + \beta^+ \Delta y_{i,t-1}^+ + \gamma' x_i + \epsilon_{it},$$

where $\Delta y_{i,t-1}^- = \Delta y_{i,t-1} I(\Delta y_{i,t-1} < 0)$, $\Delta y_{i,t-1}^+ = \Delta y_{i,t-1} I(\Delta y_{i,t-1} \geq 0)$, $I(\cdot)$ is the indicator function, and x_i is a vector of fixed firm characteristics that do not depend on time.

We include the three characteristics studied in the previous section: trend, turbulence and size. Trend and turbulence are again coded as dummies. As the unit of observation is now a firm quarter pair, we measure the size of the firm as the number of employees at the end of the previous calendar year. We then form three size dummies: small firms have 10-50 employees, medium sized firms 51-250 employees and large firms more than 250 employees. The baseline “tiny” firm has fewer than 10 employees. While size therefore does vary over time, change is so slow such that the size dummies are essentially time-invariant. We observe only 48 jumps from one size category to another in our sample.

5.1 Time variation in subjective uncertainty and growth

Table 4 reports regression results. As a benchmark, we start in columns (1) and (2) with a simple linear regression and a piecewise linear regression with a break at zero, respectively. The two columns provide formal counterparts to the scatter plot Figure 3. The next four columns augment the piecewise linear specification with dummies for fixed characteristics, first adding size, trend and turbulence separately and then in column (6) adding all characteristics together.

The main result from Table 4 is that a strongly significant asymmetric V remains even if we control for fixed characteristics. Indeed, the coefficients on both negative and positive past sales growth are statistically significant in all specifications. Column (6) says that, holding fixed all characteristics, after one percent lower negative sales growth, next quarter’s span is 30 basis points wider. Similarly, a one percent higher positive sales growth is followed by a 18 bp wider span. To put these numbers in perspective, note that the volatility of sales growth for the average firm is about 11 percentage points, whereas the volatility of span for the average firm is about 6. Responses to past growth thus account for a considerable part of time variation in subjective uncertainty.

The impact of firm characteristics is also significant. First, introducing dummies improves the fit of the regression: for example, the R-squared improves from 0.19 in column (2) to 0.29 in column (6). Coefficients on the dummies reproduce the cross sectional effects discussed in the previous section. For example, firms with more than 250 employees are more than 6 percentage points less uncertain on average than very small firms. Firms that experience more than median turbulence are at least 4.5 percentage points more uncertain than those with low turbulence. The impact of trend is asymmetric: firms on a bad trend are more than 2 percentage points more uncertain than those on a normal trend, whereas a good trend has no significant effect on span.

It is natural to conjecture that fixed characteristics other than size, trend and turbulence matter for subjective uncertainty. We thus rerun the regression in column (7) with firm fixed effects. As expected, we find a large increase in R^2 . Remarkably, however, there is virtually no change in the coefficients on past growth. We can thus conclude that size, trend and turbulence dummies exhaustively control for the impact of firm characteristics on the uncertainty-growth relationship.

In column (8), we include time sector dummies. This neither alters our coefficient estimates nor markedly improves the fit of the regression. This finding is consistent with the fact that variation in subjective uncertainty is largely firm-specific. We conclude that our results are neither driven by industry-composition effects, industry-specific or aggregate trends and cycles.

Table 4: Subjective uncertainty and past sales growth


Dependent variable: Span between best and worst case sales growth rate for quarter t	(1) POLS	(2) POLS	(3) POLS	(4) POLS	(5) POLS	(6) POLS	(7) FE	(8) POLS
Sales growth rate in quarter $t - 1$	-0.0598** (0.0259)							
Negative sales growth rate in quarter $t - 1$		-0.498*** (0.0612)	-0.470*** (0.0602)	-0.436*** (0.0626)	-0.351*** (0.0661)	-0.306*** (0.0675)	-0.272*** (0.0786)	-0.304*** (0.0686)
Positive sales growth rate in quarter $t - 1$		0.279*** (0.0329)	0.266*** (0.0317)	0.280*** (0.0327)	0.166*** (0.0335)	0.180*** (0.0314)	0.159*** (0.0319)	0.173*** (0.0303)
Dummy small firms			-4.480* (2.560)			-3.959* (2.178)		-3.508* (1.897)
Dummy medium sized firms			-6.677*** (2.516)			-5.452** (2.141)		-5.157*** (1.972)
Dummy large firms			-7.858*** (2.570)			-6.295*** (2.170)		-5.970*** (2.014)
Dummy 'bad' sales growth trend				3.711*** (0.951)		2.248*** (0.856)		2.300*** (0.858)
Dummy 'good' sales growth trend				0.410 (0.658)		-0.434 (0.645)		-0.618 (0.667)
Dummy medium low turbulence					1.699*** (0.578)	1.388** (0.591)		1.340** (0.649)
Dummy medium high turbulence					5.157*** (0.752)	4.560*** (0.764)		4.572*** (0.772)
Dummy high turbulence					7.702*** (1.035)	6.748*** (0.969)		6.525*** (0.979)
Intercept	12.22*** (0.392)	8.428*** (0.435)	14.76*** (2.532)	7.695*** (0.435)	6.206*** (0.425)	11.37*** (2.154)	10.06*** (0.480)	9.774*** (3.014)
Time-sector dummies								YES
No. of observations	2762	2762	2762	2762	2762	2762	2762	2762
No. of firms	400	400	400	400	400	400	400	400
No. of parameters (excl. intercept)	1	2	5	4	5	10	401	199
R-squared	0.0079	0.19	0.22	0.21	0.26	0.29	0.57	0.34

Note: Sales growth rates are in the interval (-100,100). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Comparing cross section and time series variation in subjective uncertainty

The coefficients on positive and negative growth in column (6) in Table 4 effectively isolate large and asymmetric *time series* responses of firms to past growth. Interestingly, the asymmetric V-shaped response induced by time series responses is remarkably similar to that induced purely by cross sectional heterogeneity. To see this, consider Figure 4, where we show the fitted line from before and several nonparametric regression lines to compare the two forces.

Figure 4: Cross-sectional and time series relationships between uncertainty and past sales growth



Note: Besides a linear fitted line with break at zero that corresponds to column (2) of Table 4 (red line), the chart presents three non-parametric regression lines. Respectively, the nonparametric regression lines are based on the full sample (blue line), the cloud of predicted values of column (6) of Table 4 (yellow line), and the cloud of predicted values from a model with size, trend, turbulence, and time-sector dummies as regressors (green line). The thin vertical lines mark the interdecile range that extends from -15% to 15% , see Table 1.

As a benchmark, the solid blue line is the nonparametric regression line fit to the data, reproduced from Figure 1. The yellow and green lines are fit not to the data, but to clouds of predicted values from two parametric regressions reported in Table 4.⁵ In particular, the dotted red line is the fitted line from the regression model in column (2) of Table 4, already shown in Figure 3; the dashed yellow line is fit through the predicted

⁵The collection of precise values comprise a cloud rather than, say, a line because growth is not the only regressor; there are also fixed firm characteristics as well as the time sector dummies. For example, for two firms that experienced the same sales growth rate in the previous quarter, the model in column (6) of Table 4 will predict different spans depending on the firms' size, trend, and turbulence.

values from the regression model in column (6) of Table 4. Finally, the dashed green line is fit through the predicted values from a model with *only* the three classes of dummies, thus reflecting only cross sectional variation.

The main takeaway from Figure 4 is that all lines lie effectively on top of each other, especially within the interdecile range. In other words, time series and cross sectional variation induces the same V-shape, albeit through very different mechanisms. For the time series response, the V follows directly from the difference in coefficients on positive and negative growth. For the cross section, the effect is more subtle and comes from the comovement of span and volatility documented in Table 3: firms with higher span also see higher absolute values of their growth rates (due to differences in size, trend or turbulence).⁶

5.3 Subjective uncertainty and volatility in the time series

Section 4.2 showed that subjective uncertainty and volatility vary differently across firms. How do subjective and statistical uncertainty compare in the time series dimension? Table 5 compares our baseline regression of span on past growth and fixed characteristics – column (1) here reproduces column (6) of Table 4 – to an analogous regression for the absolute value of the firm’s subjective forecast error, shown in column (3).

Controlling for fixed characteristics, a firm that observes one percent worse negative growth in the previous quarter not only increases its span by 31 basis points, but also experiences, on average, a subjective forecast error that is 34 basis points higher in absolute value. In contrast, one percent higher positive growth increases span by 18 bp and the absolute value of the forecast error by 12 bp. The asymmetric V that emerges in the time series of firms’ uncertainty is thus also present in firms’ experience of shocks.

The differences between subjective uncertainty and volatility observed in the cross section appear to be largely orthogonal to the time series dynamics of span and growth. Indeed, the regression coefficients in columns (1) and (3) of Table 5 display the same patterns as columns (4) and (6) of Table 3: large firms and very turbulent firms are less uncertain than one might expect given the size of the shocks they face, and the asymmetric relationship between trend growth and uncertainty is a subjective phenomenon.

⁶In Appendix D we also show that the V-shaped relationship between sales growth and subjective uncertainty holds separately, and in a quantitatively similar manner, for all firm-level subgroups: the four firm size groups, the four turbulence groups, and the three growth trend groups.

Table 5: Regression of subjective uncertainty and the absolute forecast error on past sales growth, fixed firm characteristics, and additional controls

Dependent variable:	(1) span	(2) span	(3) firms' abs(FE)	(4) firms' abs(FE)
Negative sales growth rate in $t - 1$	-0.306*** (0.0675)	-0.253*** (0.0777)	-0.337*** (0.0689)	-0.341*** (0.0682)
Positive sales growth rate in $t - 1$	0.180*** (0.0314)	0.173*** (0.0322)	0.123** (0.0529)	0.147*** (0.0505)
Dummy small firms	-3.959* (2.178)	-3.734** (1.805)	-0.632 (1.545)	0.445 (1.477)
Dummy medium sized firms	-5.452** (2.141)	-5.796*** (1.909)	-1.672 (1.460)	-0.934 (1.374)
Dummy large firms	-6.295*** (2.170)	-6.655*** (1.924)	-1.810 (1.598)	-1.218 (1.535)
Dummy 'bad' sales growth trend	2.248*** (0.856)	1.668** (0.845)	1.340* (0.811)	0.516 (0.937)
Dummy 'good' sales growth trend	-0.434 (0.645)	-0.418 (0.676)	1.417* (0.826)	1.651* (0.938)
Dummy medium low turbulence	1.388** (0.591)	1.416** (0.656)	1.802*** (0.356)	1.817*** (0.484)
Dummy medium high turbulence	4.560*** (0.764)	4.448*** (0.766)	4.160*** (0.472)	4.048*** (0.498)
Dummy high turbulence	6.748*** (0.969)	5.864*** (0.959)	9.161*** (0.953)	8.240*** (0.977)
Expected sales growth rate for quarter t		0.0751** (0.0322)		0.0152 (0.0513)
State of business (VAS) at the time of the survey		-0.0575*** (0.0182)		0.00269 (0.0244)
Intercept	11.37*** (2.154)	32.70*** (7.961)	3.782** (1.484)	0.520 (4.583)
Other firm controls		YES		YES
Time-sector dummies		YES		YES
No. of observations	2762	2607	1664	1597
No. of firms	400	384	389	374
No. of parameters (excl. intercept)	10	222	10	206
R-squared	0.29	0.38	0.25	0.37

Note: span is our measure of subjective uncertainty, firms' abs(FE) denotes firms' absolute forecast error, and statistical abs(FE) denotes the absolute forecast error generated by a simple statistical forecasting model of sales growth as explained in Section 6.1. All equations are estimated by pooled OLS. Sales growth rates are in the interval (-100,100). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Other firm controls include the capacity utilization rate and dummies for the state of business, orders, stock of inventory, production changes, demand changes, order changes, price changes, credit appraisal, capacity utilization, and constraints to production.

5.4 Controlling for firm actions and expectations

What mechanism is behind the time series relationship between uncertainty and growth? One possible explanation is that uncertainty shocks are reflected in firm actions and hence growth. In particular, it may be that we can simply extend recent models of uncertainty-driven recessions to firm dynamics. The key feature of such models is that the stochastic process driving uncertainty is orthogonal to innovations to growth conditional on uncertainty, as in example (3)-(4) above.

An alternative interpretation of the data invokes learning or signal extraction on the part of firms, another special case of our general model (1)-(2). Suppose firms view growth itself as a signal of future business conditions. Many models of belief updating imply that a more unusual realization should increase uncertainty. For a concrete example, consider a regime switching model: the state of business follows a finite state Markov chain, and growth is given by a regime-dependent mean plus iid noise. A more unusual growth rate observation will typically lead the firm to perceive more uncertainty about its current estimate of the regime.

The key difference between the two candidate mechanism is the role of firm actions for growth. Under the uncertainty shock story, growth responds to persistent changes in uncertainty. Comovement between current uncertainty and past growth is thus induced by firms' past responses to uncertainty. In (3)-(4), there is no additional signal value to growth conditional on volatility last period. Under a learning story, in contrast, there may be *no* effect of current uncertainty on past firm actions.

The different role of firm actions for the two stories suggests a simple check: if the relationship between growth and uncertainty is induced by firms' actions, then controlling for firm actions in our regressions should alter the coefficients we find on past growth. We perform the check in column (2) of Table 5: here we augment the baseline regression in column with (1) with both time sector dummies as well as additional set of controls that vary *over time at the firm level*.

The new controls are firms' expectations about their future sales growth, their assessment of the state of their business and their capacity utilization, as well as categorical variables that report orders, the stock of inventory, demand, production, prices, and credit conditions. The difference between these additional controls and the firm characteristics already included above is that the former vary at high frequency at the firm level. In particular, they capture actions the firm could have taken at the beginning or during the course of the quarter such as scale down production or increasing prices.⁷

The main result from column (2) is that the coefficients on growth are virtually unchanged, as predicted by a learning model. This is not because the new variables are uncorrelated with subjective uncertainty: on the contrary, the R-squared improves by 9 percentage points between columns (1) and (2). About half of this improvement comes from the inclusion of time varying firm level controls – this follows from a comparison with column (9) of Table 4. Nevertheless, the impact of the new controls appears to be orthogonal to the role of past growth in predicting subjective uncertainty. We conclude that the timing of growth, uncertainty and firm actions favors a learning interpretation of our time series facts, and provides little support for an uncertainty shock mechanism.

⁷Appendix F provides a detailed list of the additional firm-level controls and their timing relative to quarter t , when span was uttered.

6 The dynamics of subjective uncertainty and volatility

In this section, we further explore the dynamics of subjective uncertainty. Our approach is motivated by two properties of many common learning rules: we would expect higher forecast errors to increase uncertainty, and changes in uncertainty to propagate over time. In principle, these properties alone could induce a V-shaped relationship between uncertainty and growth. It is thus interesting to ask whether controlling for lagged forecast errors as well as lagged uncertainty affects the relationship between growth and uncertainty. The answer will help guide the choice among alternative models of belief updating.

The second goal of this section is to compare the dynamics of subjective uncertainty perceived by firms to that of “statistical” uncertainty experienced by firms. We have already seen in Section 5.3 that the projection of the size of forecast errors on past growth – controlling for fixed characteristics – yields coefficients that are quite similar to the coefficients of our baseline span regression. We now ask to what extent firms’ updating of subjective uncertainty resembles the dynamic behavior of the conditional volatility of shocks. In terms of the framework of Section 3.2, how similar are the dynamics of the subjective and statistical uncertainty measures σ and $\hat{\sigma}$?

The section proceeds as follows. As a preliminary step, it makes sense to distinguish between the two sources of firm forecast errors discussed in Section 3.2: bias and conditional volatility. In Section 6.1, we thus “clean” subjective forecast errors by removing firm fixed effects of forecast bias already documented in Section 4.2. We also provide an alternative benchmark set of forecast errors from a statistical model. That model is by construction unbiased, although it may also be based on less information than what is available to firms. Its purpose is show whether firm’s subjective forecasts are special, or whether they just reflect generic properties of the data.

We then estimate dynamic models of subjective uncertainty and conditional volatility. For subjective uncertainty, which we directly observe, Section 6.2 runs regressions of span on lagged span, past growth, past cleaned forecast errors as well as fixed firm characteristics. To provide a statistical counterpart, we must first specify a model of conditional volatility. In Section 6.3, we thus estimate such models for both the cleaned subjective errors and for the statistical errors. Section 6.4 finally compares the dynamics of all three measures of firm uncertainty.

6.1 Cleaned subjective and statistical forecast errors

In order to clean observed forecast errors of firm-specific bias, we run regressions of forecast errors on fixed characteristics and use the resulting residuals as our cleaned errors. In terms of the representation (5), this removes the part of the bias $b(s_t^i, x^i)$ that depends on the fixed characteristics x^i . We do this in order to focus on belief dynamics. In particular, we are interested in the response of span to temporary surprises experienced by firms, not surprises they routinely experience because they make biased forecasts.

We regress the forecast error on the size, trend, and turbulence dummies as well as their pairwise interactions. To prevent overfitting, we apply the LASSO estimator (Tibshirani, 1996) to select a subset of relevant regressors.⁸ We choose the LASSO tuning

⁸The LASSO is a standard shrinkage estimator popular in “big data” analysis (Varian, 2014) as it

parameter τ by minimizing Mallows’s C_p statistic as suggested by [Efron et al. \(2004\)](#). The LASSO then selects 11 predictors. In particular, denote the high and low growth dummies by gd_1 and gd_3 , the medium low, medium high and high volatility dummies by vd_1 , vd_2 , and vd_3 , and the small, medium and large size dummies by sd_1 , sd_2 , and sd_3 , respectively. The selected predictors are then gd_1 , gd_3 , as well as the interactions $sd_1 \cdot gd_1$, $sd_2 \cdot gd_1$, $sd_2 \cdot gd_3$, $sd_1 \cdot vd_2$, $sd_1 \cdot vd_3$, $sd_2 \cdot vd_4$, $sd_3 \cdot vd_4$, $gd_1 \cdot vd_4$, $gd_3 \cdot vd_4$.

To obtain a bias-adjusted forecast error we construct the OLS residuals of a regression of the forecast error on these predictors. In doing so we follow [Belloni and Chenrozhukov \(2013\)](#) who show that it is advantageous in terms of bias to let the LASSO select the relevant regressors and subsequently apply OLS to estimate the regression coefficients, see [Lehrer and Xie \(2017\)](#) for a related application. We find that the R-squared of the OLS regression is 0.04 which indicates that biased forecasting is not pervasive but rather an issue for small groups of firms.

The cleaned forecast errors may still reflect bias that varies with the firm’s information set s_t^i . To provide an alternative benchmark, we construct a set of statistical forecast errors by using our own statistical forecasting models. In particular, we regress sales growth on its own lag as well as size, growth, and turbulence dummies. We allow for an asymmetric response to past growth, as in our span regressions. We also note that since trend and turbulence are defined using as sample averages, they are strictly speaking not part of the information set of a firm. At the same time, firms have longer samples than we have that speak to their trend growth and volatility. Our assumption here is that trend and turbulence reflect medium term prospects known to firms.

The regression coefficients are reported in column (1) of Table 6. For robustness, we try a number of specifications. In column (2), we leave out the growth trend dummies from the set of regressors as they might be very dominant. In column (3) and we also estimate a simple asymmetric AR(1) model. In columns (4)-(6), we replicate the first three specifications but restrict the effect of past sales growth to be symmetric. Forecast errors from all six specifications are highly correlated, with correlation coefficients at 0.93 or above. Moreover, both AIC and BIC favor specification (1). In what follows, we report only results based on that specification.

6.2 What moves subjective uncertainty?

Table 7 digs more deeply into firms’ time series response to growth. For the subsample of responses for which we observe a forecast for the previous quarter, we can decompose growth into forecast plus forecast error. We can therefore distinguish firms’ responses to unanticipated change – as captured by their forecast error – and anticipated change already captured by their forecast. Formally, we proceed by including the forecast error as an additional regressor.

Column (1) replicates column (6) of Table 4, our baseline result, now on the somewhat smaller sample for which we observe firm forecast errors: the V-shape of subjective uncertainty in previous-quarter sales growth is again present. Then, in column (2), we replace sales growth with previous-quarter forecast errors in sales growth: we find again

recovers the correct (sparse) model with high probability ([Hastie et al., 2017](#)). By requiring that the L_1 norm of the coefficient vector does not exceed a certain threshold, say, τ , the LASSO restricts many coefficients to zero and thereby helps to balance the bias-variance tradeoff seen in forecasting. This is why the LASSO and related estimators are widely applied for economic forecasting in data-rich environments ([Bai and Ng, 2008](#); [Manzan, 2015](#); [Elliott and Timmermann, 2016](#)).

Table 6: Regression of sales growth (t to t+1) on past sales growth and dummies

	(1) POLS	(2) POLS	(3) POLS	(4) POLS	(5) POLS	(6) POLS
Negative sales growth in quarter $t - 1$	-0.269* (0.141)	-0.141 (0.145)	-0.130 (0.128)			
Positive sales growth in quarter $t - 1$	-0.0357 (0.0907)	0.0392 (0.0978)	0.0495 (0.0856)			
Dummy small firms	5.314*** (1.815)	7.798*** (2.255)		5.298*** (1.747)	7.776*** (2.192)	
Dummy medium sized firms	4.996*** (1.669)	7.408*** (2.136)		5.082*** (1.607)	7.432*** (2.073)	
Dummy large firms	5.000*** (1.689)	7.228*** (2.127)		5.013*** (1.632)	7.159*** (2.057)	
Dummy 'bad' sales growth trend	-7.143*** (0.967)			-6.726*** (0.936)		
Dummy 'good' sales growth trend	8.005*** (1.095)			8.253*** (1.112)		
Dummy medium low turbulence	-1.717*** (0.572)	-2.068** (0.959)		-1.434*** (0.500)	-1.827* (0.929)	
Dummy medium high turbulence	0.427 (0.833)	-0.882 (1.117)		0.973 (0.642)	-0.411 (0.968)	
Dummy high turbulence	-0.143 (1.561)	-0.274 (1.691)		1.343 (1.149)	0.954 (1.384)	
Sales growth in quarter $t - 1$				-0.132*** (0.0414)	-0.0365 (0.0455)	-0.0247 (0.0462)
Intercept	-3.243** (1.610)	-5.059** (2.092)	1.164 (0.811)	-2.881* (1.534)	-4.663** (2.011)	2.015*** (0.439)
No. of obs.	1329	1329	1329	1329	1329	1329
No. of firms	292	292	292	292	292	292
R-squared	0.12	0.021	0.0047	0.12	0.018	0.001
AIC	10696.7	10837.8	10848.1	10702.2	10839.9	10851.5
BIC	10753.8	10884.5	10863.7	10754.1	10881.5	10861.9

Note: Sales growth rates are in the interval (-100,100). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

a V-shape with somewhat smaller coefficients. We note that the R-squared of this regression declines somewhat compared to column (1). Column (3) presents the results from a regression specification where both sales growth and the forecast error, in both cases allowing for asymmetry, are included. It is now the asymmetric V in sales growth that wins the horse race between change and unanticipated change, the latter losing asymmetry and (almost) statistical significance. Notice again that the R-squared in column (3) hardly improves compared to column (1).

Columns (4) to (6) repeat the same steps, but this time with lagged span included in the regression. The result is the same, even starker: in a horse race between sales growth and unanticipated sales growth to determine subjective uncertainty, it is the former that enters with an asymmetric V, whereas the data do not ask for the latter over and above sales growth. Finally, we note that subjective uncertainty displays a mild persistence because it depends also on its own lag, in addition to anticipated sales growth.

Table 7: Regressions of span on past sales growth and past forecast error, with dynamic models

	(1) POLS	(2) POLS	(3) POLS	(4) POLS	(5) POLS	(6) POLS
Subjective uncertainty in $t - 1$				0.274*** (0.0701)	0.271*** (0.0748)	0.276*** (0.0695)
Negative sales growth in $t - 1$	-0.347*** (0.0742)		-0.286*** (0.0750)	-0.329*** (0.0750)		-0.333*** (0.0910)
Positive sales growth in $t - 1$	0.144*** (0.0378)		0.0973** (0.0437)	0.145*** (0.0381)		0.0936* (0.0478)
Negative past forecast error in $t - 1$		-0.240*** (0.0553)	-0.0839* (0.0496)		-0.166*** (0.0586)	0.0167 (0.0607)
Positive past forecast error in $t - 1$		0.128*** (0.0382)	0.0719 (0.0476)		0.116*** (0.0363)	0.0641 (0.0498)
Intercept	11.22*** (2.601)	11.14*** (2.629)	10.97*** (2.588)	8.573*** (2.039)	8.719*** (2.039)	8.465*** (2.008)
Size, trend, and turbulence dummies	YES	YES	YES	YES	YES	YES
No. of observations	1548	1548	1548	1513	1489	1489
No. of firms	375	375	375	372	367	367
No. of parameters (excl. intercept)	10	10	12	11	11	13
R-squared	0.33	0.32	0.34	0.40	0.38	0.41

Note: Sales growth rates are in the interval (-100,100). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Measuring conditional volatility

How does subjective uncertainty compare to conditional volatility, or more generally, uncertainty measured by an econometrician? We have identified subjective uncertainty with span, the difference between best and worst case scenarios. A natural “objective” counterpart is then length of a forecast interval constructed by the econometrician, for example the difference between an upper and lower quantile of the conditional distribution of forecast errors. For distributions of the location-scale family, that length is a multiple of the standard deviation.

We study the conditional volatility of both types of forecast errors defined in Section 6.1. In both cases, we select and estimate a volatility model that optimally describes the data as indicated by information criteria. Since our approach requires successive forecast errors, we work with a restricted sample of 949 firm quarters plus 380 pre-sample observations and thus 1,329 observations altogether.

Let $\tilde{\varepsilon}_{it}$ be the bias-adjusted forecast error of firm i in quarter t , and denote its conditional standard deviation by σ_{it} . Our choice of functional form mirrors our analysis of subjective uncertainty in Section 6.2: we write σ_{it} as a function of past growth, past forecast errors and fixed firm characteristics. We thus use a restricted version of the power GARCH model (Ding et al., 1993; Ding and Granger, 1996; Karanasos and Kim, 2006). While the unrestricted power GARCH model conditions σ_{it}^p on past information, $I_{i,t-1}$, where p is a power coefficient to be estimated, we impose the restriction $p = 1$ and thus model the conditional standard deviation.

Our conditional volatility model then has the general form

$$\tilde{\varepsilon}_{it} = \mu + \varepsilon_{it}, \quad \varepsilon_{it} | I_{i,t-1} \sim N(0, \sigma_{it}^2) \quad (7)$$

with a conditional standard deviation equation

$$\sigma_{it} = \exp(\beta_0 + \beta_1' x_{it}) + \alpha_1 (|\varepsilon_{i,t-1}| + \gamma \varepsilon_{i,t-1}) + \alpha_2 \sigma_{i,t-1}. \quad (8)$$

The mean equation (7) includes an intercept to account for a nonzero sample mean that arises because we apply the bias adjustment of the forecast error to all 1,329 observations but estimate the volatility model on an effective sample of those 949 observations for which a lag is available. The conditional volatility equation (8) allows for an asymmetric effect of the past absolute forecast error measured by the coefficient γ , as asymmetry was found to be relevant in explaining subjective uncertainty. However, we do not expect a strong effect here because the empirical unconditional distribution of the bias-adjusted forecast errors is essentially symmetric with a sample skewness of 0.1.⁹

We add two types of explanatory variables to the conditional volatility equation through an exponential link function which ensures effects are always positive. The first type consists of size, trend and turbulence dummies which are essentially time-invariant and thus control for different levels of conditional volatility for subgroups of firms. Our analysis above indicated that these dummies are sufficient to capture the bulk of time-invariant heterogeneity in subjective uncertainty. The second type includes positive and negative sales growth in the previous quarter which we found to be highly relevant to explain subjective uncertainty.

To find a reliable parsimonious specification, we estimate several restricted versions of (7)-(8) by maximum likelihood. Specification (1) adds no additional control variables ($\beta_1 = 0$), specification (2)-(4) allow only for size, trend and turbulence dummies, respectively, specification (5) allows only positive and negative sales growth rate in the previous quarter, and (6) adds all variables together. All specifications are estimated either assuming symmetric effects of past forecast errors ($\gamma = 0$) or allowing for asymmetry (γ unrestricted).

To select among these specifications, we use two information criteria, AIC and BIC, which are commonly used in applied papers (Nelson, 1991; Zivot, 2009). In finite samples the BIC typically favors overly sparse models, while the AIC picks models with a more generous number of parameters, see Efron et al. (2004, pp. 230-235) for a general discussion and Lütkepohl (2005) for asymptotic and simulation evidence in a time series context. Hence, the models chosen by AIC and BIC may be thought of giving upper and lower bounds in terms of parametrization.

The results reported in Table 8 suggest that the inclusion of turbulence dummies,

⁹A test of the null hypothesis that the population skewness is zero cannot be rejected (p -value of 0.2).

Table 8: Model selection criteria for different specifications of the conditional volatility model

Specification	Symmetric $\gamma = 0$			Asymmetric $\gamma \neq 0$		
	k	AIC	BIC	k	AIC	BIC
(1) no controls	4	7,256.39	7,275.81	5	7,253.13	7,277.41
(2) only size dummies	7	7,249.95	7,283.94	8	7,248.10	7,286.94
(3) only growth trend dummies	6	7,173.30	7,202.43	7	7,174.46	7,208.45
(4) only turbulence dummies	7	6,901.40	6,935.38	8	6,903.38	6,942.22
(5) only sales growth rate	6	7,106.90	7,136.03	7	7,104.63	7,138.62
(6) all controls	14	6,871.59	6,939.57	15	6,870.96	6,943.79

Notes: k denotes the number of parameters. Sales growth rates are in the interval (-100,100). All specifications are estimated by maximum likelihood using 949 observations.

whether by themselves in specification (4) or jointly with the other control variables in specification (6), is essential for model fit: all other specifications generate much larger information criteria. Deciding between specifications (4) and (6) is more tricky. In both the symmetric and the asymmetric case, the AIC favors the inclusion of all controls while the BIC picks the turbulence dummies alone. However, the differences in terms of AIC are large (29.81 and 32.42) while the differences in terms of BIC are small (4.19 and 1.57). Given that the BIC tends to select overly parsimonious models and based on the classification of [Kass and Raftery \(1995\)](#) that only BIC differences of more than 6 are “strong”, on balance specification (6) is preferred.

Whether to allow for symmetric or asymmetric effects of past forecast errors is also difficult to decide as the differences in AIC and BIC are small. We therefore report the coefficient estimates of both specifications in columns (1) and (2) of Table 9. It turns out that the asymmetry parameter γ is not statistically different from zero while the estimates of the other coefficients are largely unaffected by restricting it to zero. This is not unexpected as the unconditional skewness is small. We thus conclude that the symmetric specification (6) is a sufficient description of the conditional volatility process that drives the data.

We refer to the OLS residuals from our preferred forecasting regression as “statistical” forecast errors and fit the same symmetric and asymmetric volatility models as for the firms’ forecast errors, see columns (3) and (4) of Table 9. Again, the asymmetry parameter is not significantly different from zero, and restricting it to zero leaves the other parameter estimates fairly unchanged. Therefore, we take the symmetric specification—which is unanimously favored by the information criteria—as a sufficient description of the conditional volatility process that characterizes the statistical forecast errors.

6.4 Comparison of subjective uncertainty and conditional volatility

We are now ready to compare subjective uncertainty as measured by the span between the best case and worst case scenarios with our two statistical measures of conditional volatility. For the latter, we use the predicted conditional standard deviation of our preferred symmetric specification (6). Descriptive statistics for all three measures are displayed in Table 10. First note that the number of usable observations shrinks to 932

Table 9: Conditional volatility specification (6) estimated by maximum likelihood

Dependent variable	Firms' forecast errors		Statistical forecast errors	
	(1)	(2)	(3)	(4)
<i>Mean equation</i>				
Intercept (μ)	0.298 (0.252)	0.291 (0.251)	0.171 (0.270)	0.122 (0.281)
<i>Volatility equation: baseline parameters</i>				
Lagged absolute FE (α_1)	0.0852* (0.0511)	0.102* (0.0538)	0.00830 (0.0682)	0.00918 (0.0579)
Lagged volatility (α_2)	0.235*** (0.0874)	0.215*** (0.0799)	0.236** (0.0921)	0.229*** (0.0824)
Asymmetry (γ)	0 (.)	0.478 (0.317)	0 (.)	3.359 (21.85)
<i>Volatility equation: parameters of exogenous regressors</i>				
Dummy medium low volatility	0.504*** (0.0867)	0.506*** (0.0886)	0.518*** (0.0725)	0.517*** (0.0708)
Dummy medium high volatility	0.794*** (0.0873)	0.797*** (0.0874)	0.873*** (0.0684)	0.879*** (0.0674)
Dummy high volatility	1.336*** (0.110)	1.322*** (0.111)	1.519*** (0.101)	1.520*** (0.0964)
Negative sales growth in $t - 1$	-0.0286*** (0.00910)	-0.0309*** (0.00879)	-0.0259*** (0.00792)	-0.0278*** (0.00801)
Positive sales growth in $t - 1$	0.0131** (0.00515)	0.0107* (0.00576)	0.00710 (0.00522)	0.00468 (0.00605)
Dummy small firms	-0.139 (0.118)	-0.138 (0.119)	-0.142 (0.0867)	-0.127 (0.0884)
Dummy medium sized firms	-0.223** (0.105)	-0.230** (0.105)	-0.182** (0.0744)	-0.170** (0.0766)
Dummy large firms	-0.248** (0.111)	-0.269** (0.111)	-0.170** (0.0822)	-0.156* (0.0871)
Dummy 'bad' sales growth trend	0.194** (0.0882)	0.167* (0.0858)	-0.0281 (0.0652)	-0.0456 (0.0673)
Dummy 'good' sales growth trend	0.199** (0.0948)	0.209** (0.0958)	0.0648 (0.0815)	0.0818 (0.0857)
Intercept (β_0)	1.101*** (0.180)	1.125*** (0.167)	1.204*** (0.151)	1.199*** (0.146)
Number of observations	949	949	949	949
Number of firms	292	292	292	292
AIC	6871.6	6871.0	6877.1	6878.5
BIC	6939.6	6943.8	6945.1	6951.4

Sales growth rates are in the interval (-100,100). Standard errors in parentheses, clustered by firm.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

because 17 quarter-firm observations we used to construct forecast errors have either a missing upper or lower bound, or both, in the data, and thus we cannot compute a span for these observations.

Table 10: Summary statistics for measures of subjective uncertainty and conditional volatility

Variable	Mean	Std. Dev.	P10	P25	P50	P75	P90
Span between worst and best case forecast	12.1	9.8	4	5	10	15	25
Predicted conditional standard deviation of firms' forecast errors	9.7	7.2	3.5	4.8	7.5	12.3	18.3
Predicted conditional standard deviation of statistical forecast errors	9.7	7.3	3.4	4.2	7.2	12.2	19.8

A first notable result is that the distributions of the predicted conditional standard deviations of the firms' and the statistical forecast errors are remarkably similar. In fact, the correlation of the two volatility predictions is 0.97. Furthermore, we find that the average subjective forecast interval between a firm's best case and worst case scenarios is only 25 percent larger than the average of the predicted conditional standard deviations. This reflects the previous result that almost half of the realizations lie outside the subjective interval: firms probably do not think of scenarios as tail events as a, say, 95 percent statistical prediction interval should exhibit a length of four times the standard deviation.

Table 11: Comparison of subjective uncertainty and conditional volatility

	(1) Subjective uncertainty	(2) Conditional volatility of firms' forecast errors	(3) Conditional volatility of statistical forecast errors
Uncertainty/volatility in $t - 1$	0.270*** (0.0699)	0.235*** (0.0874)	0.236*** (0.0921)
Absolute forecast error in $t - 1$	0.0373 (0.0383)	0.085* (0.051)	0.008 (0.068)
Negative sales growth in $t - 1$	-0.285*** (0.0768)	-0.223*** (0.073)	-0.220*** (0.064)
Positive sales growth in $t - 1$	0.112*** (0.0422)	0.102** (0.043)	0.060 (0.044)
Dummy small firms	-3.507* (2.066)	-1.227 (1.068)	-1.316 (0.811)
Dummy medium sized firms	-3.920* (2.056)	-1.891 (0.947)	-1.656** (0.709)
Dummy large firms	-4.562** (2.022)	-2.073** (1.006)	-1.550* (0.790)
Dummy 'bad' sales growth trend	2.219*** (0.780)	1.498** (0.706)	-0.234 (0.540)
Dummy 'good' sales growth trend	-0.432 (0.530)	1.542** (0.756)	0.565 (0.727)
Dummy medium low turbulence	0.660 (0.523)	2.216*** (0.459)	2.294*** (0.361)
Dummy medium high turbulence	3.714*** (0.773)	4.094*** (0.634)	4.714*** (0.525)
Dummy high turbulence	4.528*** (0.883)	9.462*** (1.314)	12.057*** (1.552)
No. of observations	1514	949	949

Note: In the first column, pooled OLS regression coefficients are displayed, with clustered standard errors listed below the coefficients. The second and third column show average partial effects, with standard errors listed below the coefficients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In a final step, we compare the dynamics of subjective uncertainty to the dynamics of conditional volatility. Since the volatility models link the conditional standard deviation to past sales growth and the dummies via an exponential function that ensures nonnegativity, we base our comparison on average partial effects. Formally, we compute, for each rhs variable $x_{j,it}$ of the volatility equation (8), the average partial effect

$$APE_j = \sum_i \sum_t \frac{\partial \sigma_{it}}{\partial x_{j,it}}$$

and report them in columns (2) and (3) of Table 11 while column (1) replicates the coefficient estimates of the dynamic linear model for span reported in column (6) of Table 7 with the only difference that we replace the two insignificant regressors “positive past forecast error” and “negative past forecast error” by the single regressor “absolute forecast error” to conform with the specification of the volatility models.

The results indicate that the dynamics of subjective uncertainty and “objective” conditional volatility are remarkably similar. There is mild but statistically significant persistence; lagged forecast errors are largely irrelevant; and lagged absolute sales growth has an asymmetric effect with large negative realizations being roughly twice as important than large positive realizations. Hence, when forming uncertainty beliefs firms appear to have a pretty realistic impression of the dynamics that characterize the underlying data.

In contrast, subjective uncertainty and conditional volatility differ markedly in some cross sectional dimensions. Most importantly, firms operating in a highly turbulent environment report, relative to firms operating in a calm environment, a subjective uncertainty that is on average 4.5 percentage points higher while the difference in terms of the conditional volatility of their forecast errors and the statistical forecast errors amounts to 9.5 and 12.1 percentage points, respectively. Hence, volatile firms are more uncertain than others but not quite enough—they seem to underestimate the volatility they face. We also observe an opposite overestimation effect: firms on a bad sales growth trend are more uncertain than conditional volatility suggests, while larger firms feel more certain than justified by conditional volatility.

To summarize, an average firm’s updating of subjective uncertainty over time closely resembles the dynamics of conditional volatility, while its level of uncertainty appears to be fairly small on average being of roughly the same order of magnitude as the conditional standard deviation of their forecast error and in several cases fails to adequately reflect the environment the firm is operating in.

References

- Abberger, K., Sauer, S., Seiler, C., 2011. Der test des tests im ifo konjunkturtest handel. ifo Forschungsberichte 52.
- Arellano, C., Bai, Y., Kehoe, P. J., 2016. Financial frictions and fluctuations in volatility. NBER Working Paper 22990.
- Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W., Zafar, B., 2015. Inflation expectations and behavior: Do survey respondents act on their beliefs? *International Economic Review* 56 (2), 505–536.
- Augenblick, N., Rabin, M., 2017. Belief movement, uncertainty reduction, and rational updating. mimeo, Harvard University.
- Bachmann, R., Elstner, S., Sims, E. R., 2013. Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics* 5 (2), 217–249.
- Bachmann, R., Moscarini, G., 2012. Business cycles and endogenous uncertainty. mimeo, Yale University.
- Bai, J., Ng, S., 2008. Forecasting economic time series using targeted predictors. *Journal of Econometrics* 146, 304–317.
- Baley, I., Blanco, A., 2017. Firm uncertainty cycles and the propagation of nominal shocks. mimeo, University of Michigan.
- Basu, S., Bundick, B., 2017. Uncertainty shocks in a model of effective demand. *Econometrica* 85 (3), 937–958.
- Belloni, A., Chenrozhukov, V., 2013. Least squares after model selection in high-dimensional sparse models. *Bernoulli* 19 (2), 521–547.
- Ben-David, I., Graham, J. R., Harvey, C. R., 2013. Managerial miscalibration. *The Quarterly Journal of Economics* 128 (4), 1547–1584.
- Berger, D., Vavra, J., 2017. Shocks vs responsiveness: What drives time-varying dispersion? *Journal of Political Economy*, forthcoming.
- Bloom, N., 2009. The impact of uncertainty shocks. *Econometrica* 77 (3), 623–685.
- Bloom, N., 2014. Fluctuations in uncertainty. *Journal of Economic Perspectives* 28 (2), 153–176.
- Bontempi, M. E., Golinelli, R., Parigi, G., 2010. Why demand uncertainty curbs investment: Evidence from a panel of italian manufacturing firms. *Journal of Macroeconomics* 32 (1), 218 – 238.
- Ding, Z., Granger, C., 1996. Modeling volatility persistence of speculative returns: a new approach. *Journal of Econometrics* 73, 185–215.
- Ding, Z., Granger, C., Engle, R., 1993. A long memory property of stock market returns and a new model. *Journal of Empirical Finance* 1 (1), 83–106.

- Dominitz, J., Manski, C. F., 1997. Using expectations data to study subjective income expectations. *Journal of the American Statistical Association* 92 (439), 855–867.
- Dominitz, J., Manski, C. F., 2004. How should we measure consumer confidence? *Journal of Economic Perspectives* 18 (2), 51–66.
- Efron, B., Hastie, T., Johnstone, I., Tibshiranie, R., 2004. Least angle regression. *Annals of Statistics* 32 (2), 407–499.
- Elliott, G., Timmermann, A., 2016. Forecasting in economics and finance. *Annual Review of Economics* 8, 81–110.
- Fernández-Villaverde, J., Guerrón-Quintana, P. A., Kuester, K., Rubio-Ramírez, J., 2015. Fiscal volatility shocks and economic activity. *American Economic Review* 105 (11), 3352–3384.
- Gennaioli, N., Ma, Y., Shleifer, A., 2016. Expectations and investment. *NBER Macroeconomics Annual* 30, 379–442.
- Gilchrist, S., Sim, J. W., Zakrajšek, E., 2014. Uncertainty, financial frictions, and investment dynamics. NBER Working Paper 20038.
- Guiso, L., Jappelli, T., Pistaferri, L., 2002. An empirical analysis of earnings and employment risk. *Journal of Business & Economic Statistics* 20 (2), 241–253.
- Guiso, L., Jappelli, T., Terlizzese, D., 1992. Earnings uncertainty and precautionary saving. *Journal of Monetary Economics* 30 (2), 307–337.
- Guiso, L., Parigi, G., 1999. Investment and demand uncertainty. *The Quarterly Journal of Economics* 114 (1), 185–227.
- Hastie, T., Tibshirani, R., Friedman, J., 2017. *The elements of statistical learning*. Springer, 2nd edition, corrected 12th printing.
- Hurd, M. D., McGarry, K., 2002. The predictive validity of subjective probabilities of survival. *The Economic Journal* 112 (482), 966–985.
- Ilut, C., Kehrig, M., Schneider, M., 2017. Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *Journal of Political Economy*, forthcoming.
- Ilut, C., Valchev, R., 2017. Economic agents as imperfect problem solvers. mimeo, Duke University.
- Juster, F. T., 1966. Consumer buying intentions and purchase probability: An experiment in survey design. *Journal of the American Statistical Association* 61 (315), 658–696.
- Juster, F. T., Suzman, R., 1995. An overview of the health and retirement study. *The Journal of Human Resources* 30, S7–S56.
- Karanasos, M., Kim, J., 2006. A re-examination of the asymmetric power arch model. *Journal of Empirical Finance* 13, 113–128.
- Kass, R., Raftery, A., 1995. Bayes factors. *Journal of the American Statistical Association* 90 (430), 773–795.

- Koriat, A., L. S., Fischhoff, B., 1980. Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory* 6 (2), 107–118.
- Leahy, J., Whited, T., 1996. The effect of uncertainty on investment: Some stylized facts. *Journal of Money, Credit and Banking* 28 (1), 64–83.
- Lehrer, S., Xie, T., 2017. Box office buzz: Does social media data steal the show from model uncertainty when forecasting for hollywood? *The Review of Economics and Statistics* 99 (5), 749–755.
- Lütkepohl, H., 2005. *New Introduction to Multiple Time Series Analysis*. Springer.
- Manzan, S., 2015. Forecasting the distribution of economic variables in a data-rich environment. *Journal of Business & Economic Statistics* 33 (1), 144–164.
- Mietzner, D., 2009. *Strategische Vorausschau und Szenarioanalysen*. Hamburg: Springer Verlag.
- Morgan, M. G., 2017. *Theory and practice in policy analysis*. Cambridge University Press.
- Nelson, D., 1991. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* 59 (2), 347–370.
- Oskamp, S., 1965. Overconfidence in case-study judgments. *Journal of Consulting Psychology* 29 (3), 261–265.
- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B* 58 (1), 267–288.
- Varian, H., 2014. Big data: new tricks for econometrics. *Journal of Economic Perspectives* 28 (2), 3–28.
- Zivot, E., 2009. Practical issues in the analysis of univariate garch models. In: Mikosch, T., Kreiß, J., Davis, R., Andersen, T. (Eds.), *Handbook of Financial Time Series*. Springer.

Appendix A Sample Creation

In this Section, we describe the data cleaning process. The numbers of observations that remain after each step are listed in Table 12. Our original sample from 14 survey rounds consists of 5,564 observations.

Table 12: Sample creation

	No. obs.	Excluded
Original sample	5,564	
<i>Text comments</i>		
Wrong reference time excluded	5,464	100
Uncertain data quality excluded	5,431	33
<i>Outliers and Inconsistencies</i>		
Outliers excluded	5,395	36
Inconsistencies in answers to question 2 excluded	5,161	234
<i>Number of observations by firm</i>		
Firms with less than five survey answers to question 1 excluded	2,945	2,216
Answers to questions 1 and 2.a available	2,762	183
Answers to questions 1, 2.a, and 2.b available	2,710	52
<i>Adjacent forecast errors</i>		
Lag of forecast error and answer to question 1 available	949	1,761

We start by checking the free text comments respondents can give below each question, see Figure 1 for the questionnaire. We exclude 100 observations for which a comment indicates that the respondent was not able to calculate sales growth rates on a quarterly basis. For example, some firms stated that they use annual growth rates instead. Moreover, we drop 33 observations for which the comment raises doubts about the validity of the answer. For example, some firms are not able to state realized past growth rates and use estimates instead. Overall, we exclude 133 observations based on text comments leaving us with 5,431 observations.

Subsequently, we exclude outliers and inconsistencies in two steps. First, we define outliers as those quarter-firm observations for which at least one of the three following conditions are satisfied:

1. The growth rate in the previous quarter elicited in question 1 lies outside the interval $[-100; 100]$.
2. The best case and worst case sales growth rates elicited in question 2.a lie outside the intervals $[-100; 300]$ and $[-100; 100]$, respectively.
3. The expected growth rate elicited in question 2.b lies outside the interval $[-100; 100]$.

We set the upper bounds quite high because large (two-digit) growth rates typically appear to be deliberate responses as many text comments reveal. Firms give explanations such as “Many projects were moved into this quarter” and “Invoice of a major project”. We exclude 36 outliers, which leaves us with 5,395 observations.

In the second step, we check whether respondents consistently order the numbers in question 2 as worst case < forecast < best case. We exclude 234 quarter-firm observations with the orderings worst case \geq forecast \leq best case or worst case \leq forecast \geq best case because it is unclear what the respondents had in mind with these answers. However, we kept 80 quarter-firm observations with inverse ordering worst case \geq forecast \geq best case by swapping the worst case and best case numbers. Most likely inverse orderings were not intended by the respondent but they were a mistake when filling in the questionnaire.

After outlier and consistency checks, a sample of size 5,161 valid observations remains. However, in our empirical analysis we condition on time averages for each firm and therefore require a minimum number of five observations per firm. This entails a loss of 2,216 observations. Since we are interested in the relationship between past growth and current span between best case and worst case growth rates, we additionally require that firms answer questions 1 and 2.a. This reduces the number of observations by 183 to 2,762 which we define as our baseline sample.

For exercises additionally involving the answer to question 2.b, the data set reduces slightly further to 2,710 observations.

For some exercises, we need consecutive forecast errors. This reduces the sample to 1,329 observations of which 380 observations are used as lagged “pre-sample” observations so that an effective sample size of 949 firm quarters remains.

Appendix B Definition of Supersectors

Table 13 presents our definition of 14 supersectors. They are based on the 24 two-digit manufacturing sectors, which are defined by the WZ08 code of the German Statistical Office. The column with the number of observations by supersector refers to our baseline sample of 2,762 observations.

Table 13: Definition of supersectors

Supersector	Sector WZ08	Sector WZ08 name	No. of obs.
1	10, 11, 12	Food products; Beverages; Tobacco products	184
2	13, 14, 15	Textiles; Wearing apparel;	66
		Leather and related products	
3	16, 17, 31	Wood, products of wood and cork except furniture, articles of straw and plaiting materials; Paper and paper products; Furniture	286
4	18	Printing and reproduction of recorded media	191
5	19, 20, 21	Coke and refined petroleum products; Chemicals and chemical support; Basic pharmaceutical products and pharmaceutical preparations	262
6	22	Rubber and plastic products	228
7	23	Other non-metallic mineral products	133
8	24	Basic metals	120
9	25	Fabricated metal products, except machinery and equipment	324
10	26	Computer, electronic and optical products	102
11	27	Electrical equipment	201
12	28	Machinery and equipment n.e.c.	445
13	29, 30	Motor vehicles, trailers and semi-trailers; Other transport equipment	116
14	32, 33	Other manufacturing; Repair and installation of machinery and equipment	104
All			2,762

Appendix C Robustness of selecting span as preferred measure of subjective uncertainty

We construct different measures of subjective uncertainty. Our baseline measure is the span between the expected best case sales growth rate, x_b , and the expected worst case sales growth rate, x_w . It is constructed solely based on the answers to question 2.a. However, it neglects the additional information from the answer to question 2.b., the expected sales growth rate in the current quarter, x_e . To assess whether taking this additional information into account yields different measures of subjective uncertainty, we impose a distributional assumption, fit the distribution separately to each firm-quarter observation using the three available sales growth rates, x_b , x_w , and x_e , and compute two alternative implied dispersion moments—the standard deviation and the inter-quartile range—as measures of subjective uncertainty.

When fitting the distributions, it is not clear how to interpret the answer to question 2.b, x_e . It is certainly a measure of central tendency of the subjective distribution of beliefs but we deliberately do not use statistical terminology in the questionnaire. As our baseline measure, we use the mode because it characterizes, in a loose sense, the “most likely” outcome (it corresponds to the highest density). As a robustness check we also use the mean which firms would report if they had in mind an “average” outcome.

Moreover, we take into account the possibility that the respondents do not understand best and worst case sales growth rates as maximum and minimum of their underlying distribution of beliefs but rather some upper and lower percentiles. As a further robustness check, we therefore assume the best and worst case sales growth rates reflect the 95% and 5% percentiles of that distribution which leads to additional measures of subjective uncertainty.

We fit two alternative distributions with bounded support, the triangular and beta distribution. A triangular distribution is fully characterized by the upper and lower boundary of its support and its mode or mean, hence the information provided on the questionnaire is sufficient to pin down its parameters. In contrast, a four-parameter beta distribution with pdf $f(y) = c(y-l)^{\alpha-1}(u-y)^{\beta-1}$, where c is an integration constant, requires an additional restriction. Again, x_b and x_w can be used to identify the upper and lower boundaries, u and l , of its support but x_e is not sufficient to pin down the remaining two parameters α and β . Therefore, we constrain α and β such that $\alpha + \beta = r$, where r is an arbitrary number such that $\alpha > 0$ and $\beta > 0$. Using this constraint and treating x_e as the mean of the distribution, we can calculate α and β . The constraint is much less restrictive than it may appear because the implied variance of the distribution can be easily shown to be $v(r) = (x_e - x_w)(x_b - x_e)/(r + 1)$. It follows that any two variances $v(r_1)$ and $v(r_2)$ are perfectly correlated across firms. Hence, the differences in the implied measures of dispersion we use—the standard deviations and the inter-quartile range—should be small. In the robustness checks reported below, we choose $r = 10$ to ensure that the probability density function of the beta distributions is not u-shaped, which would be the case for $r \leq 2$.

Table 14 lists all resulting measures of subjective uncertainty, their correlation with the span, and the number of available observations. Since the correlation coefficients are all greater than 0.98, all subjective uncertainty measures essentially contain the same information. For two reasons, we use the span as our baseline measure. First, it is the simplest measure and does not require any distributional assumptions. Second, the sample for which we can construct the span (2,762 observations) is larger than the sample

for which we can compute any of the other measures of subjective uncertainty.¹⁰

Table 14: Comparison of span with alternative measures of subjective uncertainty

Measure of subjective uncertainty	Answer 2.a	Answer 2.b	Correlation with span	No. of observations
<i>Baseline</i>				
span	upper/lower bound	-	1.000	2,762
<i>Triangular Distribution</i>				
Standard deviation	upper/lower bound	mode	0.998	2,710
Standard deviation	95%/5% percentiles	mode	1.000	2,710
Standard deviation	upper/lower bound	mean	0.999	1,408
Interquartile range	upper/lower bound	mode	0.996	2,710
Interquartile range	95%/5% percentiles	mode	0.999	2,710
Interquartile range	upper/lower bound	mean	0.996	1,408
<i>Beta Distribution</i>				
Standard deviation	upper/lower bound	mode	0.989	2,692
Standard deviation	95%/5% percentiles	mode	1.000	2,692
Standard deviation	upper/lower bound	mean	0.992	2,395
Interquartile range	upper/lower bound	mode	0.984	2,692
Interquartile range	95%/5% percentiles	mode	1.000	2,692
Interquartile range	upper/lower bound	mean	0.988	2,395

¹⁰First, not all respondents who answered to question 2.a also answered to question 2.b. This reduces the baseline sample to 2,710 observations. In addition, the mean of the triangular distribution needs to be in the interval $[(2x_w + x_b)/3, (x_w + 2x_b)/3]$. The answer to question 2.b often does not fall in this interval, which reduces the data set to 1,408 observations. In case of the beta distribution, the parameters α and β are not defined if the three sales growth rates from question 2 are identical. This reduces the number of observations by 18 to 2,692. When treating the answer to question 2.b. as the mean, μ , of the beta distribution, we require that $x_w < \mu < x_b$, which reduces the number of observations to 2,395.

Appendix D Uncertainty is Change by Firm Characteristics

In this Section we show that the V-shaped relationship between sales growth and subjective uncertainty holds separately, and in a quantitatively similar manner, for all firm-level subgroups: the four firm size groups, the four turbulence groups, and the three growth trend groups.

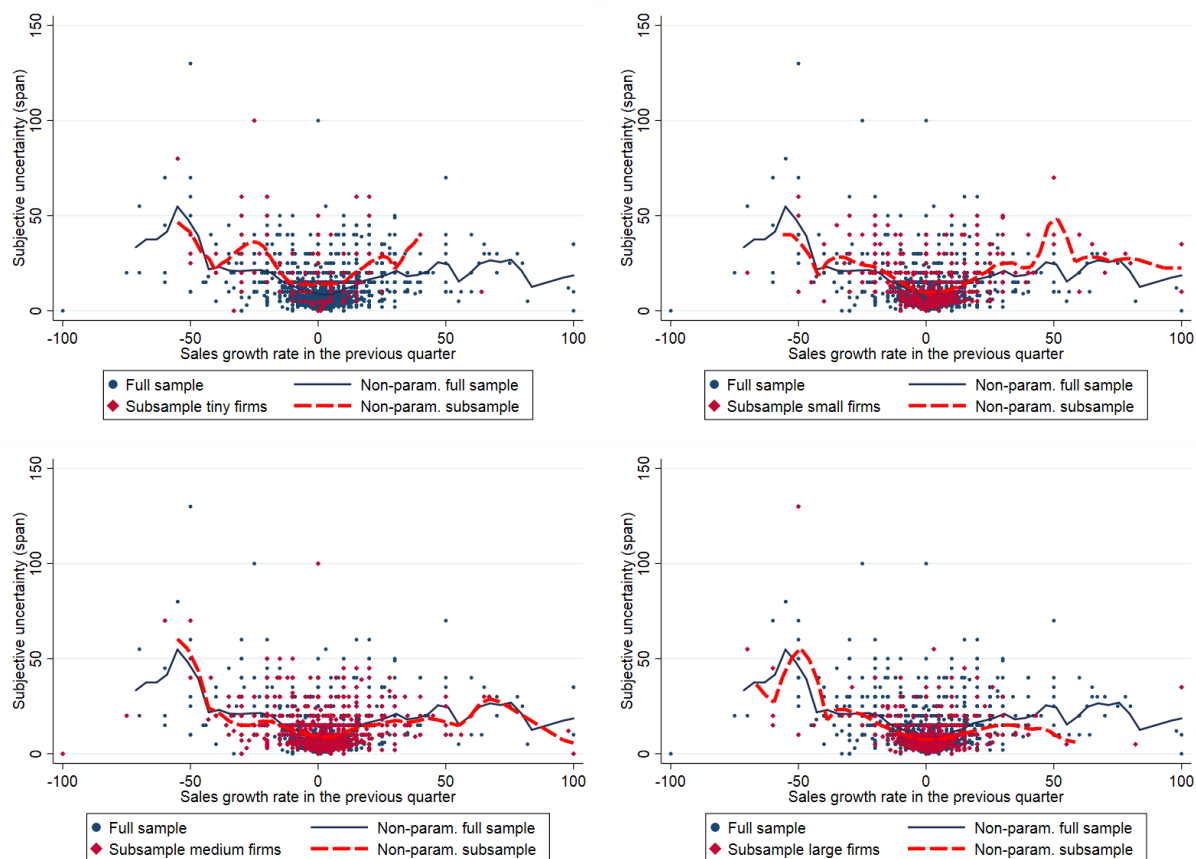


Figure 5: Relationship between subjective uncertainty (span) and sales growth in the previous quarter for four different firm size groups: tiny, small, medium, and large firms (full sample=blue, subsample=red)

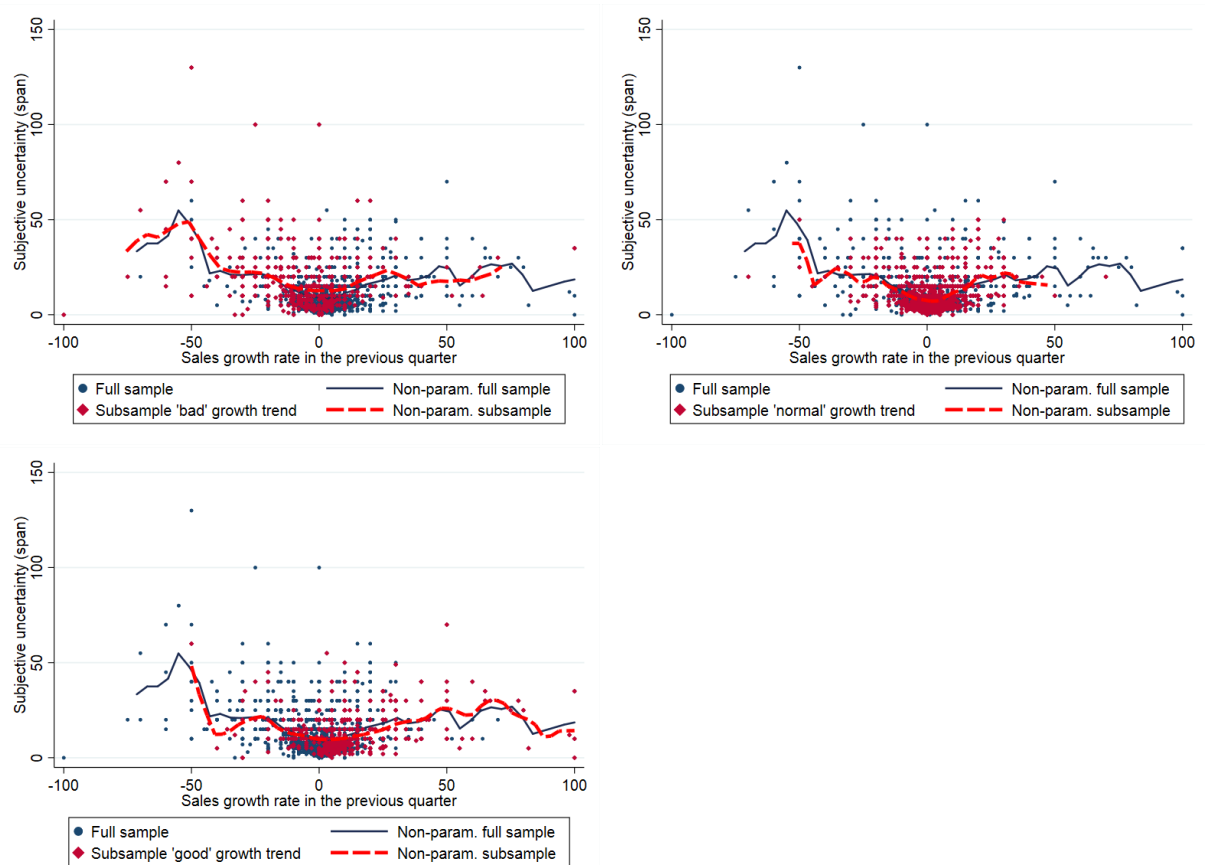


Figure 6: Relationship between subjective uncertainty (span) and sales growth in the previous quarter for three different firm trend growth groups: ‘bad’, ‘normal’, and ‘good’ trend growth (full sample=blue, subsample=red)




Figure 7: Relationship between subjective uncertainty (span) and sales growth in the previous quarter for four different firm volatility groups: low, medium low, medium high, and high volatility (full sample=blue, subsample=red)

Appendix E Robustness of Main Results

In order to test the robustness of our results, we replicate the regressions in Table 4 with the relationship between subjective uncertainty and past sales growth for two different samples.

In Table 15 we use a smaller sample, for which sales growth rates are not in the interval $(-100,100)$, but in the interval $(-15,15)$. This is the interdecile range of sales growth as depicted in table 1. The number of observations reduces from 2,762 to 2,316. Table 16 presents estimation results for a larger sample: unlike in the baseline sample, we do not require firms to have five observations with a sales growth rates, but we include all firms with at least two of such observations. This increases the sample from 2,762 to 4,120 observations. Our results are robust to these alternative sample definitions.

Table 15: Subjective uncertainty and past sales growth, for sales growth rates in the interval (-15,15)

	(1) POLS	(2) POLS	(3) POLS	(4) POLS	(5) POLS	(6) POLS	(7) FE	(8) POLS
Sales growth rate in the previous quarter	-0.0801*** (0.0269)							
Negative sales growth rate in the previous quarter		-0.582*** (0.0593)	-0.556*** (0.0577)	-0.456*** (0.0586)	-0.391*** (0.0588)	-0.308*** (0.0570)	-0.284*** (0.0424)	-0.284*** (0.0608)
Positive sales growth rate in the previous quarter		0.383*** (0.0539)	0.367*** (0.0535)	0.375*** (0.0538)	0.248*** (0.0502)	0.265*** (0.0474)	0.202*** (0.0392)	0.259*** (0.0478)
Dummy small firms			-2.838 (1.959)			-2.254 (1.496)		-2.524* (1.468)
Dummy medium sized firms			-3.999** (1.927)			-2.497* (1.480)		-2.739* (1.481)
Dummy large firms			-5.421*** (1.963)			-3.663** (1.507)		-3.991*** (1.537)
Dummy 'bad' sales growth trend				3.947*** (0.884)		2.666*** (0.839)		2.670*** (0.839)
Dummy 'good' sales growth trend				0.657 (0.640)		-0.216 (0.641)		-0.650 (0.664)
Dummy medium low turbulence					1.467*** (0.558)	1.262** (0.567)		1.200** (0.608)
Dummy medium high turbulence					4.876*** (0.689)	4.380*** (0.693)		4.287*** (0.687)
Dummy high turbulence					7.233*** (0.952)	6.512*** (0.934)		6.104*** (0.940)
Constant	10.58*** (0.315)	7.543*** (0.423)	11.63*** (1.961)	6.888*** (0.392)	5.841*** (0.489)	8.388*** (1.514)	8.984*** (0.228)	7.221*** (1.481)
Time-sector dummies								YES
No. of observations	2316	2316	2316	2316	2316	2316	2316	2316
No. of firms	394	394	394	394	394	394	394	394
No. of parameters (excl. intercept)	1	2	5	4	5	10	1	199
R-squared	0.0060	0.078	0.10	0.11	0.18	0.21	0.63	0.28

Note: Sales growth rates are in the interval (-15,15). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 16: Subjective uncertainty and past sales growth, larger sample of firms that reported at least two sales growth rates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	POLS	POLS	POLS	POLS	POLS	POLS	FE	POLS
Sales growth rate in the previous quarter	-0.0312 (0.0223)							
Negative sales growth rate in the previous quarter		-0.506*** (0.0450)	-0.480*** (0.0440)	-0.466*** (0.0475)	-0.400*** (0.0484)	-0.363*** (0.0504)	-0.290*** (0.0601)	-0.361*** (0.0514)
Positive sales growth rate in the previous quarter		0.292*** (0.0253)	0.278*** (0.0245)	0.294*** (0.0259)	0.215*** (0.0268)	0.216*** (0.0258)	0.157*** (0.0255)	0.215*** (0.0261)
Dummy small firms			-3.443* (2.091)			-2.806 (1.895)		-2.427 (1.671)
Dummy medium sized firms			-5.922*** (2.062)			-5.018*** (1.867)		-4.626*** (1.689)
Dummy large firms			-6.898*** (2.101)			-5.821*** (1.890)		-5.527*** (1.735)
Dummy 'bad' sales growth trend				2.320*** (0.740)		1.469** (0.670)		1.345* (0.695)
Dummy 'good' sales growth trend				0.452 (0.534)		0.00528 (0.515)		-0.166 (0.510)
Dummy medium low turbulence					0.344 (0.472)	0.282 (0.463)		0.455 (0.473)
Dummy medium high turbulence					3.326*** (0.582)	3.048*** (0.573)		3.206*** (0.585)
Dummy high turbulence					5.395*** (0.768)	4.826*** (0.715)		4.745*** (0.724)
Constant	12.37*** (0.292)	8.410*** (0.325)	13.93*** (2.077)	7.948*** (0.325)	7.097*** (0.378)	11.62*** (1.869)	10.14*** (0.375)	11.40*** (2.595)
Time-sector dummies								YES
No. of observations	4120	4120	4120	4120	4120	4120	4120	4120
No. of firms	940	940	940	940	940	940	940	940
No. of parameters (excl. intercept)	1	2	5	4	5	10	1	205
R-squared	0.0023	0.20	0.23	0.21	0.24	0.26	0.61	0.31

Note: Sales growth rates are in the interval (-100,100). R-squared denotes the standard R-squared. Standard errors in parentheses, clustered by firm. * p < 0.10, ** p < 0.05, *** p < 0.01.

Appendix F List of Firm-level Controls

The “**additional firm controls**” include the following variables from the ifo Business Tendency Survey:

- **Capacity utilization:** capacity utilization at the time of the survey, between 30% and 100% in pre-determined bins provided in the questionnaire; for a capacity utilization of more than 100% a free text answer is possible (in the sample, the maximum reported capacity utilization is 120%).
- **State of business:** state of business at the time of survey, coded between 0 and 100.
- **State of business dummies:** state of business at the time of survey, reported on a three-point scale. The options are “good”, “satisfactory”, or “bad”. We define two dummies, for a good and a bad state of business.
- **Orders dummies:** order book level at the time of the survey, reported on a three-point scale. The options are “relatively high”, “sufficient (for the season)”, or “too low”. We define two dummies, for a relatively high and a too low order book level.
- **Stock of inventory dummies:** Stock of inventories at the time of the survey, reported on a four-point scale. The options are “too big”, “sufficient (for the season)”, “too small”, or “no stock-keeping”. We define three dummies, for a too big and a too small stock of inventory, and for the case of no stock-keeping.
- **Production change dummies:** Change in production in the previous month, reported on a four-point scale. The options are “increased”, “roughly unchanged”, “decreased”, or “no production”. We define three dummies, for an increased and a decreased production, and for the case of no production.
- **Demand change dummies:** Change of the demand for a certain product of the firm in the previous month, reported on a three-point scale. The options are “improved”, “unchanged”, or “deteriorated”. We define two dummies, for improved and deteriorated demand.
- **Orders change dummies:** Change of the order book level in the previous month, reported on a three-point scale. The options are “increased”, “roughly unchanged”, or “decreased”. We define two dummies, for the cases of increased and decreased order book levels.
- **Price change dummies:** Change of the price for a certain product of the firm in the previous month, reported on a three-point scale. The options are “increased”, “unchanged”, or “decreased”. We define two dummies, for the cases of increased and decreased prices.
- **Credit appraisal dummies:** Willingness of banks to give credits to firms at the time of the survey, reported on a three-point scale. The options are “accommodating”, “normal”, or “restrictive”. We define two dummies, for the cases of accommodating and restrictive prices.

- **Capacity utilization change dummies:** evaluation of the remaining technical capacity at the time of the survey, given the current order book level and expected new orders in the next 12 months, reported on a three-point scale. The options are “more than sufficient”, “sufficient”, or “not sufficient”. We define two dummies, for the cases of a more than sufficient and an insufficient remaining technical capacity.
- **Constraints to production dummies:** firms report if their domestic production activity is constraint at the time of the survey. The answer options are “ yes” or “ no”. We define one dummy for the case that domestic production constraints are prevalent.