Uncertainty, Heterogeneous Expectation Errors and Economic Activity: Evidence From Business Survey Data

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2011

vorgelegt von

Steffen Elstner

Referent: Prof. Dr. Kai Carstensen
Korreferent: Prof. Rüdiger Bachmann, Ph.D.
Promotionsabschlussberatung: 08. Februar 2012

Danksagung

Zu Beginn dieser Arbeit möchte ich mich bei allen Personen und Freunden bedanken, die mich im Verlauf der letzten Jahre unterstützt haben. An erster Stelle sind hierbei meine Betreuer Prof. Dr. Kai Carstensen und Prof. Rüdiger Bachmann, PhD, zu nennen. Herrn Carstensen gilt mein Dank für die vielen Gespräche, die für das Erstellen dieser Dissertation nötig waren. Des Weiteren hielt er mir in den letzten Monaten den Rücken frei, so dass ich mich voll und ganz auf die Vollendung dieser Arbeit konzentrieren konnte. Zum Gelingen dieses Projektes hat in den letzten zwei Jahren die enge Zusammenarbeit mit Rüdiger Bachmann beigetragen. Ich kann seine überragende Betreuung nicht hoch genug würdigen. Daneben danke ich Herrn Prof. Dr. Gebhard Flaig für die Bereitschaft, meine Dissertation als dritter Gutachter zu betreuen. Danken möchte ich auch Georg Paula und Eric Sims, aus deren Zusammenarbeit zwei gemeinsame Forschungsarbeiten entstanden sind, die Teil dieser Dissertation sind.

Diese Arbeit hat stark vom Forschungsumfeld und der angenehmen Atmosphäre in meinem Arbeitsbereich "Konjunktur und Befragungen" profitiert. So führten mich Doris Hauke, Andre Kunkel, Heike Mittelmeier, Wolfgang Ruppert, Christian Seiler, Sigrid Stallhofer und Annette Weichselberger in die Welt der Befragungsdaten des ifo Geschäftsklimaindexes ein. Hierfür bin ich sehr dankbar. Daneben haben mich Gespräche mit Christian Breuer, Oliver Hülsewig, Nikolay Hristov, Michael Kleemann, Johannes Mayr, Klaus Wohlrabe und Timo Wollmershäuser deutlich in meinem ökonomischen Denken vorangebracht und einen wichtigen Beitrag zum Gelingen dieser Doktorarbeit geliefert. Dank gilt auch Alexander Ebertz und Tim Berg, die mich beim redaktionellen Teil dieser Arbeit unterstützt haben.

Ich möchte mich auch bei Julia Koller, Christina Ziegler und Christian Grimme bedanken, die mir im Verlauf der letzten Jahre enorm geholfen haben. Der größte Dank gilt meiner Familie: Liebe Lydia und lieber Frank, ohne eure emotionale Unterstützung hätte ich diese Arbeit nie vollendet.

Contents

List of Figures

List of Tables

Preface

This dissertation consists of four self-contained empirical essays that analyze the heterogeneity of firms' expectations and the effects of uncertainty shocks and oil price hikes on economic activity. Although each essay covers a different topic, the four can be classified into two broad categories.

The first part of the thesis (Chapters 1 and 2) analyzes two causes that affect the factor reallocation process of firms. To assess the importance of factor reallocation for an economy, it is crucial to understand that changes in aggregate productivity can have two causes: changes in technical efficiency and changes in the reallocation of production factors.[1] Typically, the first type of change is used to explain total factor productivity shocks in dynamic stochastic general equilibrium (DSGE) models. However, Foster et al. (2001, 2006) show that factor reallocation, including firm entry and exit, accounts for 90 percent of retail and 50 percent of manufacturing productivity growth. In a more recent study, Petrin et al. (2011) state that factor reallocation explains a much larger share of U.S. manufacturing productivity growth than do changes in technical efficiency. In short, reallocation plays a dominant role in explaining aggregate productivity growth.[2] Therefore, it is important to discover what factors affect factor reallocation.

Heightened firm-level uncertainty is one circumstance that prevents production factors from moving from low productivity firms to high productivity firms. In his seminal work, Bloom (2009) shows that large second-moment shocks trigger a sudden halt in factor movements under non-convex adjustment costs in capital and labor.[3] Put differently, an increase in uncertainty leads to a so-called "wait-and-see"-effect. The idea behind this effect is intuitive: if firms suddenly find themselves in a more uncertain environment, they stop investing

[1] Baily et al. (1992) provide a detailed decomposition equation of aggregate productivity growth.

[2] In a different context, Hsieh and Klenow (2009) show that the long-run level of manufacturing output in China and India would be 67 percent and 153 percent higher if these countries faced the same level of reallocation frictions as U.S. manufacturing.

[3] The existence of non-convex or kinked adjustment costs is well-established in the literature (see Davis and Haltiwanger, 1992, as well as Doms and Dunne, 1998).

and hiring and the economy slips into a recession. When uncertainty lessens, economic activity will eventually experience a revival and even overshoot its initial production level due to pent-up factor demand. Bloom et al. (2010) and Alexopoulos and Cohen (2009) claim that uncertainty shocks explain a sizeable proportion of business cycle fluctuations. Taking this finding as its inspiration, Chapter 1 analyzes the impact of time-varying business uncertainty on economic activity.[4] Of note is that the "wait-and-see"-effect needs only three years to play out, making it a high-frequency effect. In our study we address this point by using monthly business survey data from the United States and Germany to investigate the relationship between uncertainty and economic activity within a structural vector autoregression framework. Using business survey data has the advantage of capturing a subjective element of *decision-maker* uncertainty as opposed to that of outside experts. Specifically, these business survey data allow the construction of two complementary proxies of true ex ante uncertainty: ex ante disagreement and ex post forecast error variance. After incorporating these uncertainty measures into our empirical framework, we find little evidence of the high-frequency "wait-and-see"-effect, i.e. a large decline in economic activity after an uncertainty shock, followed by a quick rebound. To the contrary, our analysis provides evidence that increased business uncertainty leads to a slow and protracted decline in economic activity. Or, in other words, sudden increases in uncertainty have negative long-run, rather than short-run, effects on economic activity. In addition, adverse long-run "supply" shocks lead to increases in measured uncertainty. Our results are consistent with two economic environments: uncertainty shocks cause very low-frequency negative effects on activity; or high uncertainty events are merely a by-product of bad economic times: recessions breed uncertainty.

Unlike uncertainty shocks, oil price hikes *do* explain periods of high factor reallocation. Davis and Haltiwanger (2001) find that oil price shocks trigger an enormous labor reallocation process within the U.S. economy. Bresnahan and Ramey (1993), Lee and Ni (2002) and Ramey and Vine (2010) show that this is especially true for the U.S. automobile sector. In the early 1970s and 1980s, this sector was mainly specialized in producing large and fuel-inefficient cars. As a consequence, this sector was especially hard hit by the oil-price-hike-caused shifts in demand toward smaller cars. Due to the existence of large adjustment costs these shifts in demand triggered large and costly reallocation processes that ultimately contributed to a large extent to the U.S. recessions in the early 1970s and 1980s. As is the

[4] This chapter is based on joint work with Rüdiger Bachmann and Eric Sims. It is based on our paper "Uncertainty and Business Activity: Evidence From Business Survey Data," Working Paper 16143, National Bureau of Economic Research, 2010.

case in the United States, the automobile sector is very important to the German economy and, thus, the consequences of the oil price hike in 2007/08 on the German economy are worthy of study and, indeed, are the subject covered in Chapter 2.[5] To gauge the economic consequences of the 2007/08 oil price for Germany, it is crucial to understand that the oil price is not determined exogenously. As noted by Kilian (2009), the oil price responds to structural oil demand and supply shocks that, in turn, have different consequences for the German economy. The oil price hike of 2007/08 was almost entirely driven by increasing world demand (see Hamilton, 2009, Hicks and Kilian, 2011, and Kilian, 2009). This fact is especially important in the case of Germany, for which, as an export economy, the positive indirect effects on domestic production of a booming world economy can far overcompensate for the negative direct effects of an oil price increase. To address this issue, in Chapter 2, we implement the structural vector autoregression framework proposed by Kilian (2009) that distinguishes between supply shock driven and demand shock driven oil price changes. Our results show that supply shock driven and demand shock driven oil price hikes have different impacts on the German economy. We find that consumption always reacts negatively to any kind of shock, but that the impact on exports and gross investment depends on the type of oil price shock. In the cases of the oil demand shocks, favorable international price movements and shifts in global demand toward German export goods initially outweigh the negative effects on consumer' expenditures and, therefore, lead to an increase in GDP. Even though strong oil price surges do not burden German manufacturing, which primarily produces investment and export goods, we find that their effects on domestic demand become rather negative over time. Concerning the economic consequences of the 2007/08 oil price hike we find that the sustained sequence of positive world demand shocks, that triggered the oil price hike, led to a 2.3 percent reduction in German GDP in the year 2009. We thus provide evidence that this particular oil price hike made a notable contribution to the subsequent recession in Germany.

The second half of the thesis (Chapters 3 and 4) focuses on heterogeneity in firms' expectations. Zimmermann and Kawasaki (1986), Nerlove (1983) and Souleles (2004) provide the motivation for both chapters. Specifically, we use survey micro data from the IFO Business Climate Survey (IFO-BCS) to analyze firms' expectation errors. One problem that typically arises when using survey data is how to deal with the qualitative nature of these data. While useful, in that more firms are inclined to participate in a survey

[5]This chapter is based on joint work with Kai Carstensen and Georg Paula. It is based on our paper "How Strongly Did the 2007/08 Oil Price Hike Contribute to the Subsequent Recession," mimeo, 2011. This is a revised version of our working paper that circulated under CESifo-WP 3357.

when the information demands are low, qualitative information has limits, particularly when forecasting errors need to be aggregated over time so as to measure the long-run average forecasting errors of firms and possible biases therein. However, under certain assumptions, we can combine the qualitative three-month ahead production outlook from the monthly survey with the quantitative change in percentage capacity utilization from the quarterly supplement to compute idiosyncratic quarterly, one-quarter-ahead production expectation errors. We do this for the manufacturing part of the IFO-BCS from 1980 onward and thus construct a panel of quarterly production expectation errors for a period of 30 years. This data set is then used to answer two questions raised in Chapters 3 and 4.

In Chapter 3, we take a look at whether firms suffer from expectational bias.[6] If the answer is yes (and it is), we then determine the extent of welfare losses stemming from those expectational biases. Firms' expectations are unbiased if their subjective probabilities with respect to future economic states are not distorted. This implies that their long-run average expectation errors are not significantly different from zero. Therefore, we identify expectational biases by testing for each firm whether its average expectation error is significantly different from zero. If this is the case, we conclude that this firm has an expectational bias. Using this procedure, we find that, depending on the exact definition of our quantitative production expectation error, at least 6 percent and at most 34 percent of firms consistently over- or underpredict their one-quarter-ahead upcoming production. In a further step, we investigate the implications of these expectational biases by performing a simple welfare calculation. We use a frictionless heterogeneous firm model where firms decide about their factor demands before they know their idiosyncratic productivity levels. We calibrate the fractions of optimistic and pessimistic firms and the extent of their expectational biases to the distributional properties of production expectation errors in the IFO-BCS. Overoptimistic firms hire too many workers and build up capital stocks that are too high. Overpessimistic firms do not demand enough inputs. We then compare the welfare in an economy populated by firms with a distribution of production expectation errors that approximates the one in the data to a world populated only by firms with zero long-run expectation errors. We robustly find that the welfare losses from expectational errors are small, probably smaller even than conventional estimates of the welfare costs of business cycles.

The final chapter of the thesis contains a more detailed analysis of the heterogeneity of expectations and expectation errors. Weale and Pesaran (2006) point out that expectation heterogeneity arises due to differences in subjective probability densities (belief disparity)

[6] This chapter is based on joint work with Rüdiger Bachmann. It is based on our paper "Firms' Optimism and Pessimism: Evidence From the IFO Survey," mimeo, 2011.

and differences in individual-specific information sets (information disparity). This study is motivated by the idea that firms may differ with respect to their beliefs and with respect to their ability to process information. Specifically, I am interested in discovering how many firms have rational expectations in a traditional sense, i.e. their expectations are unbiased and they use all available information efficiently. In my conservative estimate of expectation errors, I find that about two thirds of all the firms in my sample have rational expectations, i.e. their expectations are unbiased and they use all available information efficiently. However, under a broader definition, this number decreases substantially to slightly more than 30 percent. Thus, there is evidence that heterogeneous firms form their expectations in heterogeneous ways, i.e. a large proportion of firms differ with respect to their beliefs and their ability to process information.

Bibliography

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

Baily, M., C. Hulten, and D. Campbell (1992): "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity*, 1, 187–249.

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

Bloom, N., M. Floetotto, and N. Jaimovich (2010): "Really Uncertain Business Cycles," *mimeo.*

Bresnahan, T. F., and V. A. Ramey (1993): "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry," *American Economic Review Papers and Proceedings*, 83(2), 213–218.

Davis, S. J., and J. Haltiwanger (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(3), 819–863.

——— (2001): "Sectoral Job Creation and Destruction to Oil Price Changes," *Journal of Monetary Economics*, 48(1), 465–512.

Doms, M., and T. Dunne (1998): "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1(2), 409–429.

Foster, L., J. Haltiwanger, and C. Krizan (2001): "Aggregate Productivity Growth: Lessons From Microeconomic Evidence," in *New Developments in Productivity Analysis.* University of Chicago Press.

——— (2006): "Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s," *The Review of Economics and Statistics*, 88(4), 748–758.

Hamilton, J. D. (2009): "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 1, 215–261.

Hicks, B., and L. Kilian (2011): "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?," *mimeo*.

Hsieh, C.-T., and P. J. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 74(4), 1403–1448.

Kilian, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99(3), 1053–1069.

Lee, K., and S. Ni (2002): "On the Dynamic Effects of Oil Price Shocks:A Study Using Industry Level Data," *Journal of Monetary Economics*, 49(2), 823–852.

Nerlove, M. (1983): "Expectations, Plans, and Realizations in Theory and Practice," *Econometrica*, 51(5), 1251–1279.

Petrin, A., T. K. White, and J. P. Reiter (2011): "The Impact of Plant-Level Resource Reallocations and Technical Progress on U.S. Manufacturing Growth," *Review of Economic Dynamics*, 14, 3–26.

Ramey, V. A., and D. J. Vine (2010): "Oil, Automobiles and the U.S. Economy: How Much Have Things Really Changed?," *NBER Macroeconomics Annual*.

Souleles, N. S. (2004): "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 28(3), 39–72.

Weale, M., and M. H. Pesaran (2006): "Survey Expectations," in *Handbook of Economic Forecasting*, ed. by G. Granger, and A. Timmermann, pp. 715–776.

Zimmermann, K. F., and S. Kawasaki (1986): "Testing the Rationality of Price Expectations for Manufacturing Firms," *Applied Economics*, 18, 1335–1347.

Chapter 1

Uncertainty and Business Activity: Evidence From Business Survey Data

Abstract[1]

What is the impact of time-varying business uncertainty on economic activity? We construct empirical measures of uncertainty based on business survey data from the U.S. and Germany. We show that measured uncertainty is robustly negatively correlated with economic activity far into the future. In particular, adverse "supply" shocks lead to large increases in measured uncertainty. In contrast, innovations in measured uncertainty uncorrelated with shocks identified as having a permanent impact on production have quantitatively small impacts on economic activity. Our results are consistent with two economic environments: uncertainty shocks cause rather low-frequency negative effects on activity, or high uncertainty events are mainly a by-product of bad economic times – recessions breed uncertainty.

[1]This chapter is based on joint work with Rüdiger Bachmann and Eric Sims. It is based on our paper "Uncertainty and Business Activity: Evidence From Business Survey Data," Working Paper 16143, National Bureau of Economic Research, 2010.

1.1 Introduction

What is the impact of time-varying business uncertainty on economic activity? The seminal contribution in Bloom (2009) has renewed interest in the aggregate effects of time-varying uncertainty and influenced a growing literature in macroeconomics, which we will discuss in detail below. In this paper we use (partly confidential) data from business surveys to investigate the relationship between uncertainty and economic activity within a structural vector autoregressions (SVAR) framework.

These business surveys contain, on a monthly basis, qualitative information on the current state of, and expectations regarding, firms' business situations. In particular, we use disagreement in business expectations for the Third Federal Reserve District Business Outlook Survey (BOS) to measure business uncertainty. Using dispersion of expectations as a measure of uncertainty has a long tradition in the literature: Zarnowitz and Lambros (1987) show with the NBER-ASA expert forecasts of output growth and inflation that disagreement and intrapersonal uncertainty are positively correlated.[2] While we do not have probabilistic forecasts of individual business situations, the confidential micro data of the German IFO Business Climate Survey (IFO-BCS) allow us to compare the disagreement-based measure of uncertainty with a qualitative index of the forecast error variance of production expectations. We find that the two uncertainty measures are positively correlated and that their impact on economic activity is qualitatively and quantitatively similar and statistically often indistinguishable.

High-frequency business survey data from narrowly defined segments of the economy are well-suited to measure the impact of uncertainty on economic decision-making for several reasons. First, business survey data capture a subjective element of uncertainty for actual decision makers, as opposed to outside experts. Second, we will show that our business uncertainty measure explains a higher fraction of the total forecast error variance of economic activity variables than volatility measures based on stock market returns. Third, the recent literature (Bloom, 2009, and Bloom et al., 2010) has highlighted the so-called "wait-and-see"-effect of uncertainty: if firms find themselves in a more uncertain environment, they stop hiring and the economy slips into a recession. Positive shocks to uncertainty can thus lead to short run fluctuations, starting with a rapid decline in economic activity, then a rebound phase and prolonged overshoot after approximately six months. As discussed more in Section 1.2, "wait-and-see"-dynamics are thus rather short-run and rely on adjustment frictions, which render high-frequency data the best candidate to detect these dynamics.

[2]Other examples in the literature that either find significant positive correlations between these two measures or use disagreement as a proxy for uncertainty are: Federer (1993), Bomberger (1996), Giordano and Soederlind (2003), Bond and Cummins (2004), Fuss and Vermeulen (2008), Clements (2008), Popescu and Smets (2010) and Bloom et al. (2010).

Readily available at a monthly frequency, business survey data have an advantage over balance sheet data, which are only available at lower frequencies. Fourth, our use of dispersion in survey responses to proxy for uncertainty rests on the assumption that respondents draw their idiosyncratic shocks from similar distributions, so that fluctuations in dispersion are the result of fluctuations in uncertainty and not merely compositional changes in the cross-section. Using data from narrowly defined segments of the economy makes this assumption more likely to hold. Finally, the confidential micro data allow us to compare expectations and realizations of economic variables and thus to construct two complementary proxies for uncertainty: ex ante disagreement and ex post forecast error variance.

We begin by estimating low-dimensional SVARs featuring the survey-based uncertainty indices and measures of economic activity within a sector. We order uncertainty first, so that innovations to uncertainty can affect economic activity immediately. We find that positive innovations to uncertainty have protracted negative effects on economic activity. The effect on impact and at high frequencies is small. This is a robust result across specifications and surveys. While they do not appear to be consistent with the aforementioned high-frequency "wait-and-see"-effect, "wait-and-see"-dynamics could be combined with an endogenous growth mechanism – R&D investment, for example – to generate the observed protracted negative implications for economic activity. In addition, we also suggest a new interpretation: the "by-product"-hypothesis. In this view, high uncertainty events are merely reflective of bad economic times, rather than their cause.

To investigate further, we then impose more structure and change the identification strategy. In systems featuring uncertainty, a measure of sectoral economic activity, and a measure of the aggregate unemployment rate, we identify three structural shocks. In the spirit of Shapiro and Watson (1988), Blanchard and Quah (1989), and Gali (1999), we use a long-run restriction to identify a shock which affects the level of sectoral economic activity in the long-run from the other two shocks, which can only have a transitory effect on output. We identify the uncertainty shock from the other "demand" shock by imposing that our measure of uncertainty not respond within period to the other shock. This identification "shuts down" the long-run influence of uncertainty in the hope of making its short-run impact shine through, while at the same time allowing uncertainty to have a strong temporary, short-lived effect on activity. In point of fact, however, shocks to uncertainty so identified have small effects on production and unemployment. Rather, consistent with the "by-product"-hypothesis, empirical measures of uncertainty appear to be largely driven by the long-run shock. Shocks which permanently lower economic activity give rise to significantly higher measured uncertainty on impact. This is true for survey-based uncertainty measures, as well as uncertainty measures based on the corporate bond spread over treasuries and uncertainty measures based on stock market volatility.

This conclusion is consistent with a general view of recessions as times of destroyed business practices and relationships, the reestablishment of which generates uncertainty. It accords with empirical work by Hamilton and Lin (1996), who find that high stock market volatility is driven mainly by bad economic times. It is also consistent with the theoretical models of Bachmann and Moscarini (2011) as well as Fostel and Geanakoplos (2011), who argue that bad economic times incentivize risky behavior – in the former through price experimentation, in the latter through increased leverage – and therefore endogenously lead to increased uncertainty.

Related Literature

There is a growing literature that studies the effects of uncertainty shocks in fully specified dynamic general equilibrium models. Bachmann and Bayer (2011), exploring data from a German firm-level panel, argue that the effects in Bloom (2009) and Bloom et al. (2010) are small and do not substantially alter unconditional business cycle dynamics. Chugh (2011), who explains the dynamics of leverage with shocks to micro-level uncertainty, also finds only a small business cycle impact of uncertainty shocks. Using a model with financial frictions, Gilchrist et al. (2010) argue that increases in uncertainty lead to an increase in bond premia and the cost of capital which, in turn, triggers a decline in investment activity. Arellano et al. (2011) show that firms downsize investment projects to avoid default when faced with higher uncertainty. Schaal (2010) uses a directed search model with uncertainty shocks to understand the recent labor market behavior. Basu and Bundick (2011) study uncertainty shocks in a sticky price environment. Fernandez-Villaverde et al. (2011) argue that positive shocks to interest rate volatility depress economic activity in several Latin American economies.

There is another literature that, like this paper, estimates the impacts of various uncertainty proxies on economic activity. Leahy and Whited (1996) is one of the first papers to document empirically a negative relationship between uncertainty and firms' investment. Bond and Cummins (2004) use data on publicly traded U.S. companies to show that various measures of uncertainty predict prolonged declines of firms' investment activities. Gilchrist et al. (2009) find a similar result for increases in the dispersion of firms' sales growth. Christiano et al. (2010), in a large-scale DSGE context, also find a strong low-frequency impact of the identified risk shock. Alexopoulos and Cohen (2009) use a narrative approach in a structural VAR framework (the incidence of the words "uncertainty" and "economy" in *New York Times* articles) and find high-frequency decline-rebound-overshoot dynamics. Popescu and Smets (2010) show, again with structural vector autoregressions and for German expert survey data, that it is shocks to risk aversion rather than innovations to uncertainty that explain roughly 10%-15% of output fluctuations.

The remainder of the paper is organized as follows. The next section discusses the "wait-and-see"-mechanism and delivers a benchmark against which we compare our empirical results. The third section describes the business survey data we use. The fourth section presents the main results and interprets them. Details and additional results are relegated to various appendices.

1.2 Uncertainty and Activity: "Wait-and-See"

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

Figure 1.1: Replication of Wait-and-See in Bloom (2009)

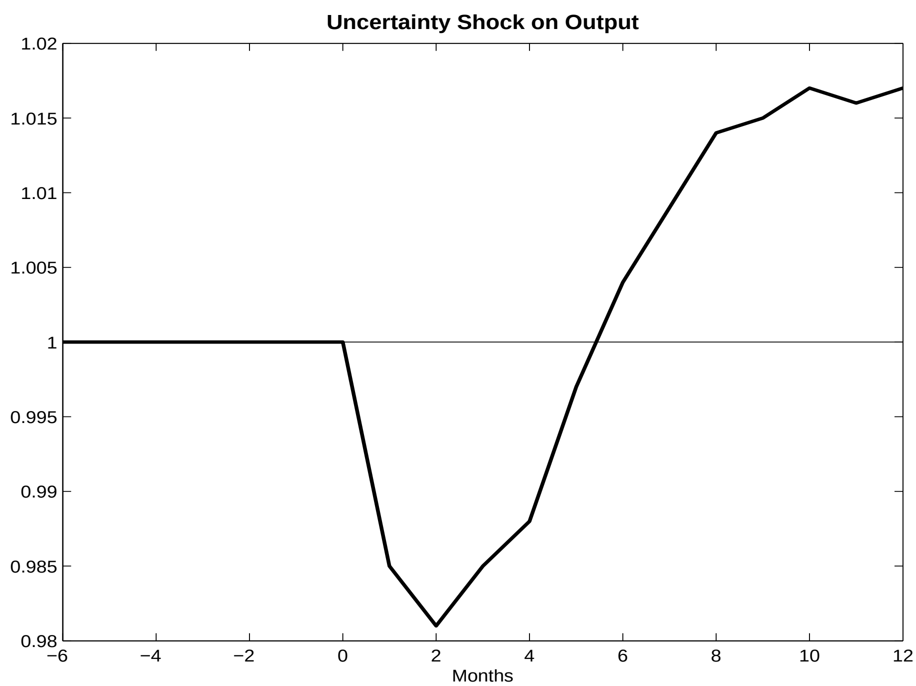

Notes: This graph is a replication of the simulated model IRF of output to an uncertainty shock, see Figure 12 in Bloom (2009).

Time-varying uncertainty at the firm level may have economic consequences when there is a degree of irreversibility to firm actions (see Bernanke, 1983, as well as Dixit and Pindyck, 1994). For a concrete example, suppose that a firm faces fixed costs to adjusting the size of its labor force and/or physical capital stock. Suppose further that there is a meanpreserving spread on the distribution of future demand for the firm's product. With fixed adjustments costs, higher uncertainty over future demand makes new hiring and investment less attractive. If a large fixed cost must be paid to adjust the firm's labor or capital, then there is reason to minimize the number of times this cost must be paid. If the future is very uncertain (in the sense that demand could be either very high or very low relative to the present), then it makes sense to wait until the uncertainty is resolved to undertake new hiring and investment. Why pay a large fixed cost now when a highly uncertain future means that one will likely have to pay the fixed cost again?

An increase in uncertainty thus makes inaction relatively more attractive. Given a reduction in hiring, employment, and hence output, will fall through exogenous separations. As the future begins to unfold, demand or productivity conditions are, in expectation, unchanged. There will be pent up demand for labor and capital. Inaction today moves firms closer to their adjustment triggers in subsequent periods, leading to expected increases in hiring, investment and a general rebound and even overshoot in economic activity, followed by a return to steady state. Figure 1.1 provides an example of an impulse response of output to an increase in uncertainty, replicated from the model in Bloom (2009).

This theoretical impulse response highlights an important aspect as pertains to our empirical work. The economic implications of uncertainty shocks in a model with "wait-and-see"-effects are decidedly high-frequency in nature. Thus, an empirical study of uncertainty that wants to detect "wait-and-see"-effects should make use of high-frequency data, which is one of the reasons why we use monthly surveys in this paper.

1.3 Measuring Business Uncertainty

We construct uncertainty measures from the Third FED District Business Outlook Survey (BOS) and the German IFO Business Climate Survey (IFO-BCS). In the next subsection we briefly describe the characteristics of each and list the main survey questions we use to measure business uncertainty. We then define the variables used in the empirical analysis, followed by a subsection on the cyclical properties of these variables.

1.3.1 Data Description

BOS

The Business Outlook Survey is a monthly survey conducted by the Federal Reserve Bank of Philadelphia since 1968. The survey design has essentially been unaltered since its inception. It is sent to large manufacturing firms in the Third FED District, which comprises the state of Delaware, the southern half of New Jersey, and the eastern two thirds of Pennsylvania. The survey questionnaire is of the "box check" variety. It asks about firms' general business expectations as well as their expectations and actual realizations for various firm-specific variables such as shipments, workforce and work hours. Respondents indicate whether the value of each economic indicator has increased, decreased, or stayed the same over the past month. They are also asked about their expectations for each indicator over the next six months. Whenever possible, the survey is sent to the same individual each month, typically the chief executive, a financial officer or other person "in the know". Participation is voluntary. The group of participating firms is periodically replenished as firms drop out or a need arises to make the panel more representative of the industrial mix of the region. Each month 100-125 firms respond. As noted by Trebing (1998), occasional telephone interviews are used to verify the accuracy of the survey responses.

The advantages of the BOS are its long time horizon, its focus on one consistent, economically relatively homogenous class of entities – large manufacturing firms in one region –, an unparalleled number of questions that are useful for our research question and the fact that for each question there is a "current change" and an "expectation" version. Its main drawback is the relatively small number of respondents. Nevertheless, given its advantages, we use the BOS for our baseline results.[3] We focus on the following two questions (the other questions we use from the BOS are documented in Appendix A2.1):

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

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

Both questions are phrased, somewhat ambiguously, about general business conditions. Trebing (1998) notes, however, that answers to these questions are highly correlated with responses to the shipments question, which is phrased as explicitly company specific. He concludes that both series are essentially indicators of firm-specific business conditions.

In addition, in order to construct an employment turnover indicator, we use the following question:

Q 3 *"Company Business Indicators: Number of Employees [LAST MONTH] vs. [CURRENT MONTH]: decrease, no change, increase?"*

[3] Appendix A4 supplements the baseline results with an analysis of the U.S. Small Business Economic Trends Survey (SBETS). There is a concern that if adjustment costs grow less than proportionally with firm size the firms in the BOS may be sufficiently large that adjustment costs do not matter for them, and therefore "wait-and-see"cannot be detected in the BOS. The SBETS also has larger cross-sections of firms compared to the BOS. We find essentially the same results.

IFO-BCS

The German IFO Business Climate Survey is one of the oldest and broadest monthly business confidence surveys available (see Becker and Wohlrabe, 2008, for more detailed information). However, due to longitudinal consistency problems and availability of micro data in a processable form only since 1980, we limit our analysis to the manufacturing sector from 1980 until the present. From 1991 on, the sample includes East-German firms as well.

One of the IFO-BCS's main advantages is the high number of survey participants. The average number of respondents at the beginning of our sample is approximately 5,000; towards the end the number is about half that at 2,500.[4] Participation in the survey is voluntary and there is some fraction of firms that are only one-time participants. However, conditional on staying two months in the survey, most firms continue on and this allows us to construct a measure of ex post forecast error uncertainty. Our final sample of continuing firms comprises roughly 4,000 respondents at the beginning and 2,000 towards the end of the sample. In terms of firm size, the IFO-BCS contains all categories. In the survey for January 2009, for example, about 12% of respondents had less than 20 employees, roughly 39% had more than 20 but less than 100 employees, 43% of the participants employed between 100 and 1000 people and less than 7% had a workforce of more than 1000 people.

The two main questions that allow us to construct a qualitative index of ex-post forecast errors are:[5]

Q 4 *"Expectations for the next three months: Our domestic production activities with respect to product XY will (without taking into account differences in the length of months or seasonal fluctuations) increase, roughly stay the same, decrease."*

Q 5 *"Trends in the last month: Our domestic production activities with respect to product XY have (without taking into account differences in the length of months or seasonal fluctuations) increased, roughly stayed the same, decreased."*

1.3.2 Variable Definitions

Survey answers fall into three main categories, *Increase*, *Decrease*, and a neutral category. We use these categories to define our expectation-based index of uncertainty and one index of current economic activity. Define $Frac_t^+$ as the fraction of "increase"-responses to a survey

[4]The IFO-BCS is a survey at the product level, so that these numbers do not exactly correspond to firms.
[5]Here we provide a translation, for the German original see Appendix A3.1.

question at time t; $Frac_t^-$ is defined analogously. We start with the uncertainty index, constructed for questions like Q 1 and Q 4:

$$Uncertainty_t \equiv sqrt\left(Frac_t^+ + Frac_t^- - \left(Frac_t^+ - Frac_t^-\right)^2\right).$$

Notice that $Uncertainty_t$ so defined is the cross-sectional standard deviation of the survey responses, if the $Increase$-category is quantified by $+1$ and the $Decrease$-category by -1 and the residual categories by 0. This is a standard quantification method for qualitative survey data. Next, we define a current index of economic activity for questions like Q 2 and Q 5. Summing up variables that essentially measure changes is intended to capture a qualitative measure of the level of economic activity:

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

1.3.3 Is Cross-sectional Dispersion a Good Proxy for Uncertainty?

Measuring the (subjective) uncertainty of decision makers is inherently difficult. Ideally, one would like to elicit a subjective probability distribution over future events from managers, as has been done in Guiso and Parigi (1999) for Italian firms. With this probability distribution it is straightforward to compute a measure of intrapersonal uncertainty for firms' decision makers. However, to the best of our knowledge such probability distributions are not available repeatedly and over long time horizons.[6] Researchers have to rely on proxies instead. Although frequently done in the literature, using the cross-sectional dispersion of firms' expectations as a proxy for firms' uncertainty is not without potential problems. First, time-varying cross-sectional dispersion in firms' survey responses might be due to different firms reacting differently to aggregate shocks even with constant uncertainty. Notice that for relatively homogenous samples such as the BOS this is likely to be less of a problem. Secondly, time variation in the dispersion of expectations might be the result of time variation in the heterogeneity of said expectations, without these expectations reflecting a higher degree of uncertainty on the part of the business managers.

We address the first concern – different firms having different factor loadings to aggregate shocks – by a variance decomposition of the IFO-based (based on Q 4, to be specific) uncertainty measure, $(Uncertainty_t^{IFO})^2$, into the average within-variance and the between-variance of the 13 manufacturing subsectors contained in the IFO-BCS (see Appendix A3.2

[6] Bontempi et al. (2010), using the same Italian data sets as Guiso and Parigi (1999), construct eight years of annual uncertainty measures from the max-min range of firms' one-year ahead sales forecasts.

for details). The idea behind this decomposition is that such differences in factor loadings to aggregate shocks might be due to industry-specific production and adjustment technologies. Figure A9 in Appendix A3.2, however, shows that the time series of $(Uncertainty_t^{IFO})^2$ is not explained by the between-variance of the manufacturing subsectors. This means it is not explained by the manufacturing subsectors getting more or less different over the business cycle.

To address the second concern – the relationship between (time-varying) dispersion, uncertainty and cross-sectional shock variance – we present in Appendix A1 a simple and highly stylized two-period model where firms receive signals about their uncertain future business situations. We show for this model that if signals are neither perfectly informative nor perfectly uninformative, under Bayesian updating both the dispersion of firms' expectations and the average subjective uncertainty in the cross-section increase in response to an increase in the cross-sectional variance of firms' future business situations.

Furthermore, the confidential micro data in the IFO-BCS and its panel structure allow us to construct a qualitative index of the ex post forecast error standard deviation, which by construction excludes heterogeneous, but certain, changes in expectations.[7] The basic idea is that we can compare firms' answers about their production expectations, Q 4, with their answers on past production realizations, Q 5, and thus construct a measure of firm-specific production expectation errors. The cross-sectional standard deviation of these expectation errors, $Uncertainty_t^{fe^{IFO}}$, is a dispersion index for the ex post forecast errors. In Appendix A3.3 we describe the construction of $Uncertainty_t^{fe^{IFO}}$ in detail.

The advantage of $Uncertainty_t^{fe^{IFO}}$ over $Uncertainty_t^{IFO}$ is that it is based on actual "uncertain-at-time-t" innovations, as opposed to potentially heterogeneous expectations about the future, which could be certain. However, the raw correlation coefficient between $Uncertainty_t^{fe^{IFO}}$ and $Uncertainty_t^{IFO}$ is reasonably high for monthly data, 0.73, and when we aggregate both series up to the quarterly level the correlation is 0.77. The fact that both conceptually different proxies for uncertainty are reasonably close to each other lends some support to the widespread practice of proxying uncertainty with survey disagreement. Most importantly, the impulse responses on economic activity look qualitatively and quantitatively similar and are statistically often indistinguishable (see Section 1.4.2).

1.3.4 Cyclicality of Business Survey Variables

In this subsection, we report basic cyclical properties of the survey-based variables introduced in Sections 1.3.2 and 1.3.3: $Uncertainty_t$, $Uncertainty_t^{fe}$ and $Activity_t$. They have been

[7] Whereas the aggregate survey responses, $Frac_t^+$ and $Frac_t^-$, are publicly available for both the BOS and the IFO-BCS, individual firm responses are not. In the case of the IFO-BCS they are available to researchers on-site.

seasonally adjusted with the SAS X12 procedure, an adaptation of the U.S. Bureau of the Census X-12-ARIMA seasonal adjustment method. Table 1.1 displays the contemporaneous correlations of the various survey-based monthly uncertainty measures with, respectively, manufacturing industrial production and the corresponding survey-based activity measures. The uncertainty indices are all countercyclical. This confirms previous findings by Bloom (2009), Bloom et al. (2010), Chugh (2011) and Bachmann and Bayer (2011), who find, using different data sources, that stock market volatility and balance-sheet-based cross-sectional measures of uncertainty are all countercyclical.[8] The correlation is even more negative when we aggregate up to the quarterly frequency.

Table 1.1: Cyclical Properties of $Uncertainty_t$ and $Uncertainty_t^{fe}$

	Monthly		Quarterly	
	Correlation with		Correlation with	
Uncertainty Measure	IP_t	$Activity_t$	IP_t	$Activity_t$
General Conditions-$Uncertainty_t^{BOS}$	-0.28	-0.47	-0.33	-0.51
Shipments-$Uncertainty_t^{BOS}$	-0.27	-0.29	-0.31	-0.32
Production-$Uncertainty_t^{IFO}$	-0.10	-0.61	-0.23	-0.62
Production-$Uncertainty_t^{fe^{IFO}}$	-0.05	-0.54	-0.24	-0.59

Notes: This table displays the unconditional contemporaneous correlations between the survey-based uncertainty variables in the rows and the month-over-month/quarter-over-quarter differences of two different activity measures in the columns. Industrial production (IP) measures are logged. The General Conditions-$Uncertainty_t^{BOS}$ measure, based on Q 1, is paired with the corresponding difference of the (seasonally adjusted) manufacturing industrial production index from the OECD main economic indicators and the General Conditions-$Activity_t^{BOS}$ measure based on Q 2. The Shipments-$Uncertainty_t^{BOS}$ measure, based on Q 6 (see Appendix A2.1), is paired with the corresponding difference of the (seasonally adjusted) manufacturing industrial production index from the OECD main economic indicators and the Shipments-$Activity_t^{BOS}$ measure based on Q 9 (see Appendix A2.1). The Production-$Uncertainty_t^{IFO}$ measure, based on Q 4, is paired with the corresponding difference of the (seasonally adjusted) manufacturing industrial production index from the German Federal Statistical Agency and the $Activity_t^{IFO}$-measure based on Q 5. Production-$Uncertainty_t^{fe^{IFO}}$ is paired with the same activity measures as the Production-$Uncertainty_t^{IFO}$ measure.

Table 1.2 displays the contemporaneous correlations of the survey-based (differenced) activity measures we constructed in Section 1.3.2 with manufacturing industrial production. These activity measures are, not surprisingly, procyclical.

[8] We also find that both uncertainty measures from the IFO-BCS, $Uncertainty_t^{IFO}$ and $Uncertainty_t^{fe^{IFO}}$, are countercyclical, separately for each of the 13 manufacturing subsectors. This excludes composition effects

Table 1.2: Cyclical Properties of $Activity_t$

Activity Measure / Correlation with	Monthly IP_t	Quarterly IP_t
General Conditions-$Activity_t^{BOS}$	0.55	0.79
Shipments-$Activity_t^{BOS}$	0.46	0.70
Production-$Activity_t^{IFO}$	0.25	0.53

Notes: This table displays the unconditional contemporaneous correlations between the differenced survey-based variables in the rows and the month-over-month/quarter-over-quarter differences of industrial production indices. Industrial production (IP) measures are logged. The General Conditions-$Activity_t^{BOS}$ measure, based on Q 2, is paired with the corresponding difference of the manufacturing industrial production index from the OECD main economic indicators. The Shipments-$Activity_t^{BOS}$ measure, based on Q 9 (see Appendix A2.1), is paired with the corresponding difference of the manufacturing industrial production index from the OECD main economic indicators. The Production-$Activity_t^{IFO}$ measure, based on Q 5, is paired with the corresponding difference of the manufacturing industrial production index from the German Federal Statistical Agency.

1.4 Results

In this section we present and discuss our main empirical results. In Choleski-identified SVARs with uncertainty ordered before economic activity variables, we robustly find that innovations to business uncertainty are associated with initally small, but slowly-building reductions in economic activity. Imposing the restriction that uncertainty shocks have no long-run effects on activity renders the responses of economic activity to uncertainty statistically and economically insignificant. Both findings are difficult to reconcile with an important "wait-and-see"-channel from uncertainty to aggregate dynamics. Rather, we find that shocks adversely impacting the economy are important drivers of various empirical uncertainty measures, suggesting that uncertainty is a consequence of bad shocks.

1.4.1 Third FED District Business Outlook Survey

We begin the analysis with the Federal Reserve Bank of Philadelphia Third District Business Outlook Survey and low-dimensional Choleski-identified SVARs containing the General Conditions-$Uncertainty_t^{BOS}$ index and various economic activity variables. We order the un-

as an explanation for the countercyclicality of the overall uncertainty measure. The numbers are available on request.

certainty index first. This gives uncertainty its "best shot" of being quantitatively important for economic activity dynamics. Figure 1.2 shows impulse responses for U.S. manufacturing industrial production (upper panel) and General Conditions-$Activity_t^{BOS}$ (based on Q 2; lower panel) to an innovation in business uncertainty.[9] Both variables enter the system in levels and we include 12 lags.[10]

Figure 1.2: Uncertainty Innovations on Manufacturing Activity

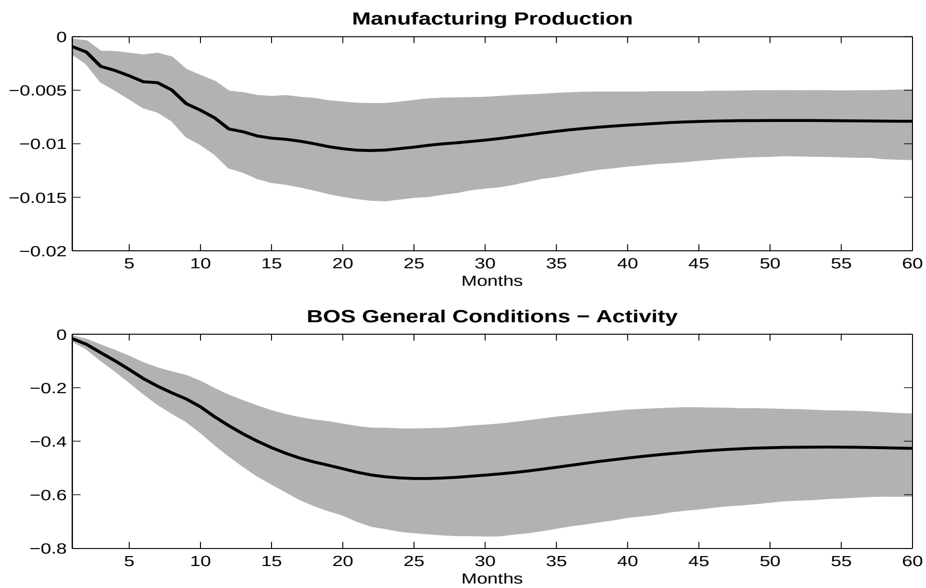

Notes: Both IRFs are based on General Conditions-$Uncertainty_t^{BOS}$, which derives from Q 1 in the BOS. The upper panel shows the response of manufacturing production to a positive uncertainty innovation in a two-variable SVAR with *Uncertainty* ordered first. Manufacturing production is the natural logarithm of the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators. The lower panel shows the response of General Conditions-$Activity_t^{BOS}$ (based on Q 2) to a positive uncertainty innovation in a two-variable SVAR with *Uncertainty* ordered first. All VARs are run with 12 lags, the confidence bands are at the 95% significance level using Kilian's (1998) bias-corrected bootstrap.

The impulse response of manufacturing production to an innovation in business uncertainty is slightly negative on impact with effects that build over time. The peak decline is at about 1 percent, occurring about two years after impact, with no tendency to revert. The lower panel of Figure 1.2 provides corroborating evidence with a different measure of

[9] One might be worried that uncertainty should not affect economic activity on impact because of various information or decision lags. For instance, one might assume that companies learn the uncertainty of their business environment only through the published surveys themselves, when they see a lot of disagreement there. Figure A3 in Appendix A2.3 presents the impulse response with economic activity ordered first. It is clear that the Choleski ordering does not drive our results.

[10] Our results are robust to alternative assumptions about how the variables enter the VAR (i.e. levels vs. differences) as well as to alternative assumptions about lag length.

sectoral economic activity. The BOS in Q 2 asks about current business conditions relative to the recent past. The impulse response of General Conditions-$Activity_t^{BOS}$ is strikingly similar to that using overall manufacturing production as the activity measure. This is particularly important, as we do not have monthly industrial production data disaggregated at the regional and sectoral level that would allow us to construct a quantitative activity measure that corresponds exactly to the BOS. The fact that the results are nearly identical across two related, but different activity measures lends credence to our finding: neither impulse response function seems to be consistent with the "wait-and-see"-dynamics as shown in Figure 1.1.[11]

Figure 1.3: Uncertainty Innovations on Manufacturing Employment

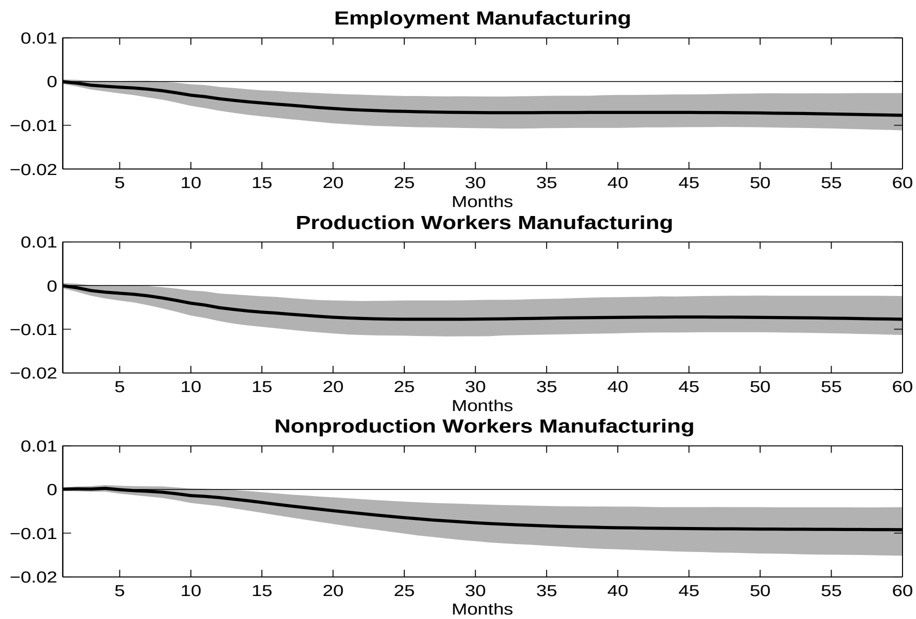

Notes: see notes to Figure 1.2. Uncertainty is ordered first. The employment measures are seasonally adjusted and logged and are taken from the BLS-CES data base.

In Figure 1.3 we show impulse responses from bivariate SVARs featuring our BOS baseline uncertainty measure and various manufacturing employment measures. The responses shown

[11] In Table A2 in Appendix A2.2 we display contemporaneous correlations of various BOS activity measures with the monthly Third FED district BLS manufacturing employment data available from 1990 on. Running the same two-variable SVAR with this employment measure as the activity variable on data from 1990 onwards results in very similar point estimates for the impulse response functions. We also compare the monthly BOS activity measures with the monthly coincident index from the Philadelphia FED, which measures overall economic, not merely manufacturing activity for the Third FED district. Using this index as the activity variable in the two-variable SVAR would yield identical results. Finally, we compare yearly averages of the BOS activity measures with the yearly NIPA manufacturing production index for the Third FED district. The BOS activity measures are positively correlated with all these other imperfect activity measures from official statistics, which shows that the BOS depicts the dynamics of real economic activity in the manufacturing sector of the Third FED district reasonably accurately.

are that of employment to uncertainty, with uncertainty ordered first. The "wait-and-see"-theory of the transmission from uncertainty shocks to business cycles emphasizes hiring and firing frictions. With these we should observe a large reduction in employment followed by a quick recovery in response to an uncertainty shock, similarly to the output response in Figure 1.1 in Section 1.2. However, the response of manufacturing employment is rather consistent with our results for production: it moves little on impact, followed by a period of sustained reductions, with no obvious tendency for reversion, even at very long horizons. Production and non-production workers, who might be subject to different adjustment costs, are affected similarly.

Another direct and related prediction of the "wait-and-see"-theory is that job turnover – defined as the sum of job creation and job destruction – should decline following an increase in uncertainty: wait and do nothing. Yet again, the survey data do not seem to support this prediction. Figure 1.4 shows the response of the extensive margin of job turnover to an innovation in uncertainty. The point estimate on and near impact is positive and insignificant from zero, turning more significant at horizons well beyond one year.

Figure 1.4: Uncertainty Innovation on BOS Job Turnover Index

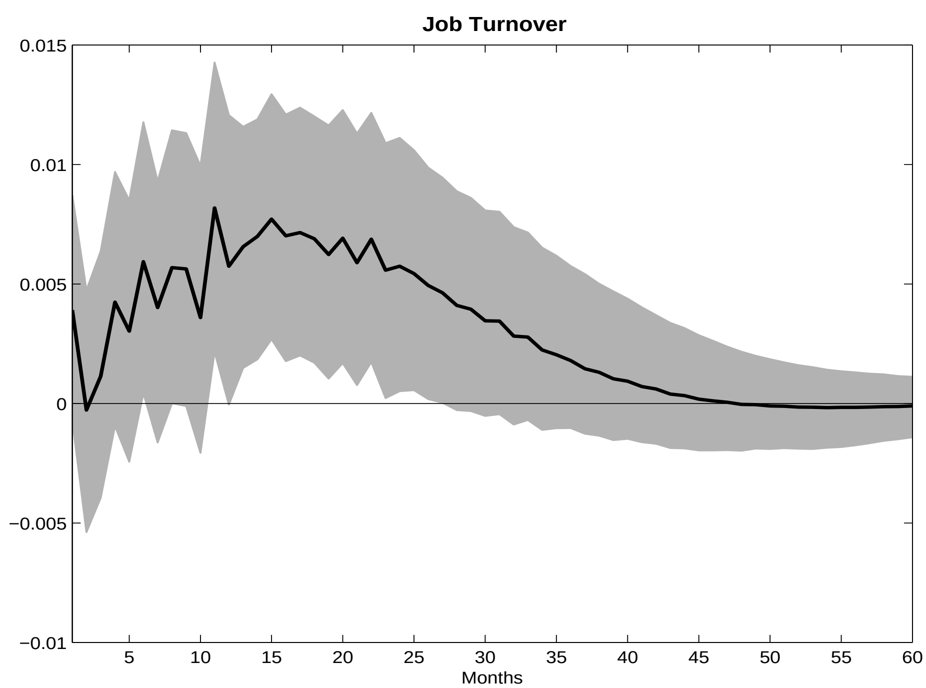

Notes: see notes to Figure 1.2. $Turnover_t \equiv Frac_t^+ + Frac_t^-$. $Turnover_t$ is based on Q 3.

For a comparison of our results with the SVAR evidence in Bloom (2009), we estimate exactly the same high-dimensional system, but replace the high uncertainty dummy variable based on stock market volatility with our General Conditions-$Uncertainty_t^{BOS}$ index. The VAR otherwise includes the S&P500 stock market index, the Federal Funds Rate, average hourly earnings, the consumer price index, hours, employment and industrial production. Uncertainty is ordered second in a recursive identification. Figure 1.5 shows the impulse response of production and employment to an innovation in General Conditions-$Uncertainty_t^{BOS}$. Although with reduced statistical significance, the pattern remains: slowly-building declines and slow recoveries of economic activity variables.

Figure 1.5: Uncertainty Innovations in the Bloom (2009) SVAR

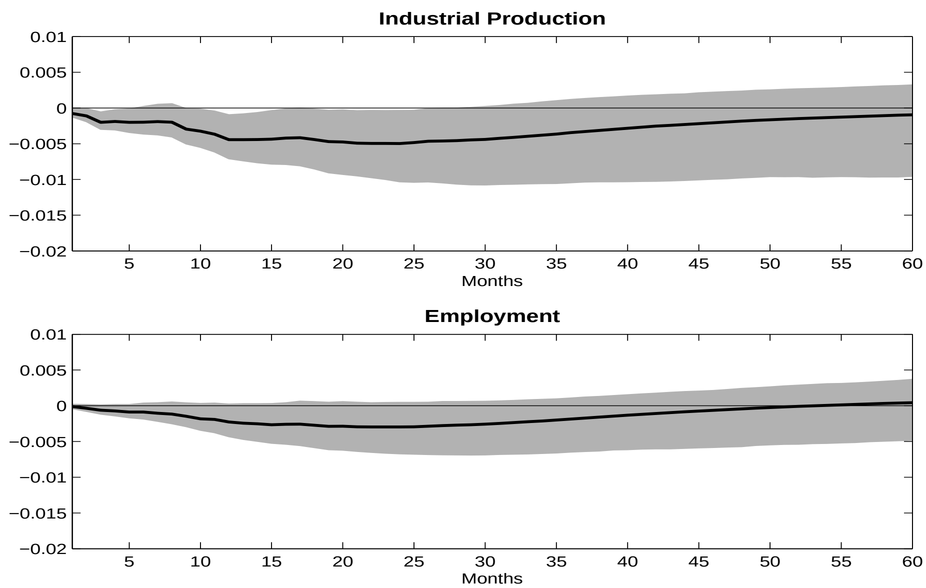

Notes: see notes to Figure 1.2. The S&P500 stock market index has been logged and is ordered first. Then follows the General Conditions-$Uncertainty_t^{BOS}$ index. Hourly Earnings, the CPI, employment and industrial production have been logged.

We also conduct a forecast error variance decomposition in this high-dimensional SVAR with uncertainty based on the BOS and compare it to the forecast error variance decomposition in the SVAR with uncertainty based on stock market volatility. On impact, the variation in production that is explained by either proxy for uncertainty is almost zero. Interestingly, the forecast error variance in production that is explained by our survey-based General Conditions-$Uncertainty_t^{BOS}$ index rises steadily to 8% at the one-year horizon, 16% at the two-year horizon and 20% at the five-year horizon. Similarly, the forecast error variance in employment that is explained by our survey-based General Conditions-$Uncertainty_t^{BOS}$ index rises steadily to 4% at the one-year horizon, 11% at the two-year horizon and 12% at the five-year horizon. In contrast, the uncertainty innovation from the high-uncertainty dummy based on stock market volatility explains never more than 3% of the forecast error variance in production at any horizon, and at most 3% of the forecast error variance in employment. These numbers are even lower when the actual volatility series is used instead of the dummy. We take this as evidence that our uncertainty measure has more explanatory power for economic activity than uncertainty measures based on stock market volatility.

We conduct many more robustness checks to our result that in Choleski-identified SVARs uncertainty innovations trigger prolonged declines in economic activity. For example, we vary the economic activity variable used in the baseline SVAR, while keeping General Conditions-$Uncertainty_t^{BOS}$ (based on Q 1) as the uncertainty measure: the BOS shipments, employment and "work hours" based activity indices and overall labor productivity in manufacturing. We also vary the uncertainty measure: an indicator variable for high uncertainty to capture uncertainty spikes as opposed to general uncertainty fluctuations, an uncertainty measure based on entropy, and uncertainty measures derived from other expectation questions in the BOS. The results are depicted in Appendix A2.3, Figures A4 to A8. The basic qualitative patterns of these impulse responses are the same as in our benchmark systems.

There are two main results from our analysis thus far – one negative and one positive. The negative result is that there is little evidence supporting the high-frequency "wait-and-see"-mechanism with a rebound, described in Section 1.2. On the positive side we have that innovations to uncertainty contain significant predictive information for the future path of sectoral economic activity.

This, in turn, leaves open two interpretations: for one, autonomous shocks to uncertainty have long-run or even permanent effects. This would be consistent with a "wait-and-see"-story where the R&D-sector is particularly heavily hit, so that persistent, but transitory uncertainty shocks could lead to permanent effects on economic activity.[12] In this case, it could well be that the high-frequency "wait-and-see"-dynamics are simply swamped by low-frequency effects, and we need to attempt to "control" for the latter.

In any event, another interpretation opens up: uncertainty could itself be generated by bad news about the future. Under this interpretation, uncertainty events are merely a by-product of bad economic times. Figure 1.6 shows results from the Choleski-identified baseline SVAR, augmented by a measure of business confidence, ordered first. We define business confidence as the difference between the fraction of positive responses and the negative responses in the business survey. As in Figure 1.2, the two upper panels use manufacturing production as the activity variable, and the two lower panels use the survey-based activity measure General Conditions-$Activity_t^{BOS}$. The two left panels show the impulse response of the uncertainty index to a negative innovation in business confidence. They are strongly

[12] The increase of measured uncertainty to an uncertainty innovation lasts about 12 months in our baseline SVAR displayed in Figure 1.2 and then dies out.

and significantly positive. Bad news increase uncertainty. On the right hand side, we see the impulse responses of economic activity to a positive innovation in business uncertainty, orthogonalized to business confidence innovations. The impulse responses from Figure 1.2 are also depicted for comparison. While the impulse responses remain small on impact and protracted over time, albeit much less so, their permanence vanishes once uncertainty innovations are orthogonalized to confidence innovations and the responses are quantitatively much smaller.

Figure 1.6: Uncertainty Innovations Orthogonalized to Confidence Innovations

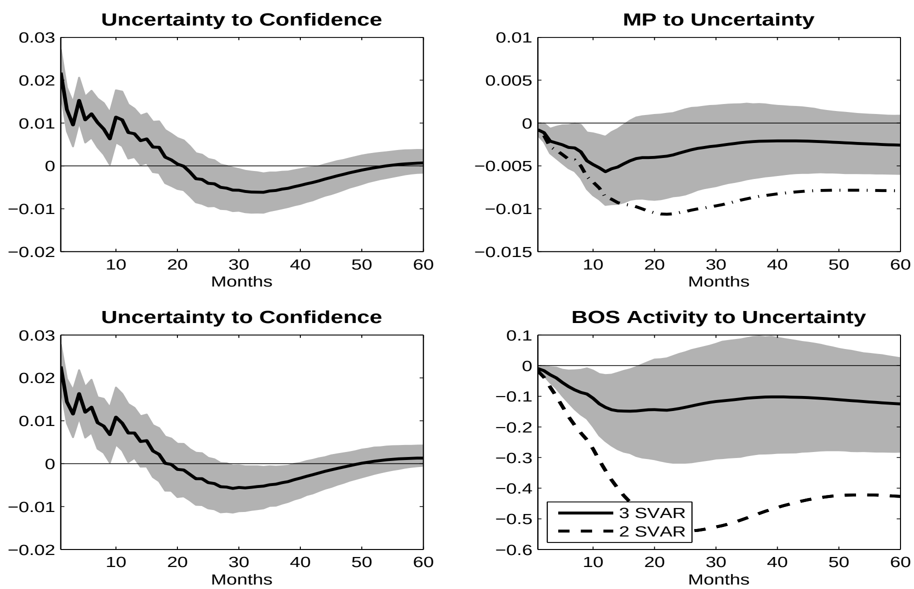

Notes: see notes to Figure 1.2. The two upper panels feature results from an SVAR with (in this ordering) General Conditions-$Confidence_t^{BOS}$, General Conditions-$Uncertainty_t^{BOS}$ and manufacturing production. General Conditions-$Confidence_t^{BOS}$ is a business confidence indicator, defined as $Confidence_t \equiv Frac_t^+ - Frac_t^-$. It is based on Q 1. In the lower panels General Conditions-$Activity_t^{BOS}$ index replaces manufacturing production as the activity variable. The two left panels show the impulse responses of the uncertainty index to a negative innovation in business confidence. The two right panels show impulse responses of economic activity to a positive innovation in business uncertainty. The dashed lines reproduce the impulse responses of activity from Figure 1.2.

To explore the "by-product"-hypothesis further, as well as to give uncertainty a better chance of leading to high-frequency "wait-and-see"-type dynamics, we now attempt to "control" for any information about long-run economic activity contained in the uncertainty measures. We do so by adopting an identification approach in the spirit of Shapiro and Watson (1988) as well as Blanchard and Quah (1989) in a three-variable VAR with General Conditions-$Uncertainty_t^{BOS}$, manufacturing production as a sectoral activity measure and the aggregate unemployment rate. We identify three structural shocks – one which can have a long-run effect on production and two which cannot. Notice that the corresponding long-run shock in our case, unlike in Blanchard and Quah (1989) who use aggregate and not sectoral production, need not literally be a productivity shock. Rather, it is any shock that permanently affects sectoral output. We identify the uncertainty shock as a shock that does not impact activity in the long-run, but can influence uncertainty and unemployment. The long-run impact of uncertainty is shut down by construction to let short-run effects of uncertainty shine through. Third, we identify a more conventional aggregate demand shock separately from the short-run uncertainty shock, where we assume that the conventional demand shock does not affect uncertainty on impact.[13]

Figure 1.7: A Three-Variable Blanchard-Quah-Type SVAR

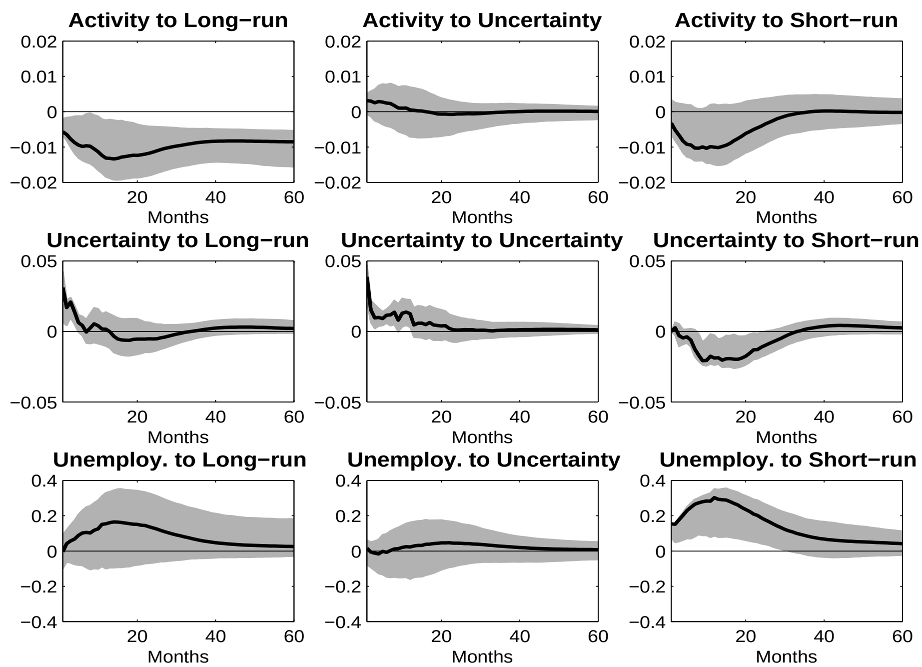

Notes: see notes to Figure 1.2. We use manufacturing production as the activity measure, and the General Conditions-$Uncertainty_t^{BOS}$ index as the uncertainty measure. The unemployment rate is the (seasonally adjusted) monthly civilian unemployment rate from the BLS. The uncertainty innovation and the conventional short-run shock are identified as shocks that do not impact manufacturing production in the long-run. The conventional short-run shock is identified as the shock that does not affect the uncertainty index on impact. The long-run shock and the conventional short-run shock have a negative sign.

Figure 1.7 shows the impulse responses in such a three-variable SVAR, and Table 1.3 the corresponding forecast error variance decomposition for horizons ranging from one month

[13]We also tried an alternative specification which identifies the uncertainty shock as the shock leading to no long-run impact on output which maximally explains variation in our uncertainty measure over various horizons (as opposed to just on impact, which is what the recursive assumption does). The results are very similar.

to five years. Two results are important: first, once the long-run impact of uncertainty is "controlled" for, there is little significant impact of uncertainty on output or unemployment left. The forecast error variance for activity is mainly driven by the long-run and the conventional short-run shock, whereas the contribution of the uncertainty shock after three months drops below 10 percent. The contribution of the uncertainty shock to the fluctuations of the unemployment rate is even smaller. Secondly, a shock which permanently lowers sectoral production is associated with an increase in uncertainty. This is consistent with the Choleski-identified results in Figure 1.6 and precisely what our "by-product"-hypothesis with respect to uncertainty implies. The forecast error variance decomposition shows that the long-run shock accounts for a significant fraction of the fluctuations in the uncertainty index, particularly in the first six months.

Table 1.3: Forecast Error Variance Decomposition - BOS

	Shock	1M	3M	6M	1Y	2Y	5Y
Activity	Long-run	62%	55%	52%	53%	64%	77%
	Uncertainty	19%	10%	6%	3%	1%	1%
	Short-run	20%	34%	42%	44%	34%	22%
Uncertainty	Long-run	39%	48%	47%	28%	21%	21%
	Uncertainty	61%	52%	51%	43%	29%	27%
	Short-run	0%	0%	2%	30%	51%	52%
Unemployment Rate	Long-run	0%	6%	11%	15%	21%	23%
	Uncertainty	1%	1%	0%	0%	1%	2%
	Short-run	99%	93%	89%	85%	77%	75%

Notes: see notes to Figure 1.7.

1.4.2 IFO Business Climate Survey

We now turn to results using the IFO Business Climate Survey, which gives us the advantage of being able to compare uncertainty measures based on ex-ante disagreement with uncertainty measures based on ex-post forecast error variance. Figure 1.8 shows the activity responses for the Choleski-identified baseline SVAR to the innovations in the two types of uncertainty we are considering: uncertainty based on the ex post forecast error standard deviation – $Uncertainty_t^{fe^{IFO}}$ – and uncertainty based on ex ante disagreement – $Uncertainty_t^{IFO}$. The activity variable is based on Q 5, the IFO current production question. The SVARs here include a dummy variable from 1991 on to account for structural breaks associated with the German reunification, though our results are insensitive to alternative ways of dealing with that event. There are two important results: first, the responses of activity to the two different measures of uncertainty are quite similar to each other, in fact statistically indistinguishable. This serves as support for our use of a disagreement measure as an uncertainty proxy. Second, the results are also similar to those from the BOS, with somewhat more evidence of reversion at longer horizons when $Uncertainty_t^{fe^{IFO}}$ is used. The impact effects on activity are small, with the trough of the negative response occurring roughly two years subsequent to the shock. This provides corroboration of the results from U.S. data in another country.

Figure 1.8: Uncertainty Innovations on Production-$Activity_t^{IFO}$

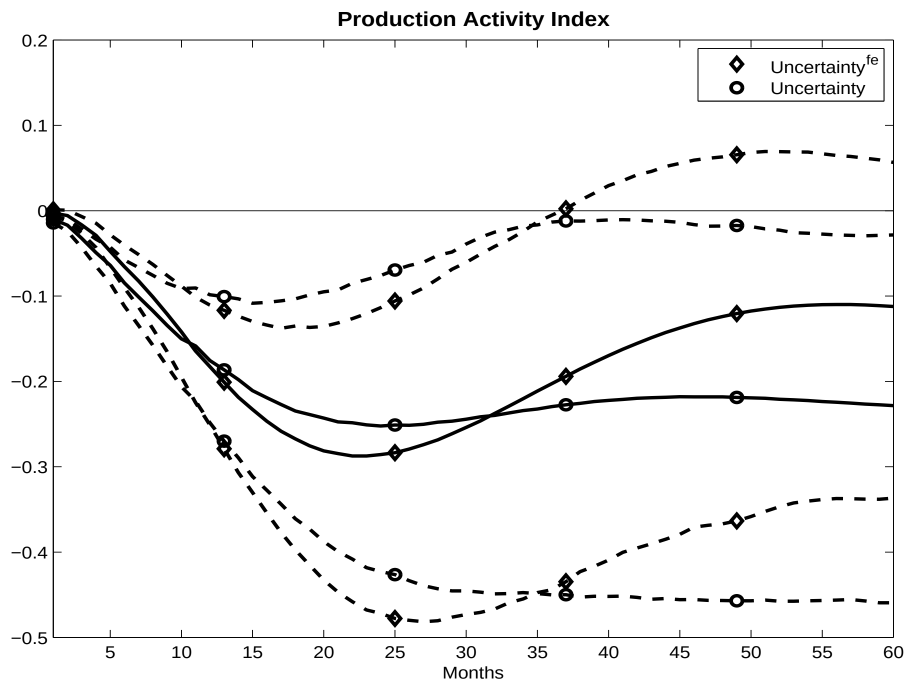

Notes: $Uncertainty_t$ is based on Q 4. $Uncertainty_t^{fe}$ is based on Q 4 and Q 5. The activity variable is based on Q 5. Uncertainty is ordered first. We include a dummy variable from 1991 to account for the German reunification. We run the VARs with 12 lags. All confidence bands are at the 95% significance level using Kilian's (1998) bias-corrected bootstrap.

We conclude by also confirming the BOS results from the three-variable Blanchard-Quah-type SVAR with Production-$Activity_t^{IFO}$, $Uncertainty_t$ and $Uncertainty_t^{fe}$, and the unemployment rate in Figure 1.9 and Table 1.4 (for the sake of readabililty, we leave out the confidence bands). We find that uncertainty measured either way has a lower impact on sectoral economic activity than in the BOS and somewhat more impact on the unemployment rate, especially for the disagreement measure $Uncertainty_t$. The impulse response to either uncertainty measure does not look like high-frequency "wait-and-see"-dynamics. We again find that a negative long-run shock has a sizeable positive impact on the uncertainty index.

The similarity between the BOS and IFO-BCS results suggests that the negative findings in Popescu and Smets (2010) as well as Bachmann and Bayer (2011) with regards to the role of uncertainty shocks as a major driving force of short-run fluctuations are not driven by their use of German data.

Figure 1.9: A Three-Variable Blanchard-Quah-Type SVAR - IFO-BCS

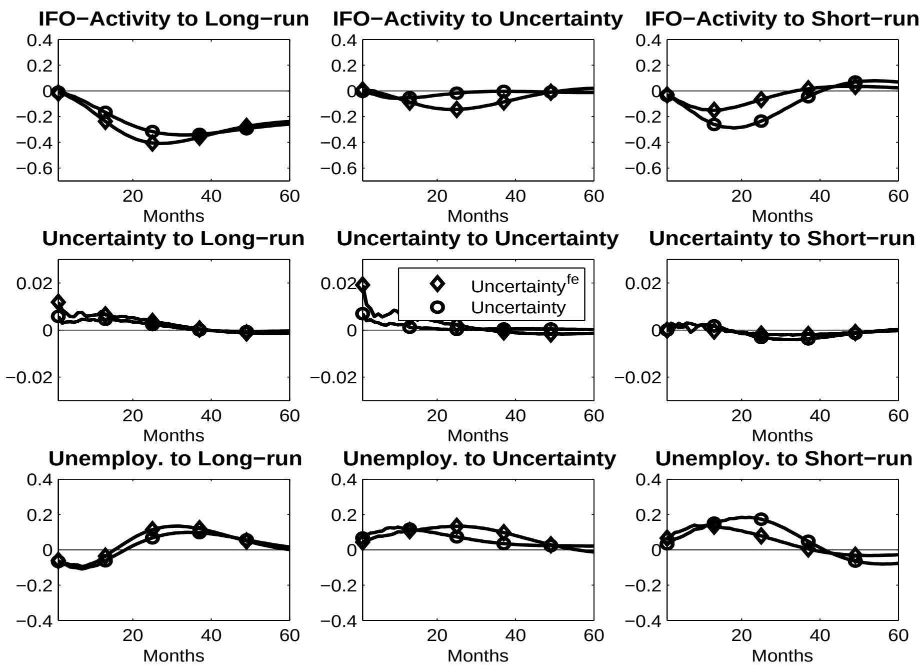

Notes: see notes to Figure 1.8. The unemployment rate is the (seasonally adjusted) monthly unemployment rate from the Bundesanstalt für Arbeit. The uncertainty shock and the conventional short-run shock are identified as shocks that do not impact manufacturing production in the long-run. The conventional short-run shock is identified as the shock that does not affect the uncertainty index on impact.

Table 1.4: Forecast Error Variance Decomposition - IFO-BCS

	Shock	1M	3M	6M	1Y	2Y	5Y
$Uncertainty_t^{fe}$							
Activity	Long-run	22%	22%	32%	51%	74%	87%
	Uncertainty	5%	1%	2%	6%	10%	7%
	Short-run	73%	77%	66%	43%	16%	6%
Uncertainty	Long-run	28%	31%	36%	40%	45%	44%
	Uncertainty	72%	67%	63%	59%	53%	51%
	Short-run	0%	2%	2%	2%	2%	5%
Unemployment Rate	Long-run	37%	35%	31%	23%	17%	33%
	Uncertainty	19%	19%	21%	25%	37%	39%
	Short-run	45%	45%	48%	52%	46%	28%
$Uncertainty_t$							
Activity	Long-run	8%	8%	13%	21%	40%	73%
	Uncertainty	1%	5%	8%	6%	2%	1%
	Short-run	91%	87%	80%	73%	58%	27%
Uncertainty	Long-run	41%	39%	41%	52%	62%	44%
	Uncertainty	59%	60%	49%	36%	25%	17%
	Short-run	0%	1%	10%	13%	13%	39%
Unemployment Rate	Long-run	44%	40%	38%	28%	14%	23%
	Uncertainty	44%	45%	41%	40%	30%	23%
	Short-run	12%	15%	20%	32%	56%	55%

Notes: see notes to Figure 1.9.

1.4.3 Discussion

In Choleski-identified SVARs with sectoral business uncertainty and sectoral economic activity variables we find protracted negative impulse responses of activity to uncertainty innovations. Job turnover reacts positively to the same shocks. A different SVAR identification identifies uncertainty shocks as having no long-run effect on production, but affecting production and unemployment on impact. An uncertainty shock so identified has little significant effect on economic activity. In contrast, a shock identified as having a permanent effect on activity is associated with significant increases in uncertainty.

Figures 1.10 and 1.11 show both these results from survey based uncertainty measures in a condensed form and compare them to results based on other uncertainty proxies used in the literature. To do so, we replace General Conditions-$Uncertainty_t^{BOS}$ with, respectively, the corporate bond spread as in Gilchrist et al. (2009), and stock market volatility as in Bloom (2009), in the three-variable Blanchard-Quah-type SVAR that leads to Figure 1.7.

Figure 1.10: Uncertainty Shock on Activity

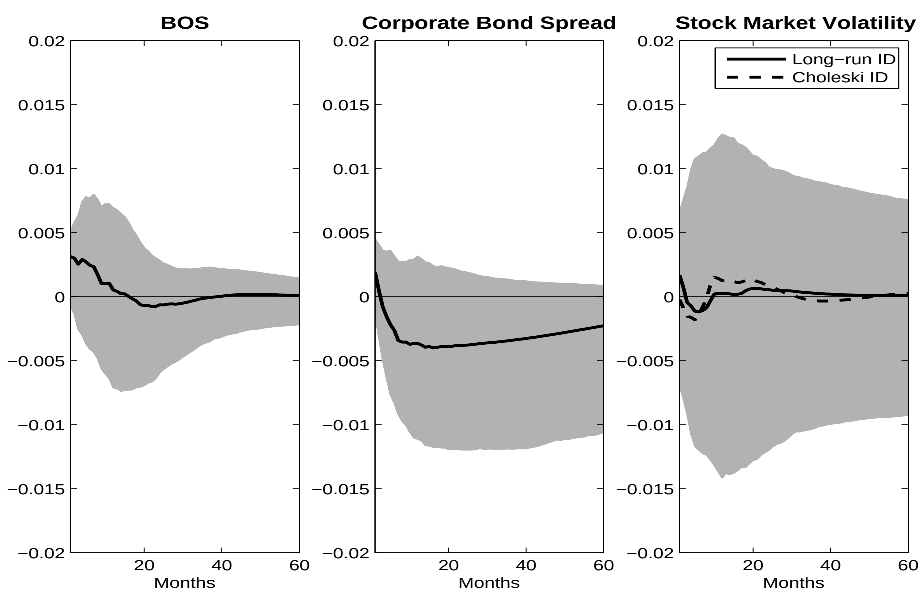

Notes: see notes to Figure 1.7. The first panel is simply a replication of the 'Activity to Uncertainty' impulse response from Figure 1.7. The second panel displays the 'Activity to Uncertainty' response of a three-variable Blanchard-Quah-type SVAR with 'Corporate Bond Spread' as the uncertainty measure, total industrial production as the activity measure and the civilian unemployment rate. 'Corporate Bond Spread' refers to the spread of the 30 year Baa corporate bond index over the 30 year treasury bond. Where the 30 year treasury bond was missing we used the 20 year bond. Data source for the bond data is the Federal Reserve Board. The third panel displays the 'Activity to Uncertainty' response of the same SVAR with the stock market volatility dummy from Bloom (2009) as the uncertainty measure. The Choleski-identified impulse response (dashed line) from Bloom (2009) is included for comparison.

Figure 1.10 compares the effects of surprise movements in various uncertainty proxies on production. The leftmost panel is simply a replication of the result in Figure 1.7, i.e. where we use the survey-based General Conditions-$Uncertainty_t^{BOS}$-index as the uncertainty measure. The middle panel uses the corporate bond spread and the rightmost panel the "high stock market volatility"-dummy from Bloom (2009). Note that the high-frequency "wait-and-see"-dynamics with a fast rebound more or less survive the long-run idenfication strategy, as far as the point estimate is concerned. This is not too surprising given that the Choleski-identified impulse response is essentially zero in the long-run. But Figure 1.10 also shows that any high-frequency impact of surprise movements in uncertainty, regardless of how it is measured, is likely to be small – much less than half a percent of production – and statistically indistinguishable from each other as well as from zero.

Figure 1.11 compares the reaction of various uncertainty indices to an adverse long-run shock. The leftmost panel is again a replication of the result in Figure 1.7. The point estimates for all three uncertainty measures are positive, signficantly so for General Conditions-$Uncertainty_t^{BOS}$ and the corporate bond spread, which is at least suggestive of the "by-product"-hypothesis.

Figure 1.11: Long-Run Shock on Uncertainty

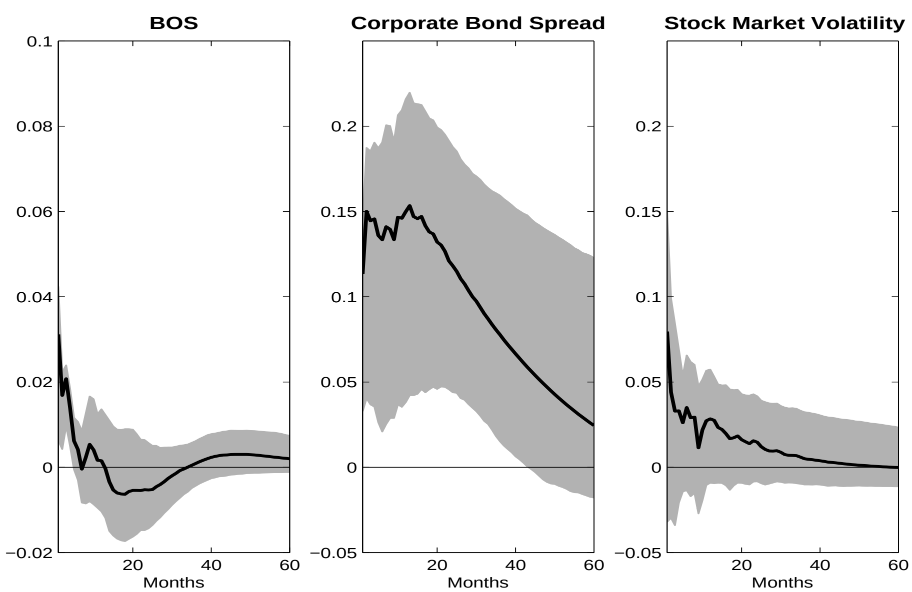

Notes: see notes to Figures 1.7 and 1.10. The first is a replication of the 'Uncertainty to Long-Run' impulse response from Figure 1.7. The second panel displays the 'Uncertainty to Long-Run' response of a three-variable Blanchard-Quah-type SVAR with 'Corporate Bond Spread' as the uncertainty measure, total industrial production as the activity measure and the civilian unemployment rate. The third panel displays the 'Uncertainty to Long-Run' response of the same SVAR with stock market volatility from Bloom (2009) as the uncertainty measure. Until 1986 this is realized monthly stock return volatility, and thereafter an implied volatility index.

As has been mentioned before our results leave open two interpretations for the role of uncertainty in economic fluctuations. The first interpretation is that uncertainty is an autonomous source of such fluctuations but has mainly long-run effects. In this case our SVARs show that structural models need a mechanism that transmits rather transitory uncertainty shocks into very persistent or even permanent output and employment declines. Alternatively, uncertainty can be viewed as an epiphenomenon that accompanies bad economic times. While we cannot strictly rule out the former, we believe that the data points in the latter direction: bad times breed uncertainty.

Table 1.5: Relation Between NBER Recessions and High Uncertainty Dates

Uncertainty Measure	High Uncertainty In Recessions	High Uncertainty Outside of Recessions
$Uncertainty^{BOS}$	7 out of 7	8.5%
Corporate Bond Spread	6 out of 8	11.2%
Stock Market Volatility	7 out of 7	8.3%

Notes: $Uncertainty^{BOS}$ refers to the BOS uncertainty measure, based on Q 1. For 'Corporate Bond Spread' see notes to Figure 1.10. For 'Stock Market Volatility' see notes to Figure 1.11. For each uncertainty proxy we construct a high uncertainty dummy, setting it unity, when the value exceeds the time series average by one standard deviation. In the first column we report how many post 1960 recessions coincide with high uncertainty events. We do not have BOS or stock market volatility data available for the 1961 recession. There are no high corporate bond spread-uncertainty events during the 1961 and the 1991 recessions. In the second column we report the fraction of months where high uncertainty events occur outside of NBER recessions.

Table 1.5 shows that almost all NBER-dated recessions were periods of high uncertainty – whether it is measured as cross-sectional forecast dispersion from business survey data, the corporate bond spread or stock market volatility. We define high uncertainty events as months when either uncertainty measure was one standard deviation above its time series average. That almost all US recessions have been times of high uncertainty is consistent with causality running in either direction – from uncertainty to economic activity or from activity to uncertainty. It is therefore interesting to note that there is a considerable fraction of months – close to 10 percent – where uncertainty spikes but the economy was not in a recession, nor did a period of economic distress soon follow. This is particularly true in the mid-late 1980s (around the time of the 1987 stock market crash) and the mid-late 1990s, well before the downturn of 2001. That such large increases in uncertainty did not lead to economic contractions is at least suggestive evidence that uncertainty is a concomitant factor of bad economic times rather than a causal factor for them.

It is beyond the scope of this paper to fully specify a model of intrinsic uncertainty as an endogenous result of bad first moment shocks. Bachmann and Moscarini (2011) do so using price experimentation as a mechanism; Fostel and Geanakoplos (2011) point to leverage. More generally, we think of recessions as times of severed business and customer relationships and of failing business models. Business and customer relationships have to be reestablished and business models altered when the economy is at trough. This generates uncertainty. In booms, in contrast, businesses have little incentive (or opportunity) to substantially change their operating practices. Customers stay with their preferred business.

As a highly stylized example, suppose there are three businesses in an economy each producing the same product, with total demand equal to 2 units of the product. Suppose initially that all three businesses have an equal share of two-thirds. In a boom demand becomes 2.5. With costs to establishing new business relationships, the customers of each business stick around and demand more. There is no uncertainty. In a recession demand becomes $2x$, where $x<1$. Assume that one of the businesses goes under and business relations are severed. The existing customers at the two remaining businesses now demand $\frac{2}{3}x$ each. What happens to the customers whose business partner vanished? Let us assume there is some uncertainty over where they are going to go, as in a location model where businesses do not know the spatial distribution of customers. On the one extreme, the allocation might be $\left[\frac{4}{3}x, \frac{2}{3}x\right]$, i.e. one business gets all the free customers, on the other extreme it might be an equal split: $[x, x]$. It is obvious that there exists a range for x, namely $\left(\frac{1}{2}, \frac{2}{3}\right)$, where even in the most equal distribution both businesses are worse off than before, but with an unequal split one business might even come out better than before in this recession. The important point is this: there is an intrinsic uncertainty due to recessions, because business structures and practices have to be re-arranged.

1.5 Final Remarks

Using two different measures of business uncertainty from high-frequency, sectoral business surveys in Choleski-identified structural vector autoregressions we find that positive innovations to business uncertainty have protracted negative implications for sectoral economic activity. This appears to be inconsistent with a high-frequency "wait-and-see"-channel being the dominant effect of suprise movements in business uncertainty. This contrasts with the results in the literature for suprise movements in stock market volatility, which trigger short-run collapses of activity and quick rebounds. We confirm this result from Bloom (2009) also in a long-run identification strategy.

This can mean two things, which are not necessarily mutually exclusive. On the one hand, perhaps stock market volatility really measures a different type of uncertainty than survey-based uncertainty and the corporate bond spread – say aggregate uncertainty versus idiosyncratic uncertainty – and these types of uncertainty have different impacts on businesses' behavior. The second possibility is that the low-frequency impact of the survey-based uncertainty measures swamps the high-frequency "wait-and-see"-dynamics. However, we show in this paper that any high-frequency impact of surprise movements in uncertainty is likely to be small, regardless of how uncertainty is measured and how its high-frequency impact is identified. This leaves open the possibility that "wait-and-see"-dynamics can be combined with an endogenous growth mechanism – R&D investment or embodied technological change – to generate the observed protracted negative implications for economic activity in Choleski-identified structural vector autoregressions. Finally, structural vector autoregression studies, by their very nature, can only make statements about the *average* effect of uncertainty shocks, which leaves open the possibility that high uncertainty events in certain episodes may have severely adverse high-frequency consequences.

But this paper also opens up another possibility, the "by-product"-hypothesis, for which we find evidence both in Choleski-identified as well as Blanchard-Quah-type structural vector autoregressions. Under this interpretation negative long-run shocks lead to high uncertainty events. Uncertainty is a concomitant phenomenon of negative first moment events in the economy. Bad times breed uncertainty. Of course, this leaves open the possibility that uncertainty and the resulting "wait-and-see" are an important *propagation and amplification* mechanism for other shocks. Businesses may invest and hire less when the outlook is bleak, but they may even be more reluctant to invest and hire when, in addition, the outlook is uncertain.

Our results suggest that research in the following three areas may prove fruitful: "wait-and-see"-mechanisms in endogenous growth environments; fully specified mechanisms that endogenously generate uncertainty in bad economic times; and the role of uncertainty as a propagation and amplification mechanism.

Acknowledgements

I am indebted to Rüdiger Bachmann and Eric Sims who are co-authors of Chapter 1. We thank seminar participants at RWTH Aachen University, Bundesbank, IAB Nuremberg, Université de la Méditerranée Aix-Marseille II, University of Michigan, 2010 Midwest Macro Meetings (Lansing), 2010 NBER-SI-ME, NYU Stern, Rochester, 2010 SED meeting (Montreal), 2010 World Congress of the Econometric Society in Shanghai, Yale and ZEW

Mannheim as well as Robert Barsky, Eduardo Engel, and Giuseppe Moscarini for their comments. We are grateful to Kai Carstensen and Sigrid Stallhofer from the IFO Institute as well as Holly Wade from the NFIB for providing us with their data and introducing us to the institutional backgrounds.

Bibliography

Alexopoulos, M., and J. Cohen (2009): “Uncertain Times, Uncertain Measures,” *mimeo.*

Arellano, C., Y. Bai, and P. Kehoe (2011): “Financial Markets and Fluctuations in Uncertainty,” *mimeo.*

Bachmann, R., and C. Bayer (2011): “Uncertainty Business Cycles - Really?,” Working Paper 16862, National Bureau of Economic Research.

Bachmann, R., and G. Moscarini (2011): “Business Cycles and Endogenous Uncertainty,” *mimeo.*

Basu, S., and B. Bundick (2011): “Uncertainty Shocks in a Model of Effective Demand,” *mimeo.*

Becker, S., and K. Wohlrabe (2008): “Micro Data at the Ifo Institute for Economic Research - The “Ifo Business Survey’ Usage and Access,” *Journal of Applied Social Science Studies*, 128(2), 307–319.

Bernanke, B. (1983): “Irreversibility, Uncertainty, and Cyclical Investment,” *The Quarterly Journal of Economics*, 98(1), 85–106.

Blanchard, O., and D. Quah (1989): “The Dynamic Effects of Aggregate Demand and Aggregate Supply Shocks,” *American Economic Review*, 79(4), 655–673.

Bloom, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77(3), 623–685.

Bloom, N., M. Floetotto, and N. Jaimovich (2010): “Really Uncertain Business Cycles,” *mimeo.*

Bomberger, W. (1996): “Disagreement as a Measure of Uncertainty,” *Journal of Money, Credit and Banking*, 28(3), 381–392.

Bond, S., and J. Cummins (2004): "Uncertainty and Investment: an Empirical Investigation Using Data on Analysts' Profits Forecasts," *mimeo*.

Bontempi, M., R. Golinelli, and G. Parigi (2010): "Why demand uncertainty curbs investment: Evidence from a panel of Italian manufacturing firms," *Journal of Macroeconomics*, 32(1), 218–238.

Christiano, L., R. Motto, and M. Rostagno (2010): "Financial Factors in Economic Fluctuations," *ECB WP 1192*.

Chugh, S. K. (2011): "Firm Risk and Leverage-Based Business Cycles," *mimeo*.

Clements, M. P. (2008): "Consensus and uncertainty: Using forecast probabilities of output declines," *International Journal of Forecasting*, 24, 76–86.

Dixit, A. K., and R. S. Pindyck (1994): *Investment under Uncertainty*. Princeton University Press.

Dunkelberg, W. C., and H. Wade (2009): *NFIB Small Business Economic Trends*. October.

Federer, P. J. (1993): "Does Uncertainty Affect Investment Spending?," *Journal of Post Keynesian Economics*, 16(1), 19–35.

Fernandez-Villaverde, J., P. Guerron-Quintana, J. Rubio-Ramirez, and M. Uribe (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review (forthcoming)*.

Fostel, A., and J. Geanakoplos (2011): "Why Does Bad News Increase Volatility and Decrease Leverage?," *Journal of Economic Theory (forthcoming)*.

Fuss, C., and P. Vermeulen (2008): "Firms' Investment decisions in response to demand," *Applied Economics*, 40(18), 2337–2351.

Gali, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Schocks Explain Aggregate Fluctuations?," *American Economic Review*, 89(1), 249–271.

Gilchrist, S., J. Sim, and E. Zakrajsek (2010): "Uncertainty, Financial Frictions and Investment Dynamics," *mimeo*.

Gilchrist, S., V. Yankov, and E. Zakrajsek (2009): "Credit Market Shocks and Economic Fluctuations: Evidence from Corporate Bond and Stock Markets," *Journal of Monetary Economics*, 56(4), 471–493.

Giordano, P., and P. Soederlind (2003): "Inflation Forecast Uncertainty," *European Economic Review*, 47(6), 1037–1059.

Guiso, L., and G. Parigi (1999): "Investment and Demand Uncertainty," *The Quarterly Journal of Economics*, 114(1), 185–227.

Hamilton, J., and G. Lin (1996): "Stock Market Volatility and the Business Cycle," *Journal of Applied Econometrics*, 11(5), 573–593.

Kilian, L. (1998): "Small Sample Confidence Intervals for Impulse Response Functions," *Review of Economics and Statistics*, 80(2), 218–230.

Leahy, J., and T. Whited (1996): "The effect of uncertainty on investment: some stylized facts," *Journal of Money, Credit and Banking*, 28(1), 64–83.

Popescu, A., and F. R. Smets (2010): "Uncertainty, Risk-Taking and the Business Cycle in Germany," *CESifo Economic Studies*, 56(4), 596–626.

Schaal, E. (2010): "Uncertainty, Productivity and Unemployment in the Great Recession," *mimeo*.

Shapiro, M. D., and M. Watson (1988): "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, 3, 111–156.

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

Zarnowitz, V., and L. A. Lambros (1987): "Consensus and Uncertainty in Economic Prediction," *The Journal of Political Economy*, 95(3), 591–621.

Appendix

A1 A Simple Model

To illustrate the relationship between concepts such as disagreement, uncertainty and cross-sectional variance, we use the following simple two-period model: tomorrow's business situation of firms is unknown today. It can move into three directions. Business situations can improve (+1), stay the same (0) or deteriorate (-1). For each firm, nature draws the change in business situation from the following probability distribution: $[0.5 * (1-p), p, 0.5 * (1-p)]$, which is assumed to be known to the firms. The cross-sectional variance of the future business situation is obviously $(1-p)$, a decreasing function of p. Furthermore, we assume that businesses receive a signal about the change in their business situation, with a structure illustrated in Table A1. For instance, if tomorrow's true state is +1, the signal can be +1 (with probability q) and 0 with probability $(1-q)$. q thus measures the informativeness of the signal.

Table A1: A Simple Two-Period Model of Firms' Business Situations

				State Tomorrow				
	$0.5*(1-p) \swarrow$			$\downarrow p$			$\searrow (1-p)*0.5$	
	+1			0			-1	
$q \swarrow$		$\searrow (1-q)$	$0.5*(1-q) \swarrow$	$q \downarrow$	$\searrow (1-q)*0.5$	$(1-q) \swarrow$		$\searrow q$
+1		0	+1	0	-1	0		-1
				Signal				

Using Bayes' Law we can compute the probabilities of the true state, conditional on a signal:

1. (a) $Prob(state = 1|signal = 1) = \frac{q*0.5*(1-p)}{q*0.5*(1-p)+0.5*(1-q)*p}$
   (b) $Prob(state = 0|signal = 1) = \frac{0.5*(1-q)*p}{q*0.5*(1-p)+0.5*(1-q)*p}$
   (c) $Prob(state = -1|signal = 1) = 0$
2. (a) $Prob(state = 1|signal = 0) = \frac{(1-q)*0.5*(1-p)}{(1-q)*0.5*(1-p)+q*p+(1-q)*0.5*(1-p)}$
   (b) $Prob(state = 0|signal = 0) = \frac{q*p}{(1-q)*0.5*(1-p)+q*p+(1-q)*0.5*(1-p)}$
   (c) $Prob(state = -1|signal = 0) = \frac{(1-q)*0.5*(1-p)}{(1-q)*0.5*(1-p)+q*p+(1-q)*0.5*(1-p)}$
3. (a) $Prob(state = 1|signal = -1) = 0$

(b) $Prob(state = 0|signal = -1) = \frac{0.5*(1-q)*p}{q*0.5*(1-p)+0.5*(1-q)*p}$

(c) $Prob(state = -1|signal = -1) = \frac{q*0.5*(1-p)}{q*0.5*(1-p)+0.5*(1-q)*p}$

From these conditional probabilities, conditional expectations and variances can be computed. And these, in turn, allow us to calculate 1) the variance of the conditional expectations over the change in business situations, which is a measure of disagreement; and 2) the average conditional variance over the change in the business situation of a firm, which is a measure of the average (subjective) uncertainty in the population of firms.

We begin with the case of perfectly informative signals: $q = 1$. In this case, obviously, survey disagreement moves one for one with the variance of tomorrow's state, but firms do not experience any subjective uncertainty about the change in their business situation. With $q = 1$ and in a two period set up disagreement and uncertainty do not comove. The fact that we find substantial forecast errors in the IFO-BCS suggests that this extreme case may not be realistic. But even if we assumed $q = 1$ and thus certainty for the immediate future, higher disagreement today indicates a higher cross-sectional variance in business situations tomorrow and thus higher uncertainty about business situations for periods beyond the immediate future, as long as the variance of future innovations to the business situation of firms has some persistence beyond the immediate period and signals are not perfectly informative about this farther future. Figure A1 plots the autocorrelograms for General Conditions-$Uncertainty_t^{BOS}$, Shipments-$Uncertainty_t^{BOS}$, Production-$Uncertainty_t^{IFO}$ and Production-$Uncertainty_t^{fe^{IFO}}$, showing that uncertainty is very persistent.

Figure A1: Autocorrelograms of Various Uncertainty Measures

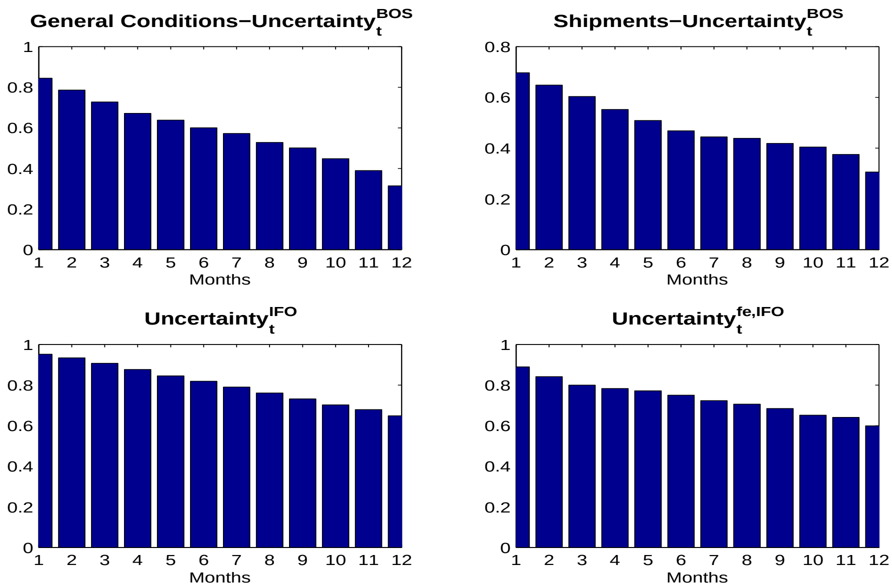

Notes: General Conditions-$Uncertainty_t^{BOS}$ is based on Q 1. Shipments-$Uncertainty_t^{BOS}$ is based on Q 6. Production-$Uncertainty_t^{IFO}$ is based on Q 4. For the construction of Production-$Uncertainty_t^{fe^{IFO}}$, based on Q 4 and Q 5, see Section 1.3.3.

Next, we look at the cases with imperfectly informative signals, i.e. $q<1$. We know from the conditional variance decomposition formula that if the variance of tomorrow's state increases either the variance of the conditional expectations over tomorrow's state (disagreement) or the average conditional variance over tomorrow's state (average subjective uncertainty) has to increase, both may increase. The following Figure A2 shows for various levels of the signal precision, q, that the latter is indeed the case in this model. The actual cross-sectional variance of tomorrow's state is given by the black solid line, the variance of the conditional expectations over tomorrow's state (disagreement) by the blue dashed line and the average conditional variance over tomorrow's state (subjective uncertainty) by the red dotted line.

Finally, in order to translate the continuous disagreement measure – the variance of the conditional expectations over the change in business situations – into discrete disagreement in survey answers, where only $[-1,0,1]$ as an answer are possible, we assume that if the firm receives zero as a signal, it will answer zero, simply because the conditional expectation is zero in this case (by the symmetry of the model). Furthermore, if it receives a signal equal to 1, the probability of answering 1 in the survey equals the expectation conditional on the signal being 1, which ranges from 1 (if $p=0$) to 0 (if $p=1$). This conditional expectation, $E[state|signal=1]$, is computed from the conditional probabilities above. This means, the

Figure A2: Cross-sectional Variance, Disagreement and Uncertainty

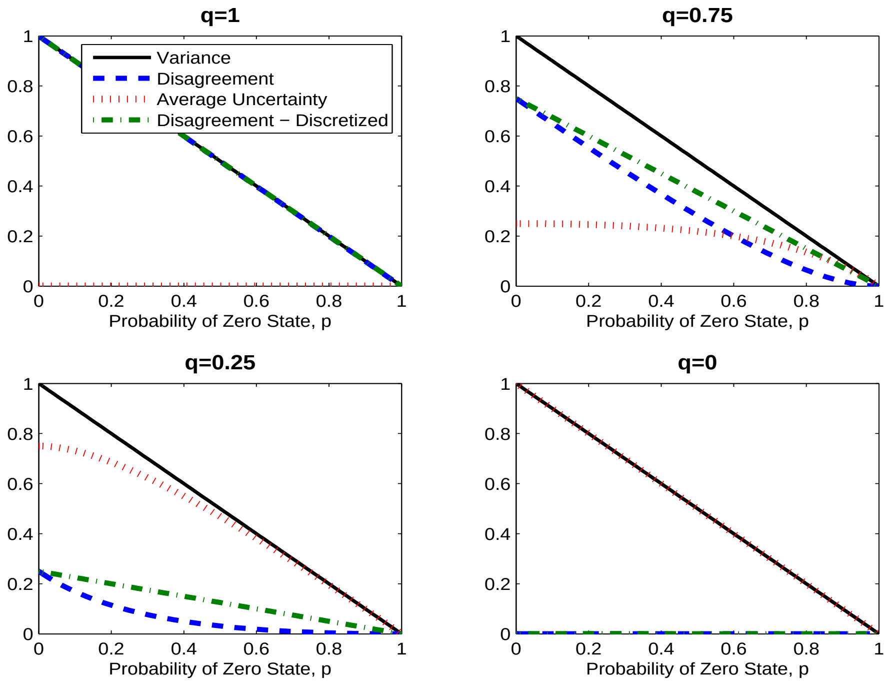

closer the conditional expectation is to unity, the more likely firms are going to respond with 1 in the survey. Symmetrically for the case of receiving a signal that equals -1. With these assumptions, the variance of the survey answers is given by ($E[answer]$ is computed analogously):

$$\begin{aligned} VAR[answer] = \quad & (1 - E[answer])^2 E[state|signal = 1] * Prob(signal = 1) + \\ & (0 - E[answer])^2 (1 - E[state|signal = 1]) * Prob(signal = 1) + \\ & (0 - E[answer])^2 Prob(signal = 0) + \\ & (0 - E[answer])^2 (1 - E[state|signal = -1]) * Prob(signal = -1) + \\ & (-1 - E[answer])^2 (E[state|signal = -1]) * Prob(signal = -1) \end{aligned}$$

This discretized version of disagreement is also shown in Figure A2, by the green dashed-dotted line. It follows closely the continuous disagreement measure. Notice that for intermediate signal qualities, both disagreement and uncertainty move in the same direction as the variance of the state tomorrow. In particular, for high values of p subjective uncertainty varies significantly with the cross-sectional variance of the change in business situations. If the signal was such that it left everybody with the same conditional expectation ($q = 0$), then of course disagreement would always be zero. Only the subjective uncertainty would then be affected.

A2 Third FED District Business Outlook Survey (BOS)

A2.1 Additional BOS Questions

Q 6 *"Company Business Indicators: Shipments six months from now vs. [CURRENT MONTH]: decrease, no change, increase?"*

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

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

Q 9 *"Company Business Indicators: Shipments [LAST MONTH] vs. [CURRENT MONTH]: decrease, no change, increase?"*

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

A2.2 Additional Information on BOS Variables

Table A2: Correlation between BOS-$Activity_t$ Variables and Official Statistics

	General Conditions	Shipments	Employment
BLS Monthly Sect. & Regio. Empl.	0.54	0.60	0.63
Philadelphia FED Coincident Index	0.71	0.68	0.60
NIPA Yearly Sect. & Regio. Prod.	0.39	0.41	-

Notes: This table displays the unconditional contemporaneous correlations of BOS-$Activity_t$ Variables, based, in column order, on Q 2, Q 9 and Q 3, with log-differences of three different measures of sectoral and regional activity measures from official statistics (in row order): 'BLS Monthly Sect. & Regio. Empl.' refers to the sum of the seasonally adjusted monthly manufacturing employment series for Pennsylvania, Delaware and New Jersey, available from the BLS from 1990 on. 'Philadelphia FED Coincident Index' refers to the GDP-weighted sum of the Philadelphia FED Coincident Indices for Pennsylvania, Delaware and New Jersey (notice that this index is regionally, but not sectorally coinciding with the coverage of the BOS). It is available from 1979 on. 'NIPA Yearly Sect. & Regio. Prod.' refers to the GDP-weighted sum of the yearly NIPA quantity indices for the manufacturing sector for Pennsylvania, Delaware and New Jersey.

A2.3 Additional BOS Results

This appendix provides various robustness checks to the results in Section 1.4.1. Figure A3 shows that the ordering between uncertainty and activity variables is irrelevant for the result that uncertainty innovations in two-variable SVARs trigger prolonged declines in sectoral economic activity. Figures A4 and A5 vary the economic activity variable used in our baseline two-variable SVAR, while keeping General Conditions-$Uncertainty_t^{BOS}$ (based on Q 1) as the uncertainty measure: the BOS shipments, employment and workhours based activity indices, and labor productivity. Figures A6 to A8, in turn, vary the uncertainty measure used: an indicator variable for high uncertainty, an entropy-based uncertainty measure and uncertainty measures derived from other expectation questions in the BOS.

Figure A3: Uncertainty Innovation on Manufacturing Production - Reverse Ordering

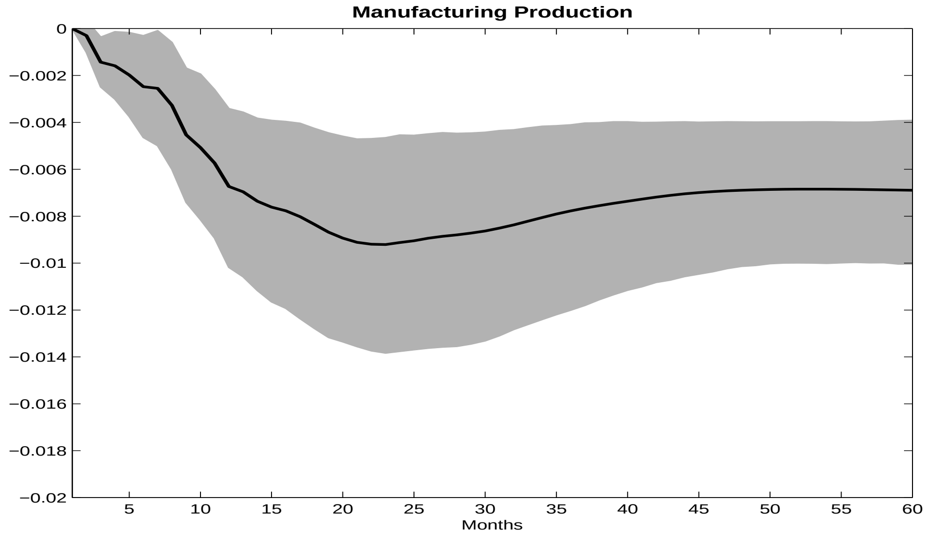

Notes: The IRF is based on a two-variable SVAR with General Conditions-$Uncertainty_t^{BOS}$ (based on Q 1 of the BOS) ordered second and 12 lags. Manufacturing production is the natural logarithm of the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators.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

Figure A4: Uncertainty Innovations on Various BOS Activity Indices

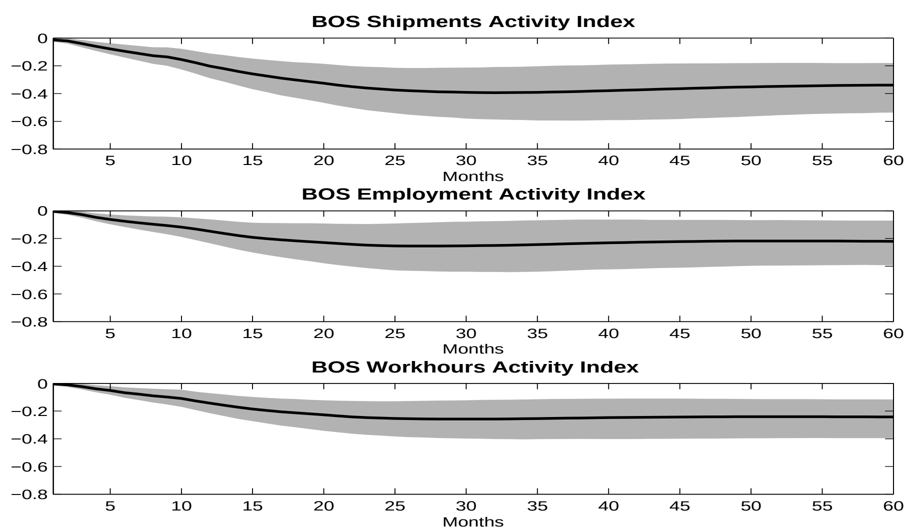

Notes: see notes to Figure A3. Uncertainty is ordered first. The activity indices for the three panels are based on Q 9, Q 3 and Q 10, respectively.

Figure A5: Uncertainty Innovation on Manufacturing Labor Productivity

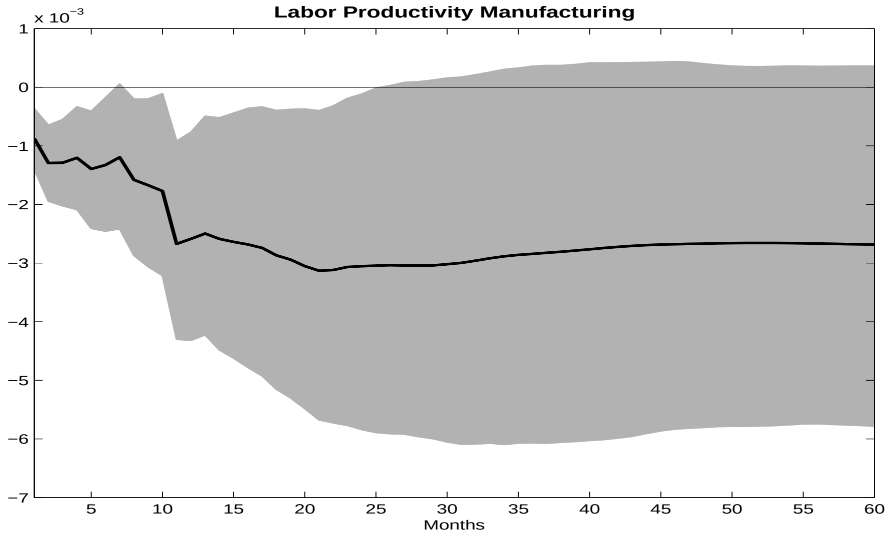

Notes: see notes to Figure A3. Uncertainty is ordered first. Labor productivity is the log-difference between the (seasonally adjusted) monthly manufacturing production index from the OECD main economic indicators and the (seasonally adjusted) monthly manufacturing total hours series, which is itself based on the manufacturing employment and weekly hours per worker series from the BLS-CES data base.

Figure A6: Uncertainty Innovation (Indicator Variable) on Manufacturing Production

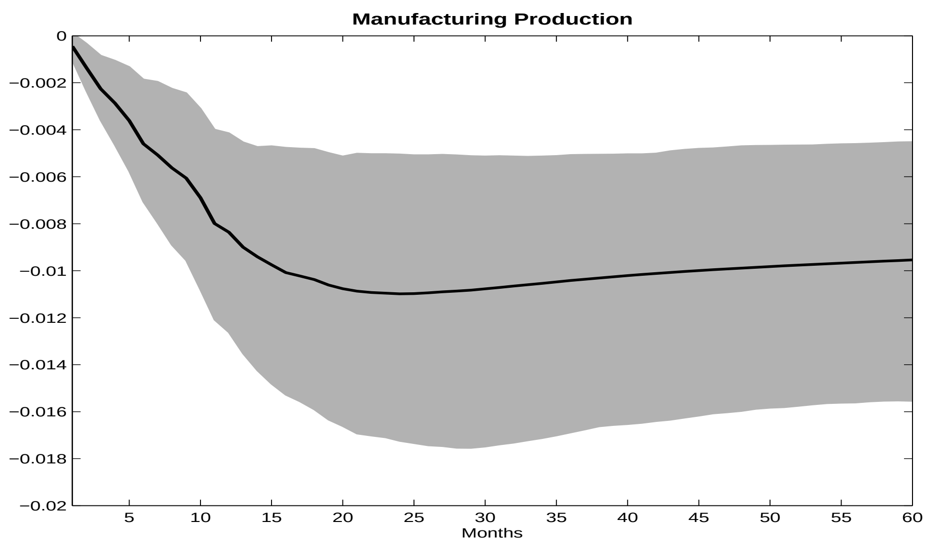

Notes: see notes to Figure A3. Uncertainty is ordered first. The uncertainty variable here is an indicator variable that takes on a value of one, if General Conditions-$Uncertainty_t^{BOS}$, the measure of uncertainty which is based on Q 1, is one standard deviation above its mean, and zero otherwise. There are 60 high-uncertainty observations, or about 12% of the sample.

Figure A7: Uncertainty Innovation on Manufacturing Production - Entropy

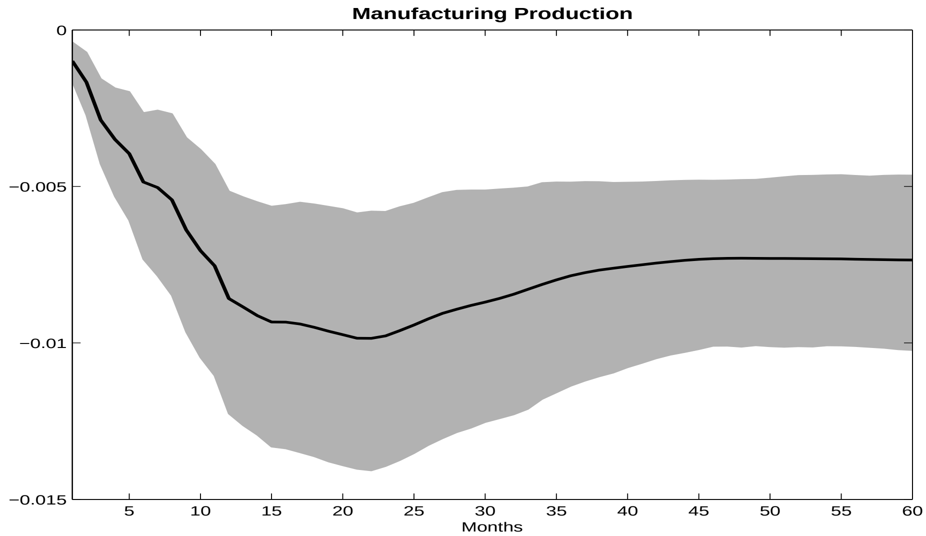

Notes: see notes to Figure A3. Uncertainty is ordered first. It is measured as $Uncertainty_t^{Entrop} \equiv Frac_t(\text{Increase})\log(1/Frac_t(\text{Increase})) + Frac_t(\text{Decrease})\log(1/Frac_t(\text{Decrease})) + Frac_t(\text{Neutral})\log(1/Frac_t(\text{Neutral}))$.

Figure A8: Uncertainty Innovations from Other BOS Activity Indices

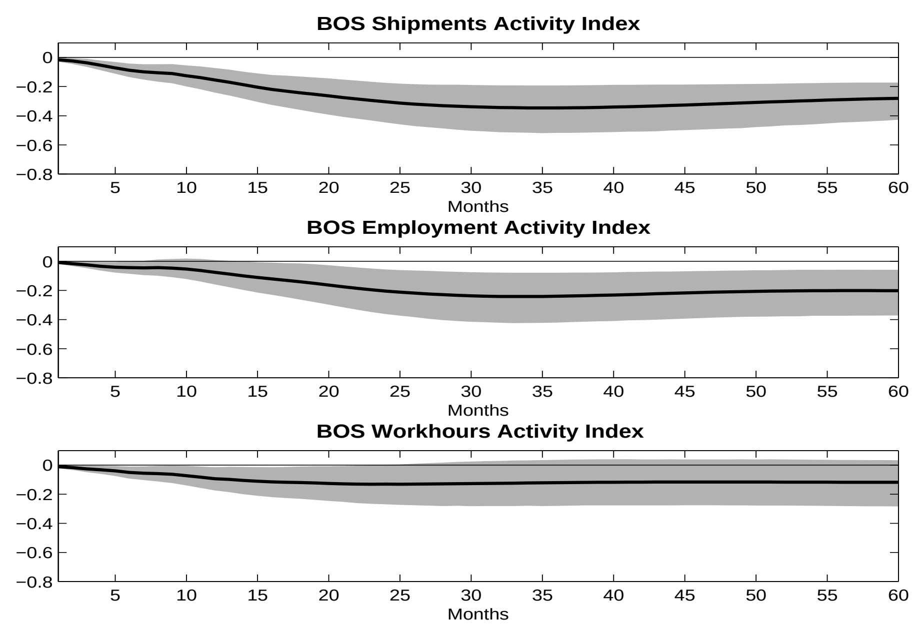

Notes: see notes to Figure A3. The uncertainty variables for the three panels are based on Q 6, Q 7 and Q 8, respectively. The activity indices for the three panels are based on Q 9, Q 3 and Q 10. Uncertainty is ordered first.

A3 IFO Business Climate Survey (IFO-BCS)

A3.1 Original German IFO-BCS Questions

Q 11 *"Erwartungen für die nächsten 3 Monate: Unsere inländische Produktionstätigkeit – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY wird voraussichtlich: steigen, etwa gleich bleiben, abnehmen."*

Q 12 *"Tendenzen im vorangegangenen Monat: Unsere inländische Produktionstätigkeit – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY ist: gestiegen, etwa gleich geblieben, gesunken."*

A3.2 Variance Decomposition of $(Uncertainty_t^{IFO})^2$

Figure A9: Variance Decomposition of $(Uncertainty_t^{IFO})^2$

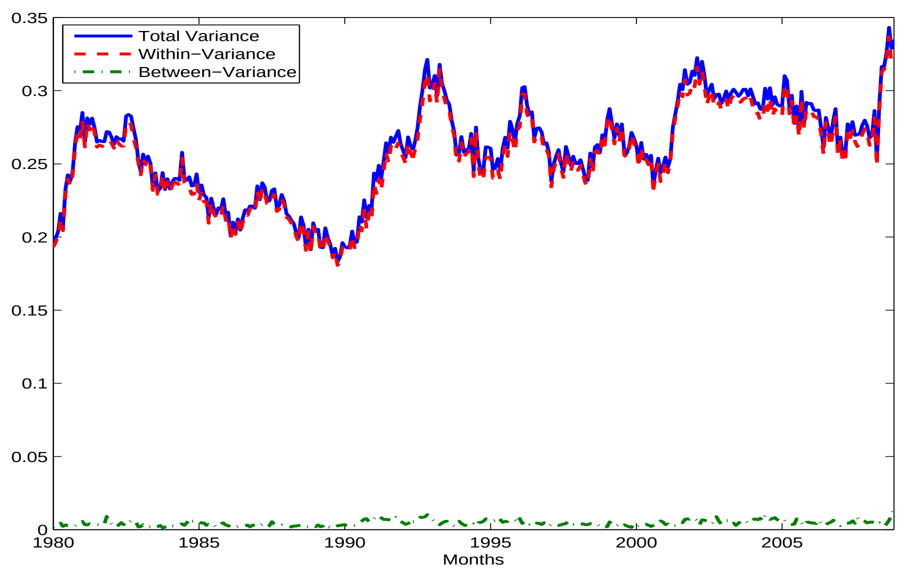

Notes: 'Total Variance' refers to $(Uncertainty_t^{IFO})^2$. 'Within-Variance' is the cross-sectional average of the industry analogs of $(Uncertainty_t^{IFO})^2$ for the following 13 manufacturing industries: transportation equipment (*Fahrzeugbau*), machinery and equipment (*Maschinenbau*), metal products (*Metallerzeugung*), other non-metallic mineral products (*Glas, Keramik, Verarbeitung von Steinen und Erden*), rubber and plastic products (*Gummi und Kunststoff*), chemical products (*Chemische Industrie*), electrical and optical equipment (*Elektrotechnik, Feinmechanik und Optik*), pulp, paper, publishing and printing (*Papier, Verlage, Druck*), furniture and jewelery (*Möbel und Schmuck*), cork and wood products except furniture (*Holz ohne Möbel*), leather (*Leder*), textiles and textile products (*Textil und Bekleidung*), food, beverages and tobacco (*Ernährung und Tabak*). We leave out the oil industry, because it has only very few observations. 'Between-Variance' refers to the cross-sectional variance of the industry-specific $Frac_t^+ - Frac_t^-$-indicators.

A3.3 Construction of $Uncertainty_t^{fe^{IFO}}$

In this section we describe the construction of the $Uncertainty_t^{fe^{IFO}}$-index. To fix ideas, we proceed at first as if the production expectation question in IFO-BCS, Q 4, was asked only for the next month instead of the following three months. In this case, when comparing the expectation in month t with the realization in month $t+1$, nine possibilities arise: the company could have predicted an increase in production and realized one, in which case we would count this as zero forecast error. It could have realized a no change, in which case, we would quantify the expectation error as -1 and, finally, it could have realized a decrease, which counts as -2.

Table A3: Possible Expectation Errors - One Month Case

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

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

Table A3 summarizes the possible expectation errors. Of course, the production expectation question in IFO-BCS is for three months ahead. Suppose that a firm stated in month t that its production will increase in the next three months. Suppose that in the next three months one observes the following sequence of outcomes: production increased in $t+1$, remained unchanged in $t+2$ and finally decreased in $t+3$. Due to the qualitative nature of the IFO-BCS we have to make some assumptions about the definition of the expectation error at the micro level. As a baseline we adopt the following steps. First, we define for every month t a firm-specific activity variable over the next three months, $t+3$, by the sum of the *Increase*-instances minus the sum of the *Decrease*-instances over that time period.[14] Denote this variable by $REALIZ_t$. It can obviously range from $[-3, 3]$. Then the expectation errors are computed as:

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

Table A4: Possible Expectation Errors - Three Month Case

		$Experror_{i,t}$
Expected $Increase_{t-3}$	$REALIZ_{i,t}>0$	0
Expected $Increase_{t-3}$	$REALIZ_{i,t} \leq 0$	$(REALIZ_{i,t} - 1)$
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}>0$	$REALIZ_{i,t}$
Expected $Unchanged_{t-3}$	$REALIZ_{i,t} = 0$	0
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}<0$	$REALIZ_{i,t}$
Expected $Decrease_{t-3}$	$REALIZ_{i,t}<0$	0
Expected $Decrease_{t-3}$	$REALIZ_{i,t} \geq 0$	$(REALIZ_{i,t} + 1)$

Notes: Rows refer to the qualitative production expectations in IFO-BCS in month t (Q 4).

Notice that the procedure in Table 4.1 is analogous to the one month case. Dividing by three is simply a normalization. $Expectationerror_{t+3}$ ranges from $[-\frac{4}{3}, \frac{4}{3}]$, where for instance $-\frac{4}{3}$ indicates a strongly negative forecast error: the company expected production to increase over the next three months, yet every single subsequent month production actually declined.

Computing the cross-sectional standard deviations of the expectation errors at each month, t, gives us a qualitative series of forecast error standard deviations. Specifically:

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

Notice the timing in the definition of $Uncertainty_t^{fe}$, which is the same as in Bloom (2009) for stock market volatility: the standard deviation of *realized* expectation errors in $t+3$ does not constitute uncertainty in $t+3$. It is the knowledge (at time t) of this standard deviation going up or down that makes decision makers more or less uncertain at time t. It should be emphasized that this timing does not require decision makers to know anything about the future, other than that it is more or less uncertain. Figure A10 depicts $Uncertainty_t^{fe^{IFO}}$ and $Uncertainty_t^{IFO}$, both at the monthly and the quarterly frequency, and shows that they strongly comove.

Figure A10: Comparison of $Uncertainty_t^{IFO}$ and $Uncertainty_t^{fe^{IFO}}$

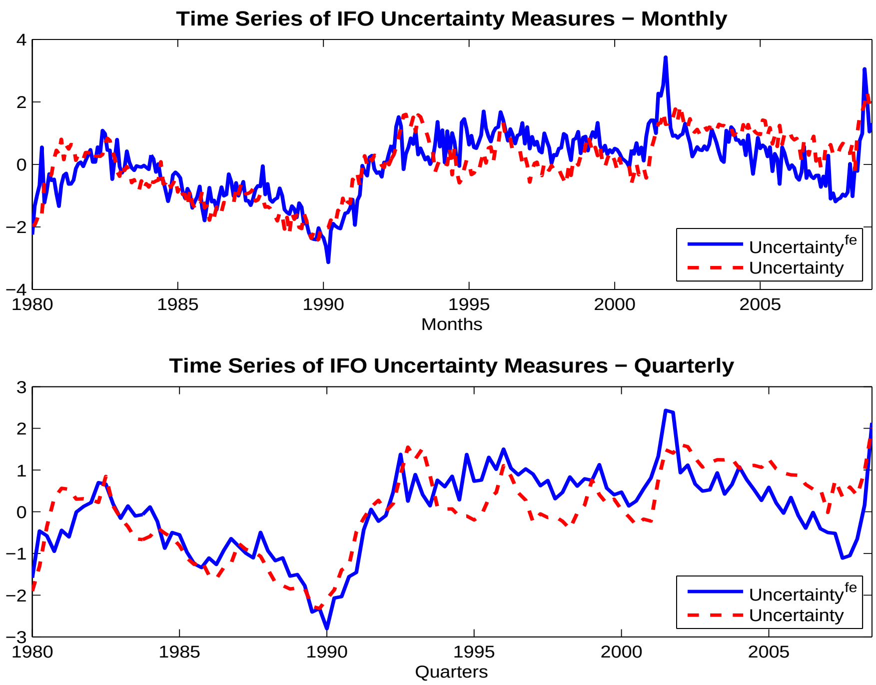

Notes: The upper panel shows the monthly time series of $Uncertainty_t^{IFO}$ and $Uncertainty_t^{fe^{IFO}}$, demeaned and standardized by their standard deviation. Their correlation is 0.73. The lower panel shows the quarterly averages of the monthly $Uncertainty_t^{IFO}$ and $Uncertainty_t^{fe^{IFO}}$ time series, demeaned and standardized by their standard deviation. Their correlation is 0.77.

A4 Small Business Economics Trends Survey (SBETS)

The Small Business Economic Trends Survey (SBETS) is a monthly survey conducted by the National Foundation of Independent Businesses (NFIB) which focuses on small companies across the U.S. and across all sectors. Thus the SBETS is a good complement to the BOS which focuses on larger manufacturing firms in the Third FED District. To the extent that the SVAR results are similar this section lends additional support to our findings. The SBETS's monthly part starts in 1986. The survey on a quarterly basis is available since the mid 1970s. We prefer the highest possible frequency to give the "wait-and-see"-dynamics the best possible chance to appear in the data. None of our results depend on that choice of frequency. In terms of participation, the October 2009 issue of the SBETS (see Dunkelberg and Wade, 2009, for more detailed information) reports that from January 2004 to December 2006 roughly 500 business owners responded, and that the number has subsequently increased to approximately 750.[15] Almost 25% of respondents are in the retail sector, 20% in construction and 15% in manufacturing, followed by services, which ranges well above 10%. All other one-digit sectors have a single digit representation fraction. In terms of firm size, the sample contains much smaller enterprises than the BOS: the modal bin for the number of employees is "three to five", to which over 25% of respondents belong, followed by the "six to nine"-category with roughly 20%. The highest category is "forty or more", which contains just under 10% of firms.

We use three questions from the SBETS. The uncertainty index is based on a question about general business conditions just like in the BOS (the box and the bold font are also used in the original):

Q 13 *"About the* $\boxed{\textit{economy in general,}}$ *do you think that* ***six months from now*** *general business conditions will be better than they are now, about the same, or worse?:* $\boxed{1}$ *Much better,* $\boxed{2}$ *Somewhat better,* $\boxed{3}$ *About the same,* $\boxed{4}$ *Somewhat worse,* $\boxed{5}$ *Much worse,* $\boxed{6}$ *Don't know."*

One advantage of this question over its BOS cousin is that it is slightly more nuanced because it allows for two "increase"- and two "decrease"-categories. We quantify the extreme categories with -2 and 2, respectively. To measure activity in the SBETS we use:

Q 14 *"During the* ***last calendar quarter****, was your dollar sales volume higher, lower, or about the same as it was for the quarter before?* $\boxed{1}$ *Much higher* $\boxed{2}$ *Higher* $\boxed{3}$ *About the same,* $\boxed{4}$ *Lower* $\boxed{5}$ *Much lower."*

And as with the BOS we construct a turnover index for employment from an actual employment change question:

Q 15 *"During the* ***last three months****, did the* ***total*** *number of employees in your firm increase, decrease or stay about the same?* $\boxed{1}$ *Increased* $\boxed{2}$ *Decreased* $\boxed{3}$ *Stayed the same."*

Figure A11 displays the analog of Figure 1.2 in Section 1.4.1. Positive business uncertainty innovations lead to long and protracted negative reactions of the economic activity of small firms. Similarly to the BOS, there is little or no high-frequency impact followed by a strong rebound of economic activity.

Figure A12 is similar to Figure 1.4 in Section 1.4.1. It shows the impulse response of the job turnover measure to an innovation to uncertainty. As before, to the extent to which job turnover reacts to business uncertainty at all, it rises (at least the point estimate), which appears to be inconsistent with the "wait-and-see"-theory of uncertainty shocks.

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

Figure A11: Uncertainty Innovations on SBETS Sales Activity Index

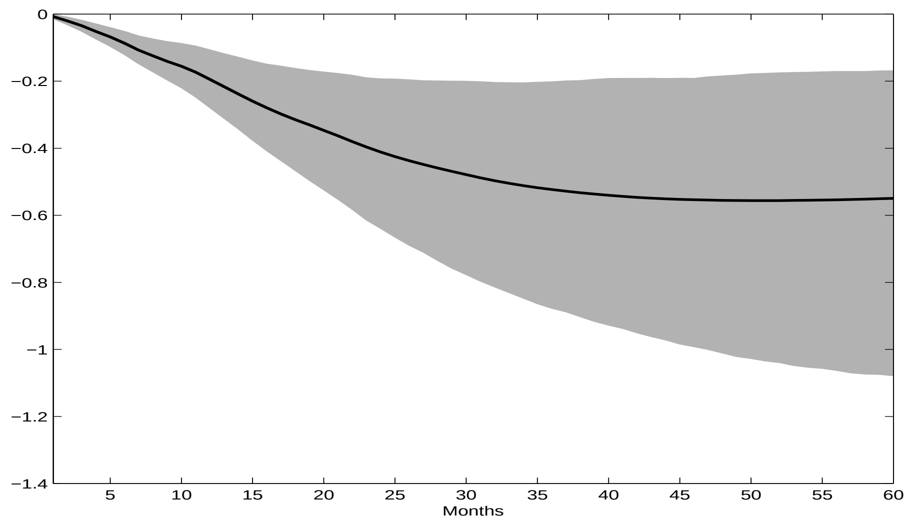

Notes: The uncertainty index is based on Q 13. The activity variable is based on Q 14. The impulse response is based on a two-variable SVAR with uncertainty ordered first, then activity, and 12 lags. It displays the response of the SBETS Sales Activity Index to a positive uncertainty innovation. All confidence bands are at the 95% significance level using Kilian's (1998) bias-corrected bootstrap.

Figure A12: Uncertainty Innovation on SBETS Job Turnover Index

Notes: see notes to Figure A11. The IRF is based on a two-variable SVAR with uncertainty ordered first and then job turnover. Job turnover is based on Q 15.

Finally, Figure A13 and Table A5 display the analogs of Figure 1.7 and Table 1.3 in Section 1.4.1. There is little, albeit compared to the BOS somewhat larger impact of uncertainty innovations to either sectoral economic activity or the economy-wide unemployment rate. There is again some impact of the long-run innovations on the uncertainty index.

Figure A13: A Three-Variable Blanchard-Quah-Type SVAR - SBETS

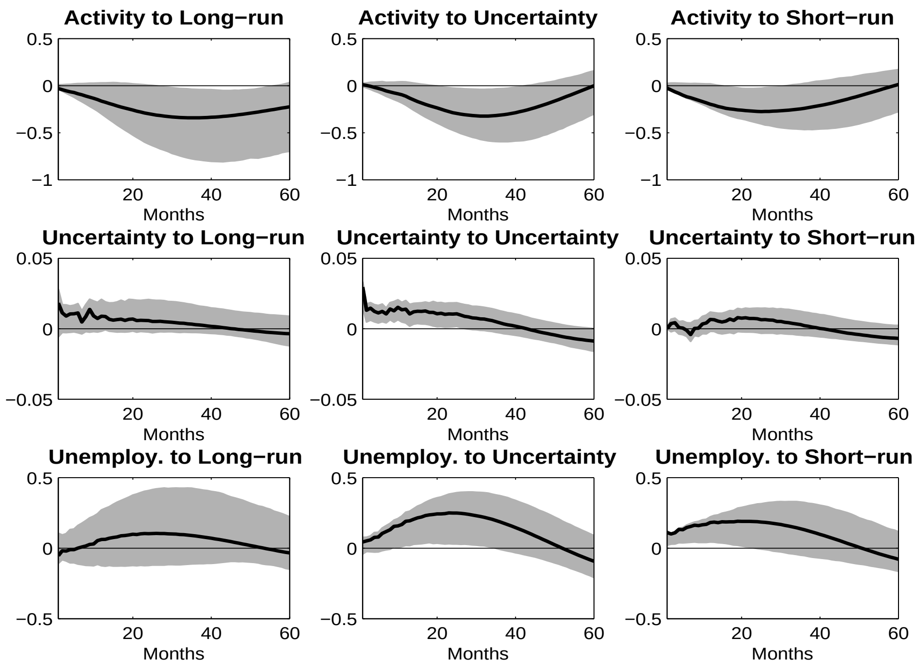

Notes: see notes to Figure A11. The unemployment rate is the (seasonally adjusted) monthly civilian unemployment rate from the BLS. The uncertainty innovation and the conventional short-run shock are identified as shocks that do not impact manufacturing production in the long-run. The conventional short-run shock is identified as the innovation that does not affect the uncertainty index on impact.

Table A5: Forecast Error Variance Decomposition - SBETS

	Shock	1M	3M	6M	1Y	2Y	5Y
Activity	Long-run	54%	45%	36%	34%	35%	46%
	Uncertainty	5%	1%	5%	12%	26%	30%
	Short-run	41%	54%	60%	54%	39%	24%
Uncertainty	Long-run	28%	30%	34%	32%	28%	24%
	Uncertainty	72%	69%	65%	65%	63%	61%
	Short-run	0%	2%	1%	3%	10%	15%
Unemployment Rate	Long-run	17%	8%	3%	3%	7%	9%
	Uncertainty	11%	17%	25%	40%	51%	54%
	Short-run	72%	75%	72%	58%	42%	37%

Notes: see notes to Figure A13.

Chapter 2

How Strongly Did the 2007/08 Oil Price Hike Contribute to the Subsequent Recession

Abstract[1]

What have been the economic consequences of the 2007/08 oil price hike for Germany? In this paper we implement a structural vector autoregression framework that distinguishes between supply shock and demand shock driven oil price changes. We find that the direct negative effects of oil price surges play a less important role for German manufacturing primarily producing investment and export goods. However, higher oil prices always considerably stifle private consumption expenditures. In a counterfactual analysis we show that the world demand driven 2007/08 oil price hike triggered a reduction of German GDP of 2.3 percent in the year 2009 and, therefore, notably contributed to the subsequent recession in Germany.

[1]This chapter is based on joint work with Kai Carstensen and Georg Paula. It is based on our paper "How Strongly Did the 2007/08 Oil Price Hike Contribute to the Subsequent Recession," mimeo, 2011. This is a revised version of our working paper that circulated under CESifo-WP 3357.

2.1 Introduction

What have been the economic consequences of the 2007/08 oil price hike for Germany? This question is particularly interesting since in the year 2009 Germany experienced with minus 5.1 percent its largest decline in GDP since 1949. While there exists much agreement that this global recession was caused by the worldwide financial crisis, Hamilton (2009) and Kilian (2009a) argue that the preceding oil price hike reaching 145 US dollars per barrel[2] at the beginning of July 2008 also significantly contributed to this recession. In this paper we take up this point and provide a comprehensive analysis concerning the effects of oil price shocks on the Germany economy. To do so, we implement the structural vector autoregression (SVAR) framework proposed by Kilian (2009b) that distinguishes between supply shock and demand shock driven oil price changes.

Since the worldwide stagflation period in the early 1970s oil price hikes have often been stated as reasons for subsequent recessions. Hamilton (1983) and Hamilton (2011) document that all U.S. postwar recessions except of the economic downturn in 1960/61 were preceded by oil price hikes. Therefore, extensive research has been carried out to analyze the effects of oil price shocks on aggregate activity. Recently, several studies have concluded that the oil price has lost its strong effect on the production level since 1984. Blanchard and Gali (2009) explain this finding with more flexible labor markets, more credible monetary policy and a smaller share of oil in the production process.[3] However, these studies typically assume that oil price innovations are homogenous over time. In a seminal contribution Kilian (2009b) highlights that the oil price is affected by structural demand and supply shocks which may have different effects on aggregate production as shown in Kilian (2009a) and Kilian (2009b). According to this view, the oft-cited structural break in the oil price-macroeconomy relationship only reflects shifts in the composition of oil demand and oil supply shocks that occurred over time.[4] This finding is important against the background that the oil price hike during the years 2002 to 2008 was mainly driven by increasing world demand as pointed out by Kilian (2009b), Hicks and Kilian (2011) and Hamilton (2009). It is even more important in the case of Germany, for which, as an export economy, the positive indirect effects on domestic production of a booming world economy can far overcompensate for the negative direct effects of an oil price increase. Thus, focussing on one structural oil price shock can be misleading for our analysis.

[2]See U.S. Energy Information Agency (EIA).

[3]Other studies discussing this point are e.g. Davis and Haltiwanger (2001), Hooker (1996), Kilian (2008b) and Kilian and Edelstein (2009).

[4]Ramey and Vine (2011) also provide evidence that no structural break in the oil price-macroeconomy relationship exists. However, their analysis focuses more on institutional settings like price controls and a complex system of entitlements that led to some rationing and shortages.

Our paper contributes to the literature in three ways: first, it studies the different effects of oil supply and oil demand shocks on the German economy at the aggregate level for the time period from 1973 until the beginning of 2011. While Germany is currently the largest European economy and fourth largest economy of the world, it has to be considered as a small open economy. According to the WTO Germany's share in world merchandise exports was in 2010 with 8.3 percent only slightly lower than the corresponding number for the United States (8.4 percent).[5] This number is noteworthy, since Germany's economic value added is less than one fourth of the production in the United States.[6] Therefore, one has to accentuate the German dependence with respect to the global business cycle. Second, this study provides additional disaggregate evidence for German manufacturing. A key feature of this sector is the great importance of the automobile industry. Bresnahan and Ramey (1993), Lee and Ni (2002) and Ramey and Vine (2011) stress the importance of this sector for the U.S. in explaining the consequences of oil price shocks. Interestingly, we observed a dramatic production decline in the German automobile sector since the middle of 2008 that ultimately lead to the introduction of a "Cash for Clunkers" program. Thus, it seems worthwhile to implement an analysis on the industrial level to understand the effects of oil price hikes. Finally, we provide an estimate by how much the 2007/08 oil price hike contributed the subsequent recession.

Even though there exists a large literature concerning the effects of oil price shocks on the United States, the evidence for Germany is scarce: Cunado and de Gracia (2003) and Jiménez-Rodríguez and Sánchez (2005) provide evidence of negative non-linear effects of oil price increases on German production. Blanchard and Gali (2009) show for the period since 1983 that an oil price increase leads to a rise in German GDP. They use a linear specification of oil price changes. Kilian (2008a) uses unexpected movements in global oil production to identify oil supply shocks. In his baseline results he finds that German GDP declines a few quarters after an oil supply shock. Peersman and van Robays (2009) use sign restrictions to identify oil demand and oil supply shocks. They find for Germany that an oil supply shock triggers an increase in German GDP.[7] Our analysis on the industrial level is motivated by Davis and Haltiwanger (2001), Herrera et al. (2011) and Lee and Ni (2002) who present results for the United States. For Germany Jiménez-Rodríguez (2008, 2011) studies the effects of oil price changes on industry sectors. However, she only analyzes the period between 1975 and 1998.

[5]In the year 2010 China was the biggest merchandise export nation with a share in world exports of 10.4 followed by the United States and Germany.

[6]This fraction is computed with nominal GDP numbers in U.S. dollar of the year 2010. Data source is the World Economic Outlook of the International Monetary Fund.

[7]More recently, Kilian and Murphy (2012) show that using only sign restrictions for identification of oil supply and demand shock leads to distorted results.

With the exception of Peersman and van Robays (2009), all studies of the German case assume there is only one structural oil price shock. This is a central difference to our analysis because we use the currently dominating approach proposed by Kilian (2009b), Kilian and Park (2009) and Kilian et al. (2009) that distinguishes between different sorts of oil price shocks. Furthermore, former studies often rely on nonlinear transformations of the oil price variable or assume asymmetric effects of oil price increases versus decreases. While we are aware of potential nonlinearities, it is beyond the scope of this paper to add to the ongoing debate on this topic that has attracted much attention in recent times (see among others Kilian and Vigfusson (2011a), Kilian and Vigfusson (2011b), Hamilton (2011) and Herrera et al. (2011)). Thus, we apply a linear framework.

Our results show that aggregate production in Germany significantly declines after an oil supply shock. For the aggregate demand shock we find an significant increase in GDP during the first year. Afterwards, however, the negative direct effects of the oil price increase prevail and German GDP falls considerably. While these two results are broadly in line with those presented by Kilian (2009b) for the United States, the dynamic response of German GDP following an oil-specific demand shock exhibits a completely different picture. German production rises persistently after an oil-specific demand shock. We explain this finding by favorable international price movements and shifts in global demand that favor German exports. To support these explanations we show at the industrial level that in particular German exporting firms experience production increases after aggregate and oil-specific demand shocks. We further show by means of a counterfactual analysis that the oil price hike in 2007/08 triggered a 2.3 percent reduction of German GDP in the year 2009 and, therefore, made a notable contribution to the subsequent recession in Germany.

The remainder of the paper is structured as follows. Section 2.2 provides statistics concerning German oil consumption and manufacturing. In Section 2.3 we outline our empirical framework that is used throughout the whole analysis. The subsequent Sections 2.4 and 2.5 present empirical results at the aggregate and disaggregate level of Germany. In Section 2.6 we ask whether the German economy would have experienced the recent recession without the preceding oil price hike. Section 2.7 concludes.

2.2 Statistics on German Oil Consumption

In 2008 Germany imported 105 million tons of crude oil and increased its oil imports by more than 17 million tons since 1991.[8] Petroleum is the single most important energy source in Germany. In 2010 it contributed roughly 34 percent to the primary energy consumption in

[8]Due to the recession crude oil imports declined to 98 millions tons in 2009. These data are taken from the national working group "Energiebilanzen" and are available at www.ag-energiebilanzen.de.

Germany. This figure has not considerably changed during the last 20 years. The following most important energy sources are natural gas (22 percent), coal (12 percent) and nuclear energy (11 percent). The pattern of oil consumption, however, differs remarkably across sectors. While petroleum is very important for households (23 percent of energy consumption in 2009), transportation (93 percent) and the trade and services sector (15 percent), its contribution to energy consumption in German manufacturing amounted to 6 percent.

Similar to countries such as the U.S. or the U.K. manufacturing in Germany has become less important over the last 40 years. Its share in gross value added declined from 37 percent in 1970 to 21 percent in 2010. Nonetheless the collapse of manufacturing production has contributed to a great extent to the recession in 2009. For our analysis we choose six industrial sectors that feature two characteristics: first, they account for a large share in German manufacturing and, second, energy is a crucial input in their production process. The chosen sectors are displayed in Table 2.1.[9]

Table 2.1: Statistics concerning German Manufacturing

Industrial sector	Share in manufacturing in percent (volume of sales)	Energy intensity (cost of energy in euro cent for each euro of sale)	Export ratio of production in percent (volume of sales)
Refined petroleum	7.5	0.6	7.2
Chemicals and chemical products	7.7	5.0	58.5
Basic metals	6.5	6.0	38.5
Fabricated metal products	5.9	1.9	32.1
Machinery	13.0	0.9	61.1
Automobile & transport equipment	19.2	0.8	62.9
Other industrial sectors	40.1	2.1	39.0
Manufacturing	100	2.1	46.3

Notes: We use sales and energy cost data from the year 2008 to compute the numbers in the first two columns. The last column presents numbers of establishments with at least 50 employees for 2010.

The largest German industrial sector is automobile & transport equipment with a production share of almost 20 percent. Machinery and the chemical industry follow on the second and third place. In total, the six selected industrial sectors represent almost 60 percent of total manufacturing production. Compared to the other sectors, “chemicals and chemical products” and “basic metals” have the highest energy intensities.

[9]The data source is the German Federal Statistical Agency. To compute the shares in manufacturing and energy intensities we use data from 2008. Using instead the latest available data from 2009 would not considerably change the results.

One distinctive feature of German manufacturing is its high export share. Establishments with at least 50 employees exported 46 percent of their production in 2010. The major export industries are automobile & transport equipment, machinery and the chemical industry, which export considerably more than half of their production.

2.3 Empirical Framework

2.3.1 The Structural VAR Model

Our empirical approach is based on the structural VAR model of Kilian (2009b) that describes the global crude oil market. This model accounts for the simultaneity between crude oil supply and demand, and it allows to decompose unexpected oil price changes into shocks to world oil supply, to global demand, and to oil-specific demand. The latter captures shifts in market concerns about the availability of future oil supply and may therefore also be called precautionary demand for oil. The VAR model has the form

$$y_t = c + A(L)y_{t-1} + u_t, \tag{2.1}$$

where c is a vector of constants and $A(L)$ denotes a lag polynomial. The vector y_t includes the percent change in world crude oil production, a measure of global real activity, and the real price of oil. A more detailed description of the data is given below.

To estimate the effects the shocks to the global oil market have on the German economy, we follow the approach outlined in Kilian and Park (2009) and add one German variable at a time to the oil market model (2.1). This yields

$$\begin{pmatrix} y_t \\ z_t \end{pmatrix} = \tilde{c} + \tilde{A}(L) \begin{pmatrix} y_{t-1} \\ z_{t-1} \end{pmatrix} + \tilde{u}_t, \tag{2.2}$$

where z_t is the additional German variable.[10] We estimate the VAR models using monthly data over the period from January 1973 to March 2011. We include 24 lags to allow for delayed effects of up to two years. To account for conditional heteroskedasticity in the monthly data we construct our confidence bands using the recursive design wild bootstrap proposed by Gonçalves and Kilian (2004). The structural shocks are identified using the

[10]One might argue that e.g. German manufacturing production is not able to influence the global oil market at all as its share of world production is rather small. However, our results do not change if we would use instead a subset VAR that does not allow the lags of the German variable z_t to affect the oil market variables summarized in y_t.

recursiveness assumption proposed by Kilian (2009b) and Kilian and Park (2009):

$$\tilde{u}_t \equiv \begin{pmatrix} \tilde{u}_{1,t}^{\textit{oil production}} \\ \tilde{u}_{2,t}^{\textit{real activity}} \\ \tilde{u}_{3,t}^{\textit{real oil price}} \\ \tilde{u}_{4,t}^{\textit{German variable}} \end{pmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \begin{pmatrix} e_{1,t}^{\textit{oil supply shock}} \\ e_{2,t}^{\textit{aggregate demand shock}} \\ e_{3,t}^{\textit{oil-specific demand shock}} \\ e_{4,t}^{\textit{German shock}} \end{pmatrix}. \tag{2.3}$$

Specifically, we assume that the short-run global oil supply curve is vertical. Hence, only oil supply shocks may lead to instantaneous changes in global oil production. The aggregate demand shock (for industrial commodities) is defined as the innovation in the global real activity index that is not explained by oil supply shocks. This implies that oil-specific demand shocks do not affect the global business cycle within the month. Therefore, oil-specific demand shocks are defined as the part of the surprise changes in the real oil price that is not explained by oil supply and world demand shocks. Finally, we assume that shocks originating in the German economy do not affect the global oil market block within a month.

To analyze the effects of the structural shocks on the quarterly macroeconomic variables q_t, we again follow Kilian (2009b) who calculates quarterly averages of the shocks, say, $\bar{e}_{j,t}$, $j = 1,2,3$, and estimates regressions of the form

$$q_t = \alpha_j + \sum_{i=0}^{12} \phi_{ji} \bar{e}_{j,t-i} + \varepsilon_{j,t}, \quad j = 1,2,3. \tag{2.4}$$

The number of lags is set to 12 quarters. Including the contemporaneous value of $\bar{e}_{j,t}$ amounts to assuming that the global shocks are predetermined to German variables which seems plausible.[11]

2.3.2 The Data

For the global oil market we use an updated version of the data analyzed by Kilian (2009b). World crude oil production is provided by the U.S. Energy Information Administration. The global real activity index is constructed from single voyage bulk dry cargo ocean shipping freight rates. Assuming that the supply of shipping capacity is fixed in the short run, changes in freight rates reflect the development in world demand for industrial commodities such as grain, coal, and scrap metal. Kilian (2009b) shows that this real activity index is superior to the industrial production index of the OECD countries.[12] Finally, for the real price of

[11] As a robustness check, we aggregated monthly German variables such as industrial production to the quarterly frequency and ran both the monthly oil market VAR (2.2) and the quarterly regression (2.4). The resulting impulse response functions were—up to the frequency—largely the same.

[12] The index is made available by Lutz Kilian at www-personal.umich.edu/~lkilian/paperlinks.html.

oil we use the U.S. refiner's acquisition cost of crude oil deflated by the U.S. CPI.[13] To analyze the impact of the various shocks on the German macroeconomy, we include five real quarterly national accounts variables, namely GDP, private consumption expenditures, gross investment, exports, and imports. All variables are seasonally adjusted and used in percent changes.[14] To avoid a break in the time series due to German reunification in 1991, post-1991 data (which refer to reunited Germany) are extended backwards by using the growth rates of the pre-1991 data (which refer to Western Germany only).

The monthly German variables are manufacturing production, several price and exchange rate indices. Further, we link currently available real production data at the industry level with discontinued earlier series that are available back to 1970.[15] We chain these time series for those sectors for which we can ensure that their definition has not changed since 1970. Finally, we adjust all these variables for seasonal and calendar effects and use them as percent changes.

A major strike by the union of metal workers to reduce the weekly workload to 35 hours affected the German automobile industry in May and June 1984. Production went down by 16 and 51 percent, respectively, just to recover with rates of 170 and 18 percent in July and August. To preclude that this exogenous event contaminates any of the estimation results, we replace the four monthly observations with forecasts of an autoregressive model with 12 lags that is fitted to the remaining sample.[16]

2.4 Empirical Results at the Aggregate Level

2.4.1 Results of the Structural Oil Market Model

In this section we briefly present the results of the structural oil market model (2.1). As already mentioned we update the data set of Kilian (2009b) until March 2011. Therefore, it is not surprising that our results are similar. Figure 2.1 summarizes the responses of the oil market variables to the structural shocks. An adverse oil supply shock leads to a sizeable and permanent decline in oil production. The real price of oil goes up slightly but without

[13]The data sources are the U.S. Energy Information Administration and FRED. The oil price series has been extended backwards by using the original time series of Kilian (2009b) that is available at www.aeaweb.org/articles.php?doi=10.1257/aer.99.3.1053. The real oil price series is expressed as deviations from the mean.

[14]The percent changes are not annualized so a cumulative impulse response can directly be interpreted as the percent difference between the initial level and the level that is triggered by the shock.

[15]All these time series are available at the German Federal Statistical Agency.

[16]Our estimation results reported below are robust to different lag orders for the autoregressive adjustment model or to using average growth rates instead. The effect of the strike is also visible in total manufacturing production as well as in GDP and exports. However, replacing the respective observations with autoregressive forecasts yields only negligible changes in our results. Therefore, we prefer to use the original data.

Figure 2.1: Responses to Structural Shocks to the Global Oil Market

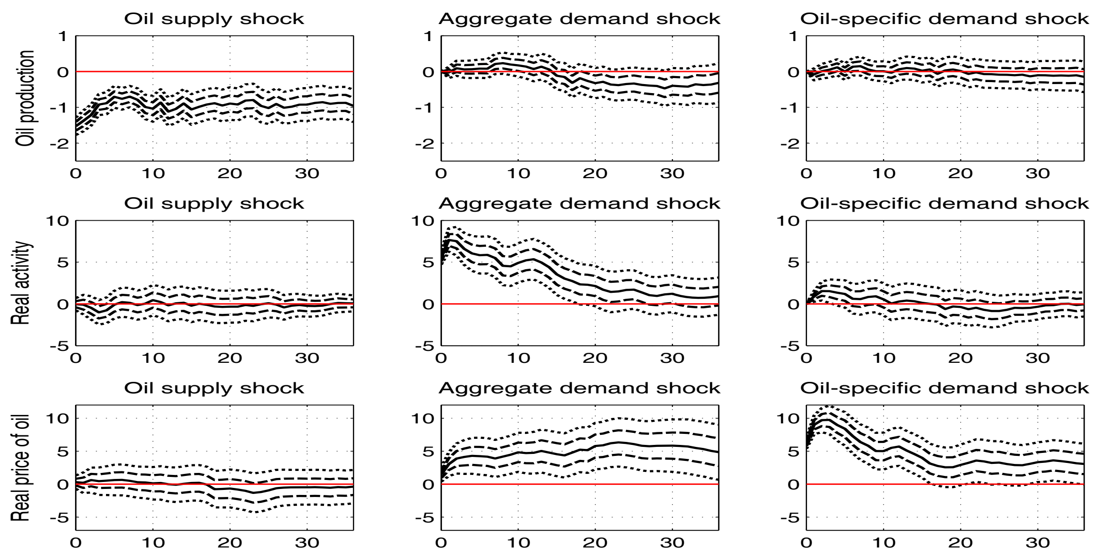

Notes: The impulse response functions are estimated with model (2.1). The time dimension (horizontal axis) is measured in months. The confidence bands (one and two standard deviations) are constructed using the recursive design wild bootstrap of Gonçalves and Kilian (2004).

statistical significance. Real activity shows almost no reaction. An expansionary aggregate demand shock immediately triggers a strong and long-lasting increase in real activity. As a consequence, the real oil price significantly increases for a sustained period of time. In contrast, the rise in oil production comes with a delay of seven months and is only marginally significant. An unanticipated increase in oil-specific demand causes a strong hike in the real price of oil while oil production does not react significantly. The shock also triggers a statistically significant short-lived increase in economic activity.

Figure 2.2 shows the cumulative contributions of the structural shocks to the historical evolution of the real oil price. We concentrate on the oil price developments since 2002. Between 2002 and 2008 the oil price hike was fueled almost exclusively by expansionary aggregate demand shocks. The cumulative effects coming from changes in oil supply or in precautionary oil demand played only a minor role in this episode. However, the rapid fall in the oil price witnessed during the recent recession was only partly due to reduced global demand for industrial commodities but mainly caused by negative oil-market specific shocks. Probably market participants were concerned about how quickly the world economy would recover and the world demand for oil would pick up again. Therefore, the precautionary demand for oil fell considerably. By the end of the sample, confidence seems to have returned as oil-market specific shocks contributed positively to the real oil price.

Figure 2.2: Historical Decomposition of the Real Price of Oil 2002:1 to 2011:3

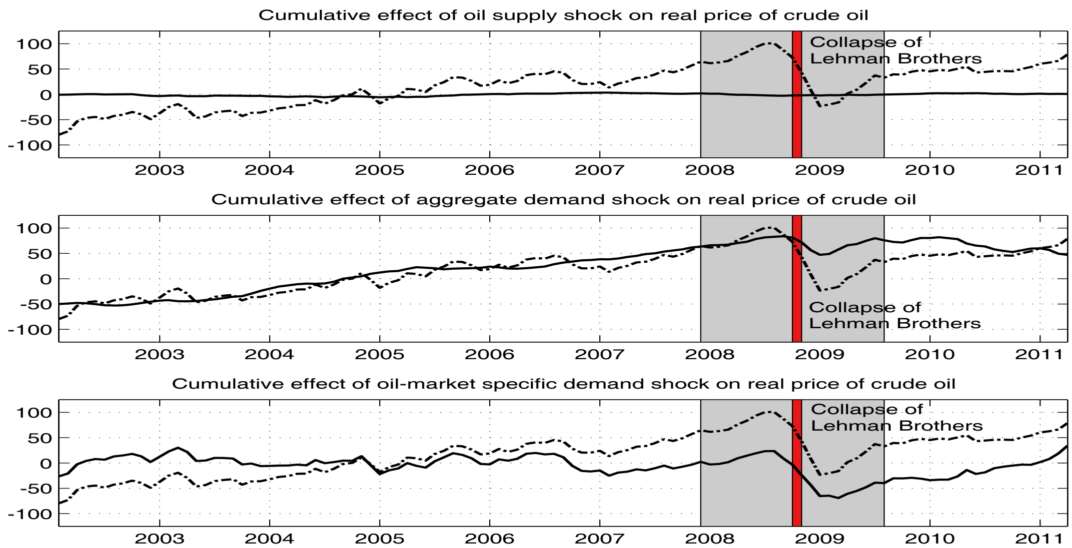

Notes: Estimates are based on model (2.1). The dot-dashed lines denote the real price of oil. The solid lines show the fluctuations in the real price of oil that are explained by the respective structural shock. The shaded grey areas define the last U.S. recession dated by the NBER. The red vertical bar represents the date of the collapse of Lehman Brothers. The time dimension (horizontal axis) is measured in years.

2.4.2 The Reaction of German Macroeconomic Aggregates

The reaction of German macroeconomic aggregates to the three structural shocks are displayed in Figure 2.3. To help interpret our results, we show in addition the responses of exchange rates and prices in Figure 2.4.

An oil supply disruption lowers GDP instantaneously.[17] However, this decline becomes statistically significant only after seven quarters. Three years after the shock, GDP is roughly 1.5 percent below the initial level which accords well with the finding of Kilian (2008a) for Germany. However, the effect is much more persistent than in the U.S. see Kilian (2009b). The decline in German GDP significantly extends to all demand aggregates.[18] Within three years, private consumption falls by 1 percent, gross investment by 4.4 percent, imports by 4 percent, and exports by 2.7 percent. Initially, however, only gross investment declines

[17] This stands in contrast to findings by Kilian (2008a) who reports an immediate increase in German GDP after a contractionary oil supply shock. He argues that this finding might be caused by a "spurious sample correlation between economic expansions in Germany and exogenous oil supply disruptions" and should disappear as more observations are added. Our result is based on an extended sample and seems to confirm this view.

[18] Using a VAR model with sign restrictions, Peersman and van Robays (2009) find that Germany—unlike all other euro area countries—experiences a persistent increase in GDP, consumption and investment after an oil supply disruption.

significantly as measured by the one-standard-error bands. Imports follow closely behind, while consumption and particularly exports react with a considerable delay.

To put these results into perspective, it is informative to consider the responses of exchange rates and prices. As also documented by Kilian (2008a), the nominal exchange rate of the euro against the U.S. dollar depreciates quickly after the shock, see Figure 2.4. The real effective exchange rate shows the same reaction suggesting that the euro prices of oil and of imported goods in general increase. This exchange rate effect may explain the swift decline in imports and the lagged reaction of exports and could account for a major difference between the responses within the U.S. and Germany. While there is an immediate drop in U.S. GDP as reported by Kilian (2009b), German GDP declines more sluggishly. As one factor behind this difference, after an adverse oil supply shock U.S. exports suffer from the immediate appreciation of the U.S. dollar while German exports are temporarily shielded by an external devaluation. Consistent with this interpretation, there does not seem to be a need for an internal devaluation in Germany as the CPI remains largely unaffected. Given the hike in import prices however, this implies that domestic prices fall, probably triggered by the reduction in consumption demand.

A positive aggregate demand shock leads to an increase in GDP within the first year that is statistically significant in terms of one-standard-error bands. This suggests that the primary effect of higher world demand initially outweighs the contractionary effect of higher oil prices. Two years after the shock, however, the price effects begin to dominate and GDP turns significantly negative. This interpretation is supported by the observation that export and investment demand show a significant positive reaction within the first seven quarters, while private consumption falls steadily and all prices climb up persistently and highly significantly. Three years after the shock, consumption is 1.3 percent below the initial level, while consumer prices have gone up by 0.5 percent.

These findings are similar to those obtained by Kilian (2009b) for the U.S. economy.[19] However, the initially positive GDP response is somewhat more pronounced for Germany which might reflect the greater export dependency of the German economy. Additionally, we find that German exports benefit from international price movements. The expansionary world demand shock leads to a deterioration of the German terms of trade, i.e., import prices increase faster than export prices. This, in turn, improves the price competitiveness of German firms.

[19]Note that the difference in magnitude is mainly caused by Kilian's approach to cumulate annualized quarter-on-quarter growth rates while we cumulate raw growth rates which seems a more natural way when one is interested in comparing levels.

Figure 2.3: Responses of GDP and its Components to the Structural Shocks

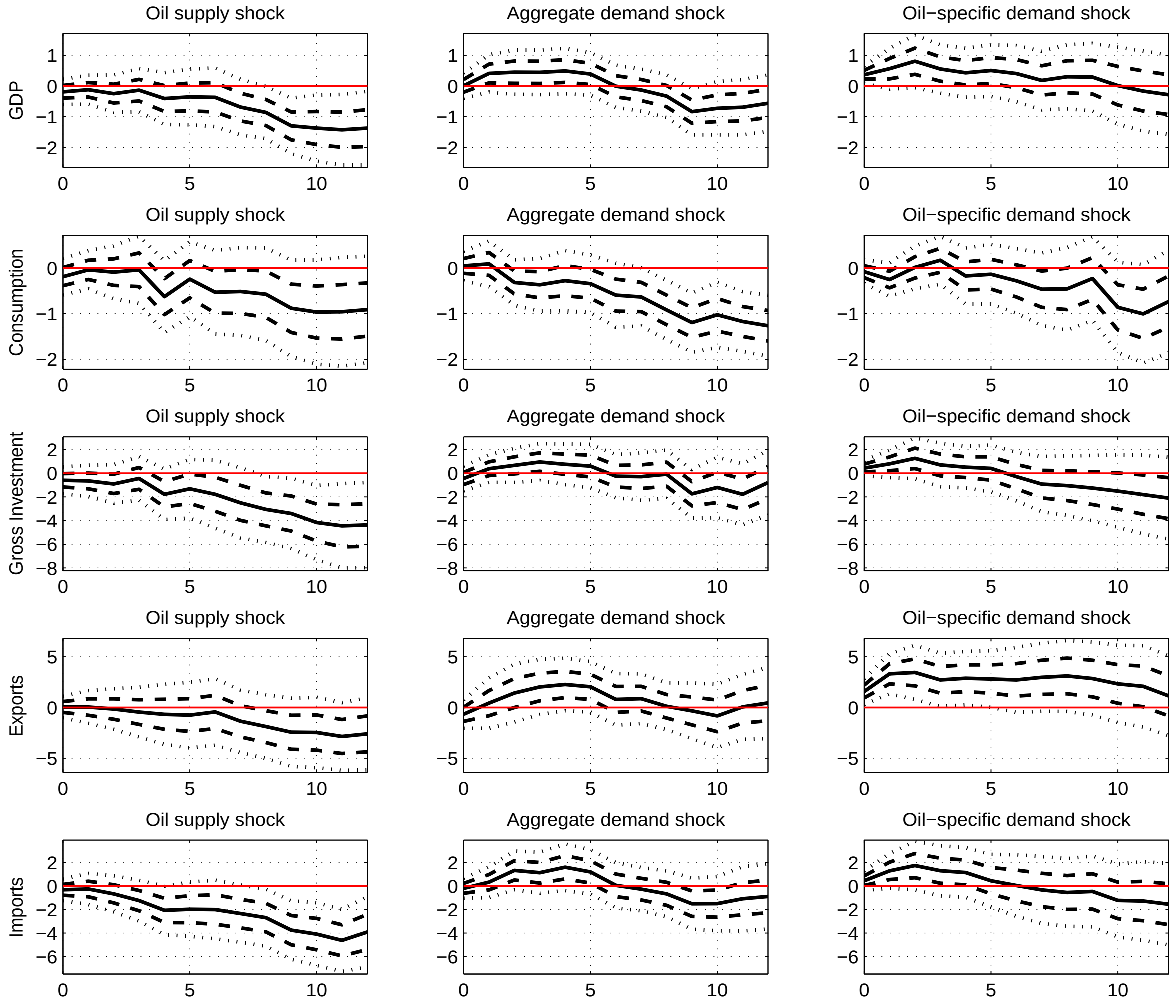

Notes: Impulse response functions are estimated with model (2.4). The time dimension (horizontal axis) is measured in quarters. The confidence bands (one and two standard deviations) are constructed using a block bootstrap method with block size 4 and 20,000 bootstrap replications.

An unexpected increase in oil-market specific demand initially raises GDP which is consistent with the previous result that world activity accelerates within the first few months. This response is mirrored by a quick and long-lasting rise in exports which is supported by a real devaluation as export prices rise by less than import prices. On the flip side, the strong hike in real oil prices significantly feeds through to consumer prices. The corresponding loss in purchasing power lowers private consumption demand and eventually pulls down GDP.

The result of a temporary rise in German GDP after an oil-specific demand shock contrasts with the U.S. experience reported by Kilian (2009b) that GDP declines steadily. Three observations may help explain this difference. First, German exporters do not completely

Figure 2.4: Responses of Exchange Rates and Prices to the Structural Shocks

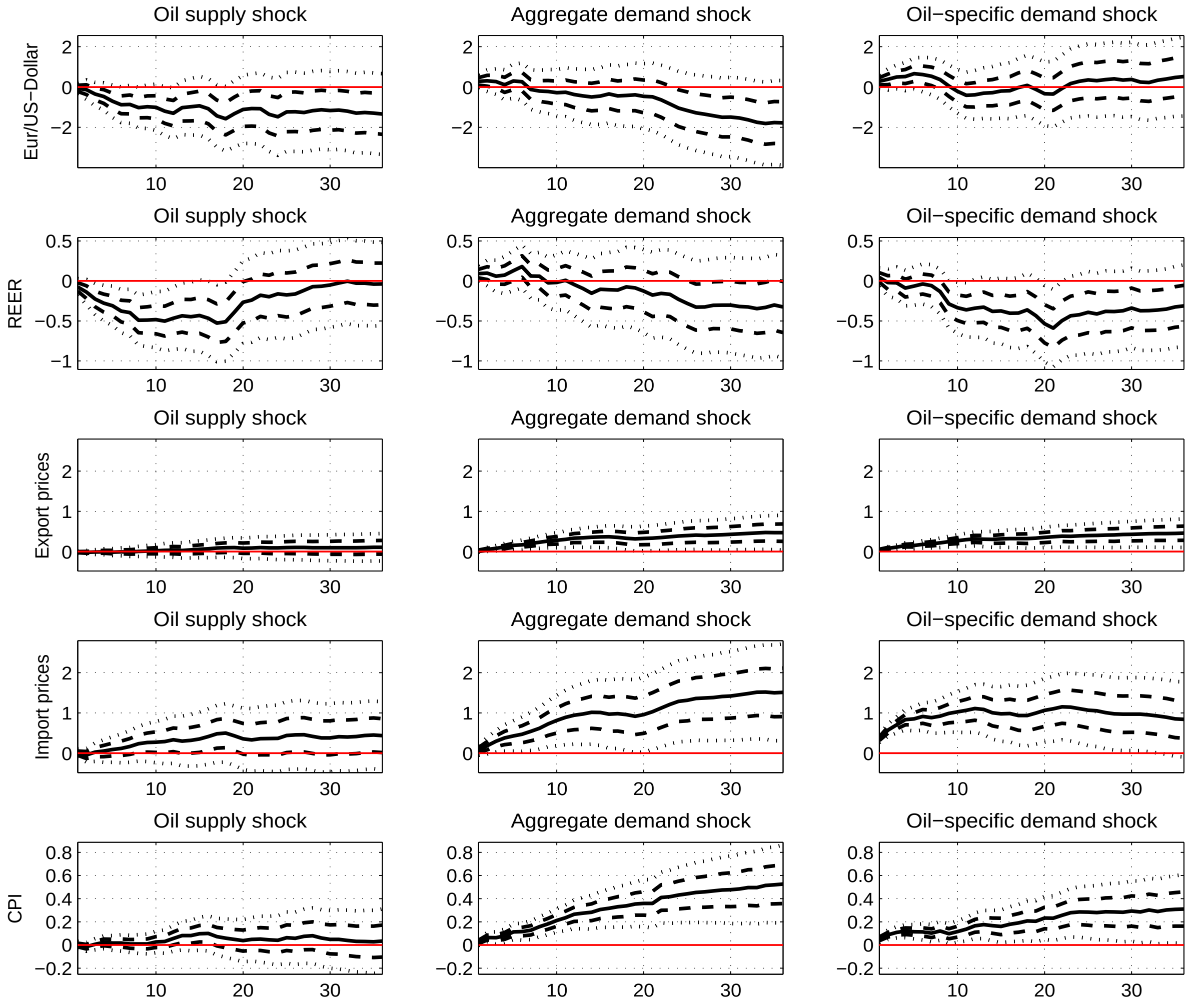

Notes: Impulse response functions are estimated with model (2.2). REER denotes the real effective exchange rate against the main trading partners, CPI is the headline consumer price index. The time dimension (horizontal axis) is measured in months. The confidence bands (one and two standard deviations) are constructed using the recursive design wild bootstrap of Gonçalves and Kilian (2004).

pass the higher oil price into their export prices. This leads to a gain in their price competitiveness and, therefore, a relatively higher demand for their products. This argument is supported by the reactions of import and export prices shown in Figure 2.4. Second, the immediate increase in the real oil price after an oil-specific demand shock triggers shifts in global demand. The most prominent example of such demand shifts is the U.S. automobile market. Kilian and Edelstein (2009) document that after energy price shocks the demand for U.S. automobiles falls, whereas the demand for foreign more energy-efficient cars evolves much more positively. Third, the German export portfolio mainly consisting of a broad range of investment goods seems to fit well the demands of many oil-exporting countries, i.e., Germany benefits from petrodollar recycling. For all these reasons, German exports react positively to an oil-specific demand shock which temporarily outweighs the negative consumption effect of higher oil prices.

Taken together we conclude that it matters which of the three shocks hits the German economy. While we find that consumption declines markedly in all cases, the reactions of exports and gross investment—and finally GDP—depend on the type of the structural shock. Not surprisingly, the primary effect of an expansionary world demand shock on GDP is positive until the oil price effect weighs in and the boom loses momentum. Somewhat more unexpectedly, however, after an oil-market specific demand shock a redirection of world demand towards German export goods seems to counteract the contractionary effect of higher oil prices on consumers' demand. This implies that the consumption demand effects stressed by Kilian and Edelstein (2009) and Kilian (2008b) could be much less important for, and thus less harmful to, the German compared to the U.S. manufacturing sector. In the following section, we analyze this interpretation in more detail.

2.5 Results at the Industrial Level

One of the leading explanations why oil price hikes have harmed the U.S. economy more than the cost share of oil suggests, is that households not only cut back on consumption due to a loss in purchasing power and an increase in precautionary savings but also shift their demand for automobiles away from U.S. products, see Kilian and Edelstein (2009). To the extent that this demand is targeted to automobiles produced in Germany, the German automobile industry should be less affected than its U.S. counterpart. Since, in addition, German production of investment goods in general may benefit from petrodollar recycling, in a next step we examine several important industrial sectors in more detail. Subsequently, we analyze whether exporting and non-exporting firms react differently as our interpretation suggests that exporters which benefit directly from the redirection effect do better after an oil price shock than non-exporters which might be more strongly exposed to the slump in domestic consumption demand.

2.5.1 The Reaction of Industrial Production

After an oil supply shock industrial production declines within a few months and persistently remains below the initial level, see Figure 2.5. This decline is statistically significant according to the one-standard-deviation band but economically moderate.

Figure 2.5: Responses of Total Manufacturing and Industry-Level Production to the Structural Shocks

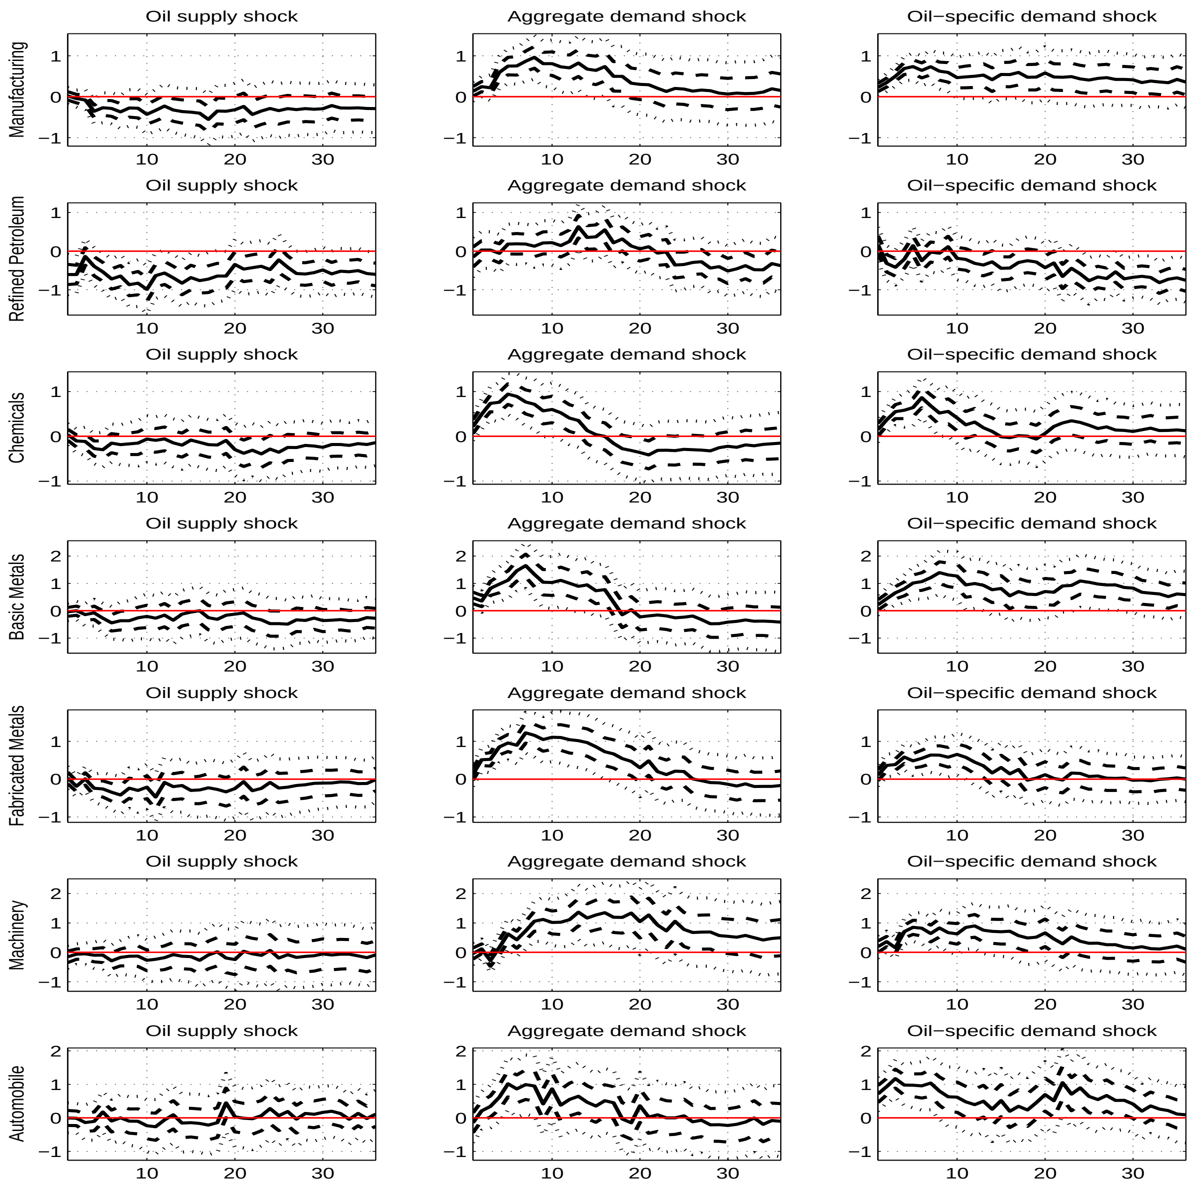

Notes: see notes to Figure 2.4. The time dimension (horizontal axis) is measured in months.

Petroleum refinery is the industrial sector which suffers most from reduced oil supply. However, compared to the GDP response even this effect is moderate, with a minimum at -1 percent after 10 months. The production of more primary goods such as chemicals and basic metals also persistently declines, even though the reactions are at most significant according to the one-standard-deviation band. In contrast, output in the machinery and automobile sectors remains largely unaffected. The visual impression that oil supply shocks seem to play no important role for German manufacturing is further supported by considering the forecast error variance in Table 2.2. Oil supply shocks are not able to explain a noticeable fraction of the forecast error variance at all horizons with the only exception being petroleum refinery. All in all, we conclude that oil supply disruptions have not been a relevant source of gross industrial output fluctuations in Germany. This implies that the transmission of oil supply shocks works mainly through changes in the exchange rate and in private demand rather than through industrial production.

An expansionary world demand shock triggers a swift and statistically significant increase of industrial production that peaks at 1 percent after eight months and subsequently phases out. In contrast to GDP, we observe no medium-term decline below the initial level. The production of chemicals, basic metals and fabricated metal products rises quickly while the automobile and particularly the machinery sector lag somewhat behind which is a typical characteristic of the German business cycle. Interestingly, after 18 months the production of more primary goods such as refined oil, chemicals and basic metals undershoots the initial output level which indicates that rising prices for oil and presumably other commodities tend to reduce the demand for these goods considerably. In contrast, machinery production is less affected by these price effects and remains persistently positive. In terms of the forecast error variance decomposition, the aggregate demand shock is much more important in explaining output fluctuations than the oil supply shock. At the one-year horizon it explains almost 17 percent of the forecast error in total manufacturing production. This share varies between 0.8 percent for petroleum refinery and 21.5 percent for fabricated metal products.

Table 2.2: Forecast Error Variance Decomposition of Manufacturing and Industry-Level Output

	Shock	1M	3M	6M	1Y	2Y	5Y
Manufacturing	Oil supply shock	0.01%	0.15%	2.07%	2.76%	4.24%	4.56%
	Aggregate demand shock	0.65%	2.16%	11.47%	16.64%	14.13%	7.99%
	Oil-specific demand shock	1.92%	7.25%	13.21%	11.35%	11.36%	10.17%
	Other shocks	97.42%	90.44%	73.24%	69.26%	70.27%	77.28%
Refined petroleum	Oil supply shock	2.41%	3.44%	5.74%	12.76%	16.44%	18.60%
	Aggregate demand shock	0.16%	0.12%	0.33%	0.79%	2.70%	8.87%
	Oil-specific demand shock	0.19%	1.39%	1.43%	1.51%	4.55%	21.11%
	Other shocks	97.24%	95.06%	92.51%	84.95%	76.31%	51.42%
Chemical products	Oil supply shock	0.03%	0.26%	1.04%	0.81%	1.88%	2.46%
	Aggregate demand shock	1.00%	8.46%	15.12%	13.55%	11.23%	6.84%
	Oil-specific demand shock	0.46%	4.68%	9.57%	8.51%	6.72%	4.17%
	Other shocks	98.52%	86.60%	74.27%	77.13%	80.17%	86.53%
Basic metals	Oil supply shock	0.02%	0.12%	0.72%	0.82%	1.18%	2.14%
	Aggregate demand shock	2.50%	5.58%	12.87%	15.09%	13.41%	8.16%
	Oil-specific demand shock	0.52%	3.67%	8.38%	12.84%	15.08%	17.87%
	Other shocks	96.96%	90.62%	78.03%	71.25%	70.33%	71.83%
Fabricated metal products	Oil supply shock	0.02%	0.35%	1.31%	1.81%	2.06%	1.62%
	Aggregate demand shock	0.26%	5.42%	13.97%	21.55%	19.71%	11.14%
	Oil-specific demand shock	0.18%	2.87%	5.45%	6.76%	4.81%	2.70%
	Other shocks	99.54%	91.36%	79.28%	69.88%	73.43%	84.54%
Machinery	Oil supply shock	0.24%	0.21%	0.46%	0.51%	0.43%	0.34%
	Aggregate demand shock	0.02%	0.67%	2.41%	9.64%	16.42%	12.92%
	Oil-specific demand shock	0.38%	1.38%	7.23%	7.75%	6.74%	4.23%
	Other shocks	99.37%	97.74%	89.90%	82.10%	76.42%	82.51%
Automobile and transport	Oil supply shock	0.12%	0.11%	0.40%	0.47%	2.09%	1.78%
	Aggregate demand shock	0.15%	0.11%	0.76%	1.92%	2.41%	2.78%
	Oil-specific demand shock	0.65%	4.95%	6.55%	6.60%	8.04%	7.44%
	Other shocks	99.08%	94.84%	92.29%	91.01%	87.46%	87.99%

Notes: see notes to Figure 2.4.

A positive oil-specific demand shocks leads to a statistically significant and persistent increase in industrial production. Even after three years, production is 0.3 percent above the initial level. This is consistent with our previous discussion of favorable international price movements and, in particular, shifts in global demand towards German products. These effects do not apply to the refinery sector which exports less than 10 percent of its production and is most directly affected by the oil price hike. Therefore, it not surprising that this is the only sector that exhibits a decline in output while all other sectors show a positive response. Most striking is the impact response of the automobile sector which, unlike all other sectors, already on impact increases production by almost 1 percent. Using the forecast error variance decompositions, it turns out that a noticeable fraction of the gross industrial output fluctuations can be explained by the oil-market specific demand shock. For total manufacturing production the maximum share is 13 percent at the six-month horizon, again with considerable variation between individual sectors.

The automobile sector receives particular attention in the discussion of the effects of oil price changes on the U.S. economy (see, e.g., Bresnahan and Ramey, 1993, as well as Ramey and Vine, 2011). Kilian (2008b) and Kilian and Edelstein (2009) document that oil price hikes do not only depress automobile demand on average but also shift demand towards foreign cars with the consequence that non-U.S. carmakers temporarily increase their U.S. sales. This is consistent with the view that European and Japanese producers have a comparative advantage with respect to smaller and more energy-efficient cars, as argued, e.g., by Lee and Ni (2002). Thus, the net effect on German automobile production is unclear. Our analysis reveals in the bottom panel of Figure 2.5 that demand driven oil price hikes go hand in hand with increases in German automobile production. However, Table 2.2 shows that all three structural oil price shocks together do not explain a sizeable part of the observed variation in production. In short, oil price shocks seem to have only a rather limited effect on German automobile production.

Altogether, we conclude that German manufacturing benefits from positive shocks to the real oil price: on the one hand, the energy cost share of most German industry sectors is small (see Section 2.2). On the other hand, the effects of favorable international price movements and shifts in global demand towards German products seem to dominate a general slump in private consumption. However, this result does not mean that the German economy as a whole is better off after an oil price shock. Our findings indicate that the reaction of German households is not much different from that documented for the U.S. This may have adverse consequences for the service sector. An analysis along these lines is left for future research.

2.5.2 The Reaction of Exporters and Non-Exporters

So far we have seen that German manufacturing is positively affected by aggregate and oil-specific demand shocks. The above explanations for this finding imply that exporting firms are affected by oil price shocks in a different way than non-exporting firms. Using data of IFO Business Climate Survey (IFO-BCS) of German manufacturing firms allows us to check this conjecture. The IFO-BCS index is a much-followed leading indicator for economic activity in Germany. It is based on a firm survey which has been conducted since 1949 and, therefore, is one of the oldest and broadest monthly business confidence surveys available (see Becker and Wohlrabe, 2008, for details). One of its main advantages is the broad coverage including approximately 5,000 respondents at the beginning of our sample and still about 2,500 towards the end.[20] Firms are asked about their business situation as well as their expectations and actual realizations for a broad set of firm-specific variables such as production, prices, demand and export situation.

In our analysis we focus on the following two questions concerning expected and current business situation:

Q 1 *"Expectations for the next six months: Our business situation with respect to XY will in a cyclical view: improve, remain about the same, develop unfavourably."*

Q 2 *"Current situation: We evaluate our business situation with respect to product XY as: good, satisfactory, unsatisfactory."*

It is noteworthy to mention that question Q1 refers to a *change* in the future business situation whereas the question Q2 concerns the actual business state of a firm. The IFO-BCS enables us to distinguish between exporting and non-exporting firms.[21] Therefore, we are able to compute and compare aggregate survey results for these two groups of firms in each month. Specifically, we calculate for each group balance statistics that are defined by the difference between the percentage shares of positive and negative responses for the respective question (Q1 or Q2). After seasonal adjustment, we add one of these time series at a time as fourth variable, z_t, to the structural oil market model (2.2). Due to data availability, the sample of this exercise starts in January 1980. Figure 2.6 presents the results.

[20] The IFO-BCS is a survey at the product level, so that these numbers do not exactly correspond to firms. The reduced number at the end of the sample reflects the declining weight of the manufacturing sector for the aggregate economy.

[21] We use the monthly IFO-BCS question concerning expected development of export trade. If a firm states that it does not export, then we regard this firm as a non-exporter in this month.

Figure 2.6: Responses of Exporting and Non-exporting Firms to Structural Oil Shocks

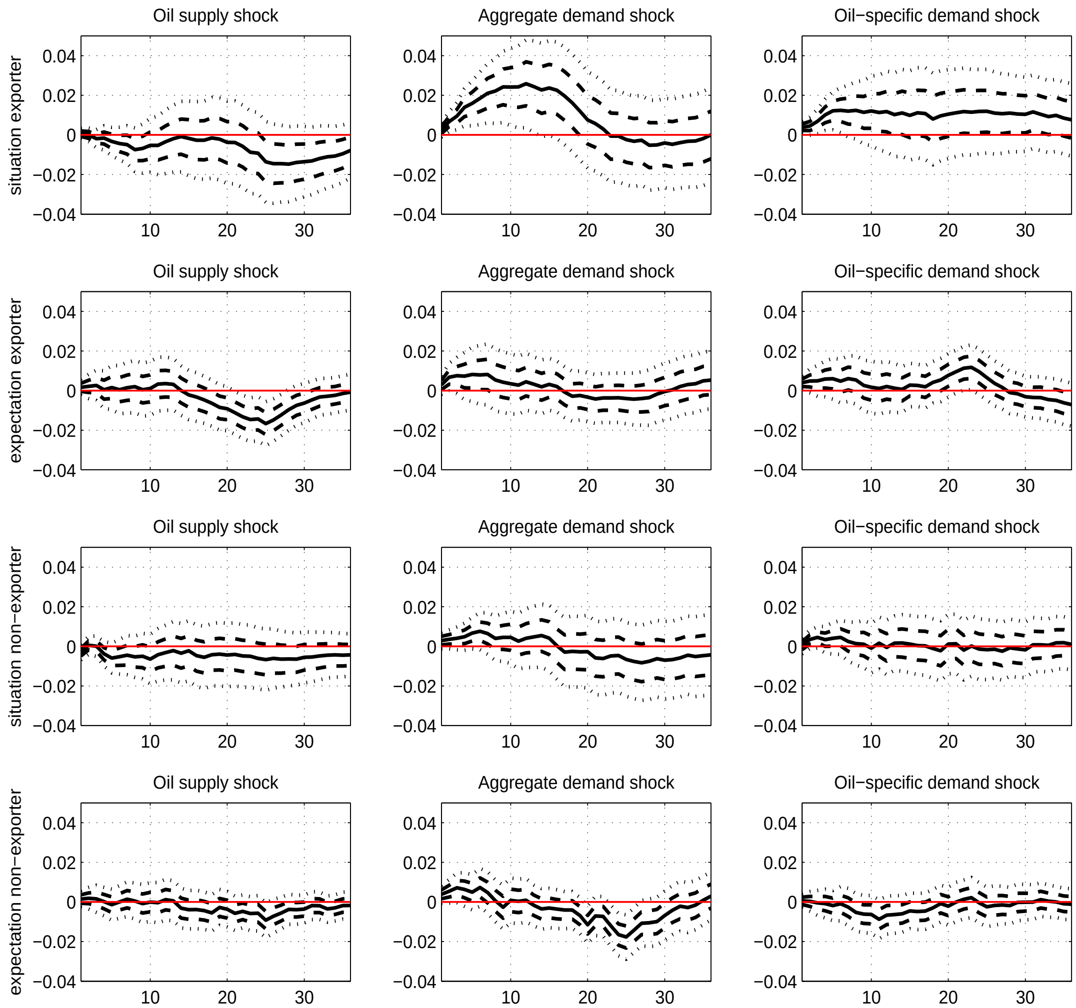

Notes: see notes to Figure 2.4. The time dimension (horizontal axis) is measured in months. The upper two panels show the reactions for the exporting firms and the lower two panels provide evidence for the non-exporting firms. We consider balances of the questions concerning expected and current business situation (Q1 and Q2). Balances are defined as the difference between the fraction of positive responses and the negative responses for the respective question. We use accumulated impulse responses for the expected business situation (Q1) as it defines a change in business situtation.

The upper two panels show the reactions of the current and expected business situation for the exporting firms. The lower two panels display the same for the non-exporting firms. There are two main results: first, the business situation of exporting firms increases significantly after both an expansionary aggregate demand shock and a positive oil-specific demand shock, while non-exporting firms report almost no change. Second, the business expectation of non-exporting firms deteriorates significantly at the one-standard-error band after an oil-specific supply shock, whereas exporting firms are rather optimistic. This suggests that only firms serving the domestic market anticipate a weakening in demand. In summary, these results support our view that it is the export industry which makes a difference. Export firms benefit from the effects of favorable international price movements and shifts in global demand towards German products.

2.5.3 Robustness Checks

In this section we provide three robustness checks: first, we include an indicator of German price competitiveness in our baseline model. This variable ensures that the structural oil price shocks are orthogonal to independent movements in the real exchange rate which could distort our results. Second, we replace in our baseline specification the real price of oil by the nominal WTI oil price in US dollars and the indicator of global activity by the industrial production index of the OECD countries plus the six major non-member economies. The use of different global activity and oil price variables serves as a sensitivity analysis. Third, we question the identification assumption concerning the oil-specific demand shock. Specifically, we replace the zero restriction of aggregate activity with respect to an oil-specific demand shock by a plausible impact elasticity. This elasticity is derived by employing a structural VAR model that is identified with sign restrictions and further inequality assumptions as proposed by Kilian and Murphy (2012). For all three exercises we obtain the same set of results: a partially significant decline in German manufacturing production after a contractionary oil supply shock and a highly significant increase of production after both an aggregate demand and an oil-specific demand shock.

The upper panel of Figure 2.7 depicts the impulse response of German manufacturing to the structural oil price shocks orthogonalized with respect to terms of trade movements.[22] Evidently, the dynamic responses are qualitatively and quantitatively very similar to our baseline findings. Thus, our results are robust with respect to controlling for independent real exchange rate movements.

The second check addresses the variable selection concerning the oil market block. We replace in our baseline specification the real price of oil by the nominal WTI oil price in US dollars and the indicator of global activity by the industrial production index of the OECD countries plus the six major non-member economies. By using these two variables we follow Fukunaga et al. (2010).

While there is no disagreement that in most economic theories the real price of oil is the preferred specification, one could motivate the use of the nominal oil price by the argument

[22] Using instead an indicator of price competitiveness does not change our results.

Figure 2.7: Robustness Checks

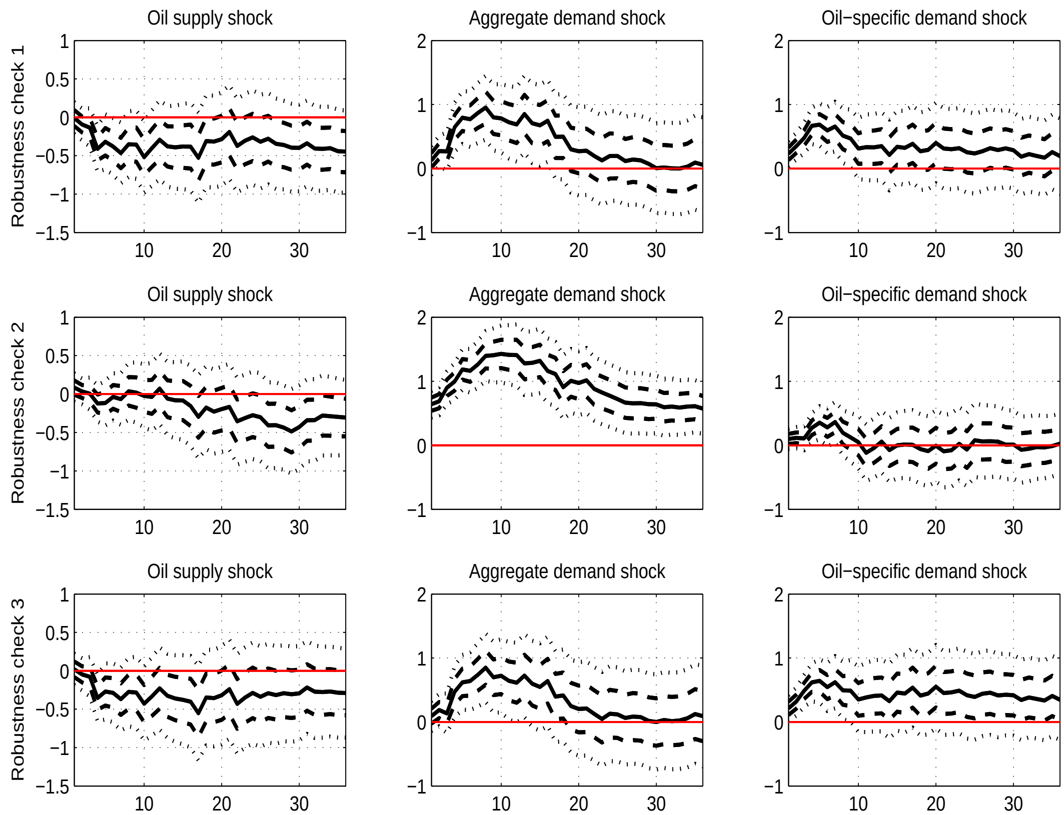

Notes: see notes to Figure 2.4. The time dimension (horizontal axis) is measured in months. In Robustness check 1 we extend our baseline model with the German terms of trade. In Robustness check 2 we use the WTI oil price and an industrial production index as proxy for aggregate activity. In Robustness check 3 we assume an impact response of real activity to an oil-specific demand shock of -0.87.

that German economic agents presumably respond more to changes in the more visible WTI oil price rather than the U.S. refiner's acquisition cost of crude oil deflated by the U.S. CPI.[23] In addition, using the aggregate industrial production of the OECD countries plus the six major non-member economies, including China and India, we provide another proxy for aggregate activity. Kilian (2009b) names three reasons why his proposed dry cargo single voyage ocean freight rate index is superior to this industrial production index: first, it is difficult to measure each country's contribution to global real economic activity. Second, technological change could break the link between industrial production and the demand for industrial commodities. Third, there exists a lack of suitable monthly data for real economic

[23] Kilian and Vigfusson (2011a) cast doubts on this behavioral argument that relies heavily on Hamilton (2011). Note, however, that our aim is not to contribute to the recent debate between Kilian and Vigfusson (2011a) and Hamilton (2011) whether the use of the nominal price is completely inappropriate.

activity in many countries. The dry cargo single voyage ocean freight rate index addresses all these issues and is therefore used in our baseline analysis. Nonetheless, it seems instructive to replace this index by the aggregate industrial production index in a robustness check.

The middle panel of Figure 2.7 shows the results for German manufacturing. In comparison to our baseline results we observe a more pronounced production increase after an aggregate demand shock and a less pronounced production increase after an oil-specific demand shock. It seems that the global demand shock identified within this VAR setup includes components of the oil-specific demand shock identified in our baseline specification. But overall, our main statements remain unaltered.

As a final check, we review the identification of the oil-specific demand shock. As mentioned by Kilian and Murphy (2012) it is surprising that the immediate oil price hike after an oil-specific demand shock induces an expansion rather than a contraction in real activity. To address this issue we reconsider the assumption that oil-specific shocks do not affect real activity within the same month. Specifically, we derive a plausible estimate for the impact elasticity of aggregate activity to an oil-specific demand shock by proceeding in three steps: first, we identify a structural VAR model with sign restrictions and further inequality assumptions proposed by Kilian and Murphy (2012). Then, we derive the median impact response of real activity with respect to an oil-specific demand shock. Finally, we use this number as the impact elasticity of aggregate activity to an oil-specific demand shock and impose otherwise the same identification assumptions as described (2.3).

Table 2.3 summarizes the sign restrictions that are used to identify the structural oil market model (2.1).[24]

Table 2.3: Sign Restrictions (Restriction Period of 1 Month)

	Oil supply shock	Aggregate demand shock	Oil-specific demand shock
Oil production	≤ 0	≥ 0	≥ 0
Real activity	≤ 0	≥ 0	≤ 0
Real oil price	≥ 0	≥ 0	≥ 0

These restrictions imply that a contractionary oil supply shock raises the price of oil and reduces real activity and oil production. A positive oil-specific demand shock induces a contraction of real activity while the price and the supply of oil increase. A positive global demand shock raises oil production, real activity and the real oil price. We impose that the sign restrictions need to hold for the impact period. We further assume that after an oil

[24] The sign restrictions are the same as in the analyzes of Baumeister and Peersman (2008, 2009) and Peersman and van Robays (2009).

price hike the short-run elasticity of oil production with respect to the real oil price is not larger than 0.0258. This upper bound is imposed for the aggregate demand and oil-specific demand shock. The last restriction concerns the impact response of aggregate activity after an oil-specific demand shock. This response is not allowed to be lower than -1.5. With these assumptions we employ the sign restriction procedure outlined in Uhlig (2005). Our results are based on 200 accepted draws. Figure A1 and Table A1 in Appendix A1 provide the results of this exercise. But more importantly, the median impact response of real activity with respect to an oil-specific demand shock is -0.87. This number replaces the zero restriction imposed in our baseline specification.

The lower panel of Figure 2.7 depicts the resulting impulse responses to this identification scheme. Somewhat surprisingly, the restriction that real activity contracts after an oil-specific demand shock does not change the response of German manufacturing with respect to an oil-specific demand shock. Therefore this check supports our main findings.

2.6 How Strongly Did the 2007/08 Oil Price Hike Contribute to the Recession in Germany?

In this section we analyze the economic consequences of the 2007/08 oil price hike for Germany. To do so, we use quarterly regression (2.4) described above to study the cumulative effects of the structural oil price shocks on German GDP during the years 2007 until 2010. Figure 2.8 shows the historical contributions of all three structural shocks to the observed variation in German GDP growth. We also depict the total effects of all structural shocks on GDP in the bottom right panel. The dashed lines display the actual demeaned GDP quarterly growth rates and the solid lines show the cumulative effects of the respective structural shock.

German GDP growth was largely unaffected by oil supply shocks in this time period. Oil-specific demand shocks contributed slightly positively to German GDP growth in the first three quarters 2008. However, these positive effects were not able to offset the negative effects originating from the aggregate demand shocks, which had a negative effect on GDP growth well before Lehman Brothers collapsed in September 2008. To assess how much of the observed GDP decline after September 2008 is due to the rapid drop in global demand we implement a counterfactual analysis. Specifically, we set all structural oil price shocks since July 2008, the month of the oil price peak, equal to zero. Note that in this case the real oil price would have been fallen from July 2008 to July 2009 by slightly more than 10 percent which is much less than actually observed. The dotted lines display the resulting cumulative effects. Clearly, the effect consisting only of structural oil shocks that occurred

Figure 2.8: Cumulative Effect of the Structural Oil Shocks on German GDP Growth

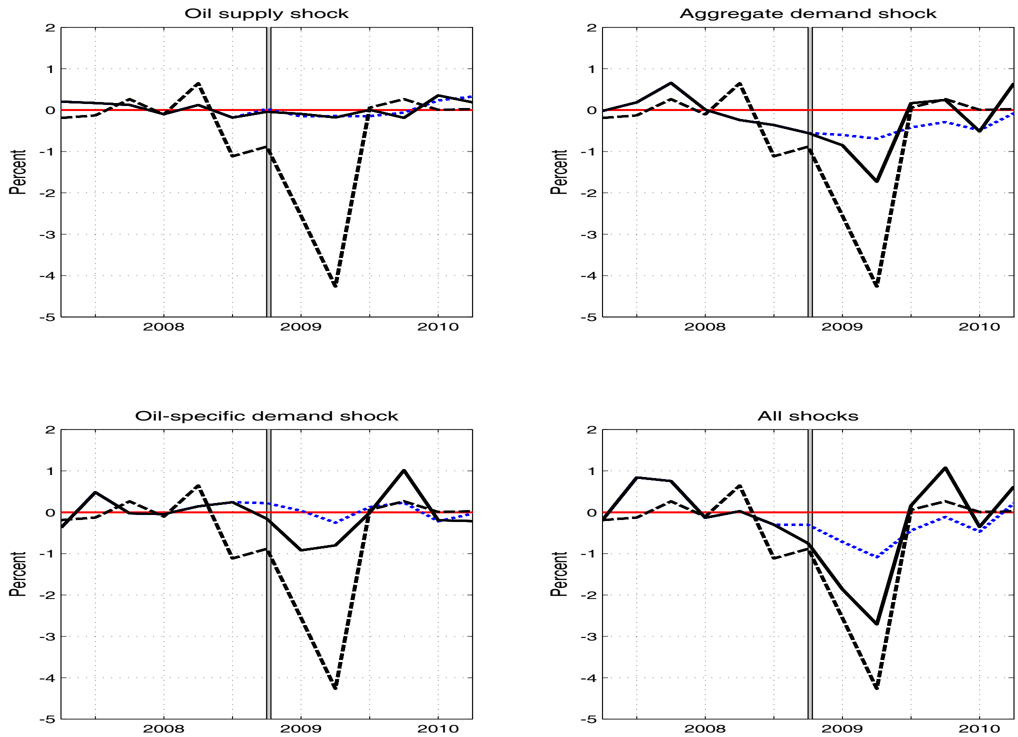

Notes: Estimates are based on model (2.4). The time dimension (horizontal axis) is measured in quarters. The dashed lines denote the actual demeaned GDP growth rates. The solid lines show the cumulative effects of the respective structural shock. The dotted lines display the cumulative effects of the counterfactual analysis in which all structural oil price shocks from July 2008 are set equal to zero. The grey vertical lines represent the date of the collapse of Lehman Brothers.

before July 2008 had a sizeable negative effect on GDP growth at the beginning of 2009. But the sequence of negative aggregate demand shocks that have shown up at the end of 2008 explain a larger part of this sharp drop.

In total, the oil price hike in 2007/08 had sizeable negative effects on GDP. We draw this conclusion by considering the counterfactual response of all shocks in the bottom right panel. In our counterfactual analysis we estimate a cumulative reduction of German GDP of 2.3 percent in the year 2009 as a result of the structural shocks that triggered the oil price hike in 2007/08. This effect alone suffices to produce a recession. Therefore, we show that this oil price hike notably contributed to the subsequent recession in Germany. Even though the main explanations for the dramatic fall in German production can be rather found in the collapse of global demand and other factors, presumably negative shocks originating from the finance sector, that are not captured by our structural oil price shocks.

2.7 Conclusion

In this paper we have shown that demand and supply shock driven oil price hikes have different impacts on the German economy. While we find that consumption always reacts negatively, the dynamic responses of exports and gross investment depend on the type of the structural oil price shock. In the cases of the oil-specific and the aggregate demand shock we observe that favorable international price movements and shifts in global demand towards German export goods initially outweigh the contractionary effects on consumers' expenditure and, therefore, lead to an increase in GDP. Even though strong oil price surges do not burden German manufacturing, that primarily produces investment and export goods, we find that their effects on domestic demand become rather negative over time.

Concerning the economic consequences of the 2007/08 oil price we find that the sustained sequence of positive aggregate demand shocks, that triggered the oil price hike in 2007/08, led to a 2.3 percent reduction in German GDP in 2009. This finding resembles the U.S. result discussed by Hamilton (2009) and Kilian (2009a) and is somewhat surprising, as the German economy is much more export orientated than the U.S. economy. However, at the end the negative effects on consumption played also in Germany the most dominant role. We thus provide evidence that this oil price hike made a notable contribution to the subsequent recession in Germany. At the same time, this result is not in conflict with the view that the oil price was not the major driver of the 2009 recession.

Acknowledgements

I am indebted to Kai Carstensen and Georg Paula who are co-authors of Chapter 2. We thank Marc Gronwald and Sarah Lein for their discussions and two anonymous referees for helpful comments. We are grateful to seminar/meeting participants at the CESifo Conference on Macroeconomics and Survey Data (2010), ESEM (Oslo) and VfS (Frankfurt). We thank Gebhard Flaig, Steffen Henzel, Hans-Ruediger Moeller, Wolfgang Ruppert and Sigrid Stallhofer for helping us collect the data.

Bibliography

Baumeister, C., and G. Peersman (2008): "Time-Varying Effects of Oil Supply Shocks on them U.S. Economy," *mimeo*.

——— (2009): "Sources of the Volatility Puzzle in the Crude Oil Market," *mimeo*.

Becker, S., and K. Wohlrabe (2008): "Micro Data at the Ifo Institute for Economic Research - The 'Ifo Business Survey' Usage and Access," *Schmollers Jahrbuch*, 128, 307–319.

Blanchard, O. J., and J. Gali (2009): "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so Different from the 1970s?," in *International Dimensions of Monetary Policy*, ed. by J. Gali, and M. Gertler, pp. 373–428.

Bresnahan, T. F., and V. A. Ramey (1993): "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry," *American Economic Review Papers and Proceedings*, 83(2), 213–218.

Cunado, J., and F. P. de Gracia (2003): "Do Oil Price Shocks Matter? Evidence for Some European Countries," *Energy Economics*, 25, 137–154.

Davis, S. J., and J. Haltiwanger (2001): "Sectoral job creation and destruction to oil price changes," *Journal of Monetary Economics*, 48(1), 465–512.

Fukunaga, I., N. Hirakata, and N. Sudo (2010): "The Effects of Oil Price Changes on the Industry-Level Production and Prices in the U.S. and Japan," Working Paper 15791, National Bureau of Economic Research.

Gonçalves, S., and L. Kilian (2004): "Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form," *Journal of Econometrics*, 123(1), 89–120.

Hamilton, J. D. (1983): "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, 91(2), 228–248.

——— (2009): "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 1, 215–261.

——— (2011): "Nonlinearities and the Macroeconomic Effects of Oil Prices," *Macroeconomic Dynamics (forthcoming)*.

Herrera, A. M., L. G. Lagalo, and T. Wada (2011): "Oil Price Shocks and Industrial Production: Is the Relationship Linear?," *Macroeconomic Dynamics (forthcoming)*.

Hicks, B., and L. Kilian (2011): "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?," *mimeo*.

Hooker, M. A. (1996): "What happened to the oil price-macroeconomy relationship?," *Journal of Monetary Economics*, 38(2), 195–213.

Jiménez-Rodríguez, R. (2008): "The Impact of Oil Price Shocks: Evidence from the Industries of Six OECD Countries," *Energy Economics*, 30, 3095–3108.

——— (2011): "Macroeconomic Structure and Oil Price Shocks at the Industrial Level," *International Economic Journal*, 25, 173–189.

Jiménez-Rodríguez, R., and M. Sánchez (2005): "Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries," *Applied Economics*, 37, 201–228.

Kilian, L. (2008a): "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries," *Journal of the European Economic Association*, 6(1), 78–121.

——— (2008b): "The Economic Effects of Energy Price Shocks," *Journal of Economic Literature*, 46(4), 871–909.

——— (2009a): "Comment on 'Causes and Consequences of the Oil Shock of 2007-08' by James D. Hamilton," *Brookings Papers on Economic Activity*, 1, 267–278.

——— (2009b): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99(3), 1053–1069.

Kilian, L., and P. Edelstein (2009): "How Sensitive are Consumer Expenditures to Retail Energy Prices?," *Journal of Monetary Economics*, 56(6), 766–779.

Kilian, L., and D. P. Murphy (2012): "Why Agnostic Sign Restrictions Are Not Enough: Understanding The Dynamics Of Oil Market VAR Models," *Journal of the European Economic Association (forthcoming).*

Kilian, L., and C. Park (2009): "The Impact of Oil Price Shocks on the U.S. Stock Market," *International Economic Review*, 50(4), 1267–1287.

Kilian, L., A. Rebucci, and N. Spatafora (2009): "Oil Shocks and External Balances," *Journal of International Economics*, 77(2), 181–194.

Kilian, L., and R. J. Vigfusson (2011a): "Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?," *Quantitative Economics (forthcoming).*

——— (2011b): "Nonlinearities in the Oil Price-Output Relationship," *Macroeconomic Dynamics (forthcoming).*

Lee, K., and S. Ni (2002): "On the dynamic effects of oil price shocks:a study using industry level data," *Journal of Monetary Economics*, 49(2), 823–852.

Peersman, G., and I. van Robays (2009): "Oil and the Euro area economy," *Economic Policy*, 24, 603–651.

Ramey, V. A., and D. J. Vine (2011): "Oil, Automobiles and the U.S. Economy: How Much Have Things Really Changed?," *NBER Macroeconomics Annual (forthcoming).*

Uhlig, H. (2005): "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 52(2), 381–419.

Appendix

A1 Model with Sign Restrictions and Additional Restrictions Imposed

Figure A1: Responses to Structural Shocks to the Global Oil Market Model Using Sign Restrictions and Additional Restrictions Imposed

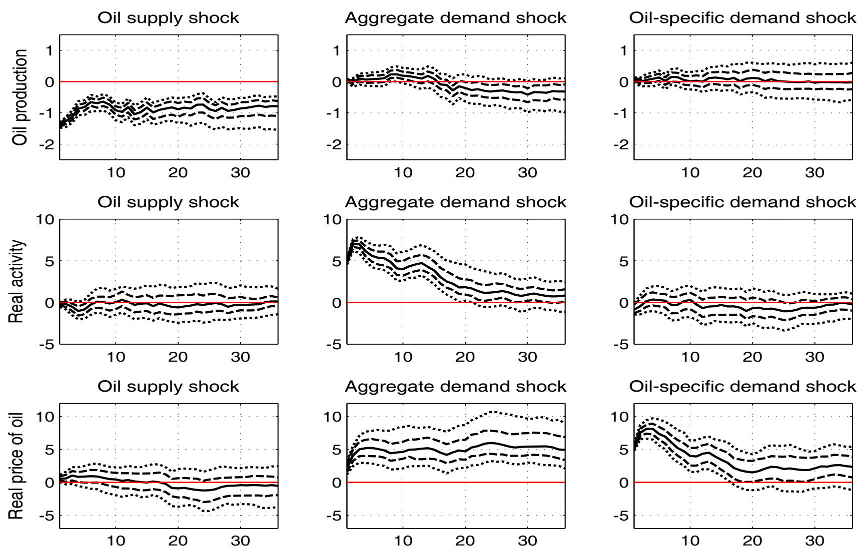

Notes: This figure is a replication of Figure 10 in Kilian and Murphy (2012) with an updated data set. The time dimension (horizontal axis) is measured in months. Results are based on 200 accepted draws of the model with sign restrictions and additional restrictions imposed. The solid lines refer to the median of the accepted draws at each horizon. The 68- and 95 percent confidence bands are the 16th/84th and 2.5th/97.5th percentiles of the accepted draws for each horizon.

Table A1: Impact Matrix—Sign Restrictions Combined with Additional Restrictions Imposed

	Oil supply shock	Aggregate demand shock	Oil-specific demand shock
Oil production	-1.3994	0.0246	0.0687
Real activity	-0.1563	4.9735	-0.8713
Real oil price	0.5514	2.2677	5.3249

Notes: see notes to Figure A1.

Chapter 3

Firms' Optimism and Pessimism: Evidence From the IFO Survey

Abstract[1]

Are firms' expectations biased? Does it matter? We use micro data on firms' production and investment expectations from the German IFO Business Climate Survey and the IFO Investment Survey and compare them to realization data from the same surveys. We then construct series of *quantitative* firm-specific expectation errors. We find that depending on the exact definition at least 6 percent and at most 34 percent of firms consistently over- or underpredict their one-quarter-ahead upcoming production. Almost 15 percent of firms over- or underpredict their upcoming yearly investment. In a simple frictionless neoclassical heterogeneous firm model these expectational biases lead to factor misallocations that cause welfare losses that in the worst case are comparable to conventional estimates of the welfare costs of business cycles fluctuations. In more conservative calibrations the welfare losses are even smaller.

[1]This chapter is based on joint work with Rüdiger Bachmann. It is based on our paper "Firms' Optimism and Pessimism: Evidence From the IFO Survey," mimeo, 2011.

3.1 Introduction

Are firms' expectations biased? Firms, based on expectations about their future business situation, decide about the allocation of an economy's capital stock and labor supply. If firms' expectations are biased, an economy's factor allocation is likely to be suboptimal compared to a case with only rational firms.

How large are the welfare losses from the resulting factor misallocations? This paper, using business survey data from Germany, is an attempt to provide quantitative answers to these questions. We find welfare losses are likely to be small, given the extent of expectational biases in the data.

Little is known empirically about firms' expectation formation.[2] There is an accounting literature (e.g. McDonald, 1973, Firth and Smith, 1992, Brown et al., 2000), which finds that managers tend to bias their public earnings and dividend forecasts upward, presumably for strategic reasons to attract investors. Malmedier and Tate (2005, 2008) using data on risk exposure in firms' investment strategies to measure CEO overconfidence, investigate the relationship between CEO overconfidence and corporate investment. These literatures rely on either publicly announced expectations or ex-post behavior to measure firms' expectations. Business survey data on expectations and realizations are less likely to suffer from strategic forecasting behavior, as they are highly confidential micro data and can only be accessed under strict non-disclosure agreements, if at all. These survey data have been used in the literature to study rationality and unbiasedness of firms' expectations: Anderson et al. (1954) conduct a very early study on qualitative expectation errors using the first few installments of the IFO Business Climate Survey. Nerlove (1983) uses German (from IFO) and French data on firms' expectations about idiosyncratic firm variables (such as prices, demand, etc.). Tompkinson and Common (1983) study expectations about idiosyncratic firm variables in the U.K. manufacturing sector and Zimmermann and Kawasaki (1986), using IFO price expectation and realization data, test whether firms are rational about the development of the market prices of their own commodities. All these studies have in common that they usually find some degree of deviation from rationality.

Progress in the empirical literature on firms' expectation formation has been slow be-

[2] Expectation formation of households is somewhat better understood: for instance, Souleles (2004), using data from the Michigan Survey of Consumers, and Bovi (2009), using the harmonized European consumer surveys, present evidence that household expectations are not rational. Agents systematically assess their current situation overcritically and form their expectations overoptimistically. On the theory side, Brunnermeier and Parker (2005) present a model where agents optimally bias their expectations about the future upwards to increase their expected life time utility. Jaimovich and Rebelo (2007) study a business cycle model with overoptimism with respect to the realization of investment-specific shocks, without, however, using micro data for calibration. Hassan and Mertens (2011) argue that small expectation errors of stock market investors can have first order effects on welfare.

cause of formidable data requirements. Ideally, researchers need high-frequency quantitative expectation and realization data on firm-specific variables for a large (and representative) cross-sections of firms, such as: "By how much do you expect your production to grow over the next quarter? By how much did your production grow during the preceding quarter?" These data need to be available over long time horizons to construct firm histories of expectation errors and to ensure that specific cyclical episodes do not bias results; after all, in booms we should expect to see upward expectation errors.

However, high frequency business survey data about idiosyncratic expectations and realizations are usually qualitative,[3] indeed trichotomous, in nature: "we expect our production to increase, decrease, stay the same over the next three months". While useful, qualitative information has its limits, in particular when forecasting errors need to be aggregated over time in order to measure the long-run average forecasting errors of firms and possible biases therein. How does one aggregate a qualitative forecasting error of $+1$ (up) today and -1 (down) tomorrow?[4] Therefore, Müller (2010) uses quantitative expectations about plants' sales and employment development over a year from the annual IAB Establishment Panel in Germany to measure whether firms are overoptimistic or overpessimistic. He finds strong evidence for the existence of both types of firms. However, the IAB Establishment Panel is still rather short (starting in 1993) and of low frequency.

In our analysis we use firm-level micro data from two IFO business surveys – the IFO Business Climate Survey (IFO-BCS) and the IFO Investment Survey (IFO-IS) – that allow us to construct quantitative expectation errors for firms' production and investment. We note that we still cannot measure surprises with respect to truly exogenous driving forces, such as technology or demand shocks. However, to the best of our knowledge there is no business survey in the world who would ask these questions directly. Nevertheless, surprises with respect to endogenous variables, such as production and investment, can be informative about truly exogenous surprises when viewed through the lens of a structural model of the firm. We will make use of this insight in the second half of the paper.

The IFO-BCS is a monthly qualitative business survey that is supplemented on a quarterly basis with quantitative questions. Under certain assumptions, we can combine the qualitative three-month ahead production outlook from the monthly survey with the quantitative change in percentage capacity utilization from the quarterly supplement to compute

[3] De Leuuw and McKelvey (1981) study quantitative annual expectation data about aggregate prices from a BEA survey of business expenditures on plant and equipment in the U.S. and find that firms do not have rational expectations.

[4] This is a problem, even when qualitative expectation data predict quantitative ex-post data rather well, which is what Lui et al. (2008) find, using survey data from the 'Confederacy of British Industry' business survey and ex-post administrative data.

idiosyncratic quarterly, one-quarter-ahead production expectation errors. We do this for the manufacturing part of the IFO-BCS from 1980 on and thus construct a panel of quarterly production expectation errors over thirty years. The IFO-IS is a biannual quantitative business survey across all sectors of the economy which asks firms in the spring and fall of every year about their investment plans for the upcoming year and the actual investments undertaken. Here the data goes back to 1970, so we can compute a panel of investment expectation errors over forty years. These (unbalanced) panel data sets allow us to analyze firm histories of expectation errors which is an unique advantage compared to previous studies.

Limiting our analysis to firms with at least eight years of observations we compute long-run averages of firms' expectation errors and study their distributions. To classify firms as optimists or pessimists, we test for each firm whether its average expectation error is significantly different from zero. Then we define optimistic firms as those firms which feature a negative average expectation error that is significantly different from zero. Analogously, pessimistic firms are defined as those firms with a positive average expectation error that is significantly different from zero. At least 6 percent and at most 34 percent of firms consistently over- or underpredict their one-quarter-ahead upcoming production. The optimist firms, for example, overpredict their production by 3.4 percent (5.8 percent) on average. Roughly 15 percent of firms over- or underpredict their upcoming yearly investment. Here the optimists overpredict by 37.3 percent on average.

To gauge the implications of these expectational biases we perform a simple welfare calculation. We use a frictionless heterogeneous firm model where firms decide about their factor demands before they know their idiosyncratic productivity levels. We calibrate the fractions of optimistic and pessimistic firms and the extent of their expectational biases to the distributional properties of production expectation errors in the IFO-BCS. Overoptimistic firms hire too many workers and build up capital stocks that are too high. Overpessimistic firms do not demand enough inputs. We then compare the welfare in an economy which is populated by firms with a distribution of production expectation errors that approximates the one in the data to a world that is only populated by firms with zero long-run expectation errors. We robustly find that the welfare losses from expectational errors are small, probably smaller even than conventional estimates of the welfare costs of business cycles. Firms seem to be pretty good in predicting their business environment.

The remainder of this paper is structured as follows. The next section describes the IFO-BCS and the construction of the production expectation errors from it. Section 3.3 performs a similar analysis for investment expectation errors for the IFO-IS. Section 3.4 introduces a simple heterogeneous firm model. Section 3.5 discusses the welfare results. Section 3.6 provides a series of robustness checks. Section 3.7 concludes.

3.2 Evidence from the IFO Business Climate Survey

3.2.1 The IFO Business Climate Survey

The IFO Business Climate index is a much-followed leading indicator for economic activity in Germany. It is based on a firm survey which has been conducted since 1949 and, therefore, is one of the oldest and broadest monthly business confidence surveys available (see Becker and Wohlrabe, 2008, for details). Due to longitudinal consistency problems in other sectors and the availability of micro data in a processable form we limit our analysis to the manufacturing sector from 1980 until the present. From 1991 on, the sample includes East–German firms.

One of the IFO-BCS's main advantages is the high number of survey participants. The average number of respondents at the beginning of our sample is approximately 5,000; towards the end it is about 2,500.[5] Participation in the survey is voluntary and confidential. Thus, there is little incentive for firms to provide overoptimistic forecasts as a signal to investors. There is some fraction of firms that are only one time participants. However, conditional on staying two months in the survey, most firms continue on and this allows us to construct an unbalanced panel data set of expectation errors. For our narrow, very conservative definition of expectation errors the final baseline panel consists of 695 firms with at least 32 quarterly observations each; for a broader definition we follow 3,859 firms for again 8 years at least.

3.2.2 Construction of Quantitative Production Expectation Errors

To construct firms' quantitative production expectation errors we would ideally need the following quantitative information about production expectations and realizations: "By how much do you expect your production to grow over the next quarter? By how much did your production grow in the preceding quarter?" To the best of our knowledge there is no firm survey that asks these questions for a long time horizon and repeatedly at underyearly frequencies. However, the quantitative quarterly supplement of the IFO survey allows us to construct - under certain assumptions - quantitative production expectations and quantitative production realizations. We are thus able to construct a panel of quarterly production expectation errors for the last thirty years.

[5]The IFO-BCS is a survey at the product level, so that these numbers do not exactly correspond to firms.

Specifically, we use the following supplementary question about capacity utilization to compute production changes:[6]

Q 1 *"Supplementary Question: The utilisation of our production equipment for producing XY (customary full utilization = 100) currently amounts to..%."*

30	40	50	60	70	75	80	85	90	95	100	more than 100 % namely

We start from the following production relationship of an individual firm i:

$$y_{i,t}^{act} = u_{i,t} y_{i,t}^{pot}, \tag{3.1}$$

where $y_{i,t}^{act}$ denotes the firm's actual output, $y_{i,t}^{pot}$ its potential output level and $u_{i,t}$ the level of capacity utilization. Only $u_{i,t}$ is directly observable in the IFO-BCS. Taking the natural logarithm and the three-month difference, we get:[7]

$$\Delta \log y_{i,t}^{act} = \Delta \log u_{i,t} + \Delta \log y_{i,t}^{pot}. \tag{3.2}$$

Under the assumption that potential output remains constant, $\Delta \log y_{i,t}^{pot} = 0$, percentage changes in actual output can be recovered from percentage changes in capacity utilization. To implement this idea we restrict the analysis to firms of which we can reasonably expect that they did not change their production capacity in the preceding quarter, making use of the following two questions in the IFO-BCS:

Q 2 *"Expectations for the next three months: Employment related to the production of XY in domestic production unit(s) will probably increase, roughly stay the same, decrease."*

Q 3 *"Supplementary Question: We evaluate our technical production capacity with reference to the backlog of orders on books and to orders expected in the next twelve months as more than sufficient, sufficient, insufficient."*

[6] Here we provide a translation, for the German original see Appendix A1.

[7] Time intervals are months. For us to construct an expectation error in t, we need an observation for capacity utilization in t and $t-3$.

Given hiring frictions in the labor market we view a firm's expectation, stated in $t-3$, that its employment level will remain the same in the next three months as highly indicative that its productive capacity did not change between $t-3$ and t. Similarly, given capital adjustment frictions a firm's statement, again in $t-3$, that its technical production capacity is sufficient for the future incoming orders suggests that this firm has no reason to change its production capacity in the near future. To be conservative we require a firm satisfy both criteria in $t-3$ for us to assume that its production capacity has not changed between $t-3$ and t. In this case, we use the quarterly percentage change in capacity utilization in t as a proxy for the quarterly percentage change in production in t. The existence of non-convex or kinked adjustment costs for capital and labor adjustment as well as time to build (see Davis and Haltiwanger, 1992, as well as Doms and Dunne, 1998) make this a reasonable assumption.

To derive production expectation errors we also need information on firms' production expectations. This allows us to compute production *surprises* out of mere production *changes*. In the IFO-BCS firms report only qualitative production expectations:

Q 4 *"Expectations for the next three months: Our domestic production activities with respect to product XY will (without taking into account differences in the length of months or seasonal fluctuations) increase, roughly stay the same, decrease."*

Qualitative expectations have a built-in asymmetry in the sense that the middle category also constitutes a quantitative expectation, zero change, whereas the increase and decrease category conveys no quantitative information. We therefore proceed in two steps. First, we consider only firms whose answer to Q4 is that their production level, $y_{i,t}^{act}$, will not change in the next three months. Under the assumption that $y_{i,t}^{pot}$ remains constant over this time period, all $\Delta \log u_{i,t}$ are automatically expectation errors. In a second step we extend our analysis to arbitrary qualitative production expectations. This will give us a broader picture of expectation errors, albeit with the added cost of more assumptions.

We also clean our sample from firm-quarter observations with extreme capacity utilization outliers, i.e. those that exceed 150%, and from firm-quarter observations with inconsistent statements. To determine the latter we consider the following monthly qualitative IFO-BCS question concerning actual production changes in the months t, $t-1$, $t-2$:

Q 5 *"Trends in the last month: Our domestic production activities with respect to product XY have (without taking into account differences in the length of months or seasonal fluctuations) increased, roughly stayed the same, decreased."*

We drop all observations as inconsistent in which firms report a strictly positive (negative) change in $\Delta \log u_{i,t}$ and no positive (negative) change in Q5 in the last 3 months. For firms that report $\Delta \log u_{i,t} = 0$, we proceed as follows: Unless firms in Q5 either answer three times in a row that production did not change, or they have at least one "Increase" *and* one "Decrease" in their three answers, we drop them as inconsistent.

3.2.2.1 Quantitative Expectation Errors under Constant Production Expectations

If the production capacity can be assumed not to have changed in the preceding quarter, and if no change in production was expected three months prior, a change in capacity utilization, $\Delta \log u_{i,t}$, is also a production expectation error of firm i in month t. We first restrict our analysis to the subset of firm-quarter observations that satisfy these assumptions. For this case, Figure 3.1 illustrates the move from capacity utilization changes to production expectation errors.

Figure 3.1: Link between Capacity Utilization and Production Expectation Errors

Prerequisite: Firm i passes the outlier and inconsistency test (Q1 and Q5)

Firm i has an observation for a change in capacity utilization $\Delta \log u_{i,t}$ (Q1)

↓ No change in potential output $\Delta \log y^{pot}_{i,t}$

Firm i answers "Constant employment expectations" (Q2)

and "Technical production capacity is sufficient" (Q3) in $t-3$

$\Delta \log u_{i,t}$ is a production change $\Delta \log y^{act}_{i,t}$

↓ No expected production change

Qualitative production expectations in $t-3$ are constant (Q4)

$\Delta \log u_{i,t}$ is a quantitative production expectation error

Notes: The time dimension t is measured in months.

Figure 3.2 illustrates the exact timing of the questions in the IFO-BCS that we use to compute production expectation errors. As a first pass we consider only firms which state in period $t-3$ that their production level, employment level and technical production capacity will remain the same in the next three months. Then we compute $\Delta \log u_{i,t}$ three months later in t. These $\Delta \log u_{i,t}$ constitute our narrow definition of production expectation errors. We denote them by $FE\_BCS^{narrow}_{i,t}$.

Figure 3.2: Derivation of Production Expectation Errors under Constant Production Expectations - Timing

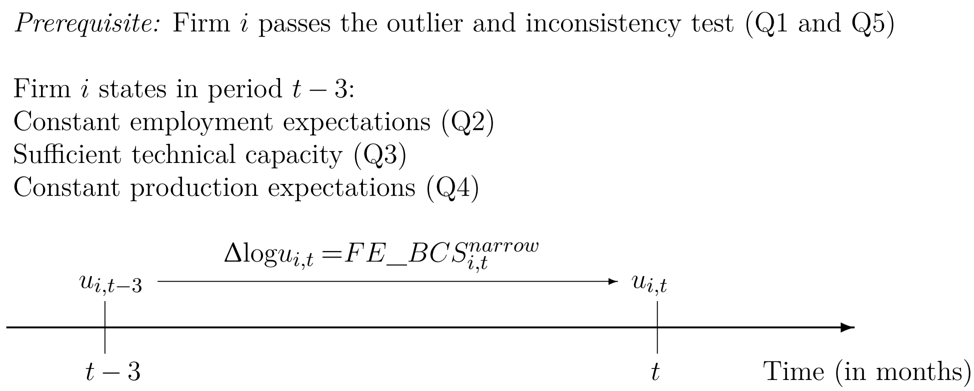

3.2.2.2 Quantitative Expectation Errors under General Production Expectations

The derivation of quantitative production expectation errors for firms with increasing or decreasing qualitative production expectations in Q4 requires additional assumptions. We admit at the outset that these assumptions may not be too palatable. However, we view this as an attempt to measure firms' quantitative production expectation errors as best as we can, *given the limited data available.* We take the following four steps. First, we define a qualitative index of production changes. Specifically, we compute a firm-specific activity variable, $REALIZ_{i,t}$, as the sum of the increase-instances minus the sum of the decrease-instances in question Q5 over the last three months going backward from t. In a second step, we use these qualitative production changes to determine qualitative expectation errors with question Q4. Then, using the conditions about employment expectations and adequacy of technical capacity, we map qualitative production changes into quantitative production changes. In a final step, we convert these quantitative production changes into quantitative production expectation errors for all firms that pass the aforementioned outlier and inconsistency tests.

The basic idea is to assign firms with large *qualitative* production expectation errors – for example a firm expecting its production to go up over the next three months, but then reporting only production declines – large *quantitative* production expectation errors, derived from a mapping between qualitative and quantitative production changes. The expectation errors for firms with constant production expectations remain the same as in the previous section. We denote this measure of quantitative production expectation errors under general production expectations by $FE\_BCS^{broad}_{i,t}$. Details of the construction can be found in Appendix A2.

3.2.3 Results

We next compute the average firm-specific production expectation error over all time periods for which we have data. We restrict our sample to firms that have at least 32 observations or eight years of expectation errors: $FE\_BCS^{narrow}_{i,t}$ or $FE\_BCS^{broad}_{i,t}$. The average firm-specific expectation errors are denoted by $AFE\_BCS^{narrow}_{i}$ and $AFE\_BCS^{broad}_{i}$.

Table 3.1: Firm-Specific Average Production Expectation Errors (IFO-BCS)

Statistics	$AFE\_BCS^{narrow}_{i}$	$AFE\_BCS^{broad}_{i}$
Obs	695	3,859
Mean	-0.0064	-0.0265
Std.Dev.	0.0161	0.0316
Percentiles		
5th	-0.0339	-0.0859
10th	-0.0242	-0.0672
25th	-0.0119	-0.0418
50th	-0.0030	-0.0212
75th	0.0013	-0.0060
90th	0.0068	0.0063
95th	0.0119	0.0140
# Optimists	37 (5.3%)	1,370 (32.3%)
# Pessimists	5 (0.7%)	49 (1.3%)
# Realists	653 (94.0%)	2,564 (66.4%)

Notes: This table provides a summary of the distributions of $AFE\_BCS^{narrow}_{i}$ and $AFE\_BCS^{broad}_{i}$. In the last three rows we report the number of firms which are classified as optimists, pessimists and realists. We define optimistic firms as those firms which feature a negative average expectation error that is significantly different from zero. Pessimistic firms are defined as those firms with a positive average expectation error that is significantly different from zero.

Table 3.1 displays the distributions of firms' average expectation errors.[8] Note that positive values of $AFE\_BCS^{narrow}_{i}$ or $AFE\_BCS^{broad}_{i}$ indicate that a firm was on average too pessimistic in the sense that its predicted production changes were on average lower than its actual production changes. Especially for $AFE\_BCS^{broad}_{i}$ the distribution is skewed towards negative values and at least 25 percent of all firms have long-run averages of expectation errors which are too optimistic by 4 percent or more. These numbers alone, however,

[8] We show in Figure A3 of Appendix A3 the corresponding histograms of the distributions.

are not sufficient to assess whether a firm has biased expectations. To provide persuasive evidence for expectational biases we need to consider the second moment of firm-specific shocks as well. Firms operate in different economic environments and face different sizes of shocks. Therefore, analyzing only average expectation errors can be misleading.

The panel structure of the IFO-BCS allows us to address this issue. We test for each firm whether its average expectation error is significantly different from zero. To this end, we regress for each firm all its observations of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ on a constant. Then we use two-sided t-tests in order to assess whether the individual firm-specific average expectation error is significantly different from zero.[9] We define optimistic firms as those firms which feature a negative average expectation error that is significantly different from zero. Pessimistic firms are defined as those firms with a positive average expectation error that is significantly different from zero. In the case of $AFE\_BCS_{i}^{narrow}$ this difference from zero is only significant for 6 percent of all considered firms. For $AFE\_BCS_{i}^{broad}$ we classify more than 30 percent as optimists.

For both definitions of forecast errors we see that the distribution of average forecast errors is skewed towards overoptimism. The optimist firms overpredict their production by 3.4 percent (5.8 percent for the broad definition) on average, which corresponds to 2.1 times (1.8 times) the standard deviation of the distribution of firms' production expectation errors. The pessismist firms underpredict their production by 1.9 percent (3.5 percent) on average, which corresponds to 1.1 times (1.2 times) the standard deviation of the distribution of firms' production expectation errors.

3.3 Evidence from the IFO Investment Survey

In this section we use the IFO Investment Survey (IFO-IS) to compute quantitative *investment* expectation errors. One advantage of the IFO-IS is that it directly provides quantitative information on expected and realized investment. A disadvantage, however, is that we can only compute expectation errors at an annual frequency. The IFO-IS is carried out twice a year and we have access to the micro data since 1970. The average number of respondents is roughly 3,000 at the beginning of our sample and declines to 1,500 towards the end. The participating firms provide quantitative information (in Euro) about their investment plans for the current and upcoming year.

Specifically, firms are asked in the spring of each year about their investment plans for the current year. In the fall, firms are asked about their investment plans for the current and the upcoming year. Thus, we have for each firm and year three different investment fore-

[9] We use a 5 percent significance level.

casts with different forecast horizons: the forecast from the fall preceding a given year, the spring and the fall forecast in a given year. In addition, firms report their actual investment expenditures of the preceding calendar year in the spring and in the fall surveys. Figure 3.3 illustrates the timing in the IFO-IS.

Figure 3.3: Firm Investment Plans and Realizations in the IFO-IS - Timing

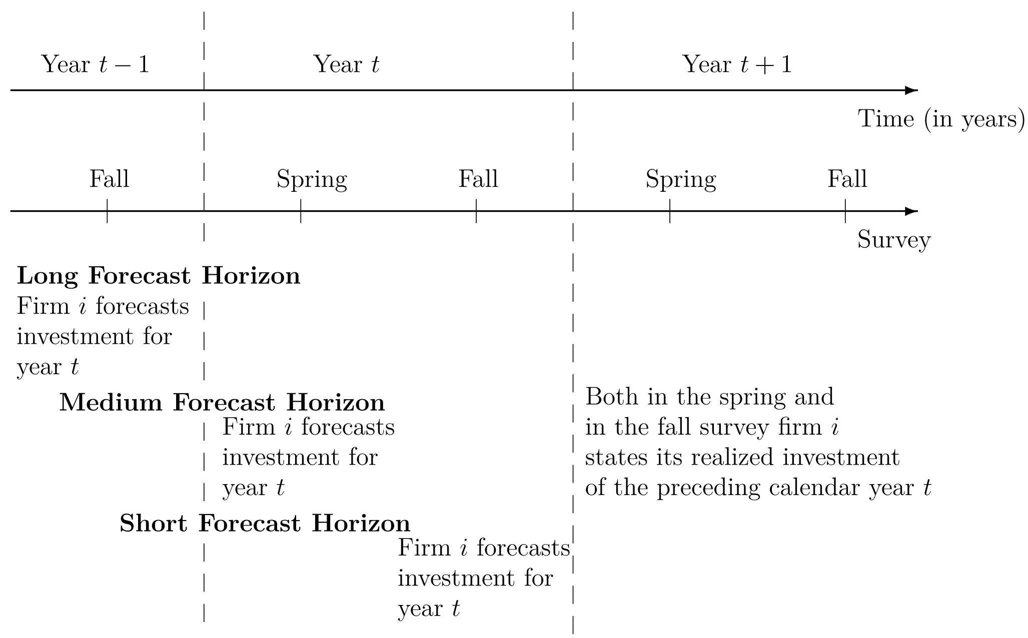

We compute the expectation error of firm i in period t as the log difference of the realized investment expenditures in period t and the predicted investment expenditures for period t.[10] Given the three forecast horizons we get three types of expectation errors. We denote the first one by $FE\_IS_{i,t}^{long}$. It uses as forecast the firms' investment plans from the fall of the previous year $t-1$. The second series of investment expectation errors, denoted by $FE\_IS_{i,t}^{med}$, uses the firm forecast given in the spring of the current year t. The last forecast error type uses the fall prediction of the current year t. It is denoted by $FE\_IS_{i,t}^{short}$.

For the final sample we perform two steps to eliminate outliers and inconsistent answers. First, we compare the spring and the fall statements concerning investment developments in the preceding calendar year. We drop all firm observations that have a 10 percent difference

[10]We use the realized investment expenditures from the fall survey in period $t+1$. Using instead the investment numbers from the spring survey or the average of both surveys does not change our results. We also compute percentage expectation errors by dividing the difference of the realized and predicted investment expenditures through the mean of the realized and predicted investment expenditures without much change to our results.

in these statements. Second, we leave out those observations which are smaller than the 1th percentile and larger than the 99th percentile of the corresponding year. We then compute for each firm the average firm-specific expectation error. Table 3.2 presents the distributions for each average investment expectation error: $AFE\_IS_i^{long}$, $AFE\_IS_i^{med}$ and $AFE\_IS_i^{short}$.[11] Again we only include firms for which we have at least 8 investment expectation errors, i.e. 8 years of observations.

Table 3.2: Firm-Specific Average Investment Expectation Errors (IFO-IS)

Statistics	$AFE\_IS_i^{long}$	$AFE\_IS_i^{med}$	$AFE\_IS_i^{short}$
Obs	1,843	1,904	2,169
Mean	0.0380	-0.0220	0.0072
Std.Dev.	0.2580	0.1844	0.1377
Percentiles			
5th	-0.3645	-0.3285	-0.2123
10th	-0.2685	-0.2507	-0.1613
25th	-0.1211	-0.1272	-0.0799
50th	0.0310	-0.0239	0.0006
75th	0.1933	0.0875	0.0875
90th	0.3538	0.2071	0.1816
95th	0.4512	0.2805	0.2332
# Optimists	131 (7.1%)	159 (8.3%)	166 (7.7%)
# Pessimists	138 (7.5%)	76 (4.0%)	130 (6.0%)
# Realists	1,574 (85.4%)	1,669 (87.7%)	1,873 (86.3%)

Notes: This table provides a summary of the distributions of the firm-specific means of $FE\_IS_{i,t}^{long}$ (second column) $FE\_IS_{i,t}^{med}$ (third column) and $FE\_IS_{i,t}^{short}$ (fourth column). Only those firms with at least eight investment expectation errors in the sample are considered. We define optimistic firms as those firms which feature a negative average expectation error that is significantly different from zero. Pessimistic firms are defined as those firms with a positive average expectation error that is significantly different from zero.

The first feature of the investment expectation errors to note is that firms get better over time in forecasting their annual investment: the standard deviation of the forecast error declines from roughly 26 percent to approximately 14 percent between the two extreme forecast horizons. The second feature is that unlike in production expectation errors, the distribution is not skewed towards "optimism" or overprediction. For each type of expectation error we observe that more than 20 percent of firms have average expectation errors that deviate at least one standard deviation of the corresponding distribution from zero. For $AFE\_IS_i^{long}$ more than half of all firms have average expectation errors larger than 10 percent. However, this bias is not significant for all of these firms. After using the same definitions for optimists and pessimists as in the IFO-BCS we find only 14 percent of firms with expectational biases.

[11] Figure A4 in Appendix A4 provides the histograms of the distributions.

For $AFE\_IS_i^{med}$ and $AFE\_IS_i^{short}$ this fraction becomes even smaller. But overall, for each type we find that at least 12 percent of all firms feature an expectational bias.

3.4 A Model

3.4.1 Firms

The model economy is populated by a unit mass continuum of ex ante identical, but ex post potentially heterogeneous firms. They produce a final generic good using a diminishing returns to scale production function with capital and labor as inputs. In addition, production is affected by idiosyncratic productivity shocks. Specifically, an individual firm i's production function is given by:

$$y_{i,t} = z_{i,t} k_{i,t}^{\theta} n_{i,t}^{\nu}, \tag{3.1}$$

where $z_{i,t}$ denotes the idiosyncratic productivity level, and $k_{i,t}$ and $n_{i,t}$ denote the firm-specific capital stock and employment level, respectively. We assume competitive factor markets. Firms pay a real wage w_t for one unit of labor input and a rental rate r_t for one unit of capital input.

To incorporate expectational biases, we assume that firms decide about their factor demands in period t before knowing their productivity levels $z_{i,t}$. Instead, firms form expectations about $z_{i,t}$ on the basis of $z_{i,t-1}$. We further assume that the natural logarithm of $z_{i,t}$ follows a three-state Markov chain on the states $[\epsilon, 0, -\epsilon]$ with the following symmetric transition matrix P^{obj}:

$$P^{obj} = \begin{pmatrix} \rho_1 & \rho_2 & 1-\rho_1-\rho_2 \\ \rho_3 & 1-2\rho_3 & \rho_3 \\ 1-\rho_1-\rho_2 & \rho_2 & \rho_1 \end{pmatrix}, \tag{3.2}$$

where $\rho_1 + \rho_2 < 1$ and $\rho_3 < 0.5$. P^{obj} is the actual transition matrix of the idiosyncratic productivity process. Rational firms would use this transition matrix to compute expectations about their idiosyncratic productivity levels. Some firms, however, which we call optimists and pessimists form their expectations with different transitions matrices, $P^{subj,opt}$ and $P^{subj,pess}$. Relative to P^{obj}, $P^{subj,opt}$ and $P^{subj,pess}$ feature expectational biases which we model parsimoniously with a parameter ϕ. This parameter is introduced additively into the true transition matrix P^{obj}. Specifically, the subjective transition matrix of an optimist looks as follows:

$$P^{subj,opt} = \begin{pmatrix} \rho_1 & \rho_2 & 1-\rho_1-\rho_2 \\ \rho_3+\phi & 1-2\rho_3 & \rho_3-\phi \\ 1-\rho_1-\rho_2+\phi & \rho_2+\phi & \rho_1-2\phi \end{pmatrix}. \tag{3.3}$$

The subjective transition matrix for a pessimist is just analogous:

$$P^{subj,pess} = \begin{pmatrix} \rho_1-2\phi & \rho_2+\phi & 1-\rho_1-\rho_2+\phi \\ \rho_3-\phi & 1-2\rho_3 & \rho_3+\phi \\ 1-\rho_1-\rho_2 & \rho_2 & \rho_1 \end{pmatrix}. \tag{3.4}$$

With this expectation formation process the optimal factor demands are given by:

$$n_{i,t} = \left[\left(\frac{\nu E[z_{i,t}|z_{i,t-1}]}{w_t} \right) \left(\frac{\theta w_t}{\nu r_t} \right)^{\theta} \right]^{\frac{1}{1-\theta-\nu}} \tag{3.5}$$

$$k_{i,t} = \frac{\theta w_t}{\nu r_t} n_{i,t}. \tag{3.6}$$

Note that if a firm expects a higher expected value of its productivity level $E[z_{i,t}|z_{i,t-1}]$, it will demand more capital and labor. This implies that overoptimistic firms hire too many workers and build up capital stocks that are too high. In the other direction overpessimistic firms do not demand enough inputs. This leads to factor misallocation and welfare losses.

3.4.2 Households

We assume a representative household with time separable preferences who maximizes the following instantaneous utility function:

$$U_t = \log C_t - \frac{A}{1+\eta} {N_t}^{1+\eta}, \tag{3.7}$$

where C_t is aggregate consumption and N_t denotes aggregate employment. η is the inverse of the Frisch elasticity of labor supply. The budget constraint of the household is given by:

$$w_t N_t + (1-\delta+r_t)K_t + \Pi_t = K_{t+1} + C_t, \tag{3.8}$$

where δ is the depreciation rate, K_t the aggregate capital stock and Π_t denotes aggregate profits. We assume that all firms are owned by the representative household who does not know the types of the firms, for instance whether they are realists, optimists or pessimists.

After solving the intertemporal optimization problem of the household we get the usual first-order conditions:

$$w_t = AC_t N_t^{\eta}, \tag{3.9}$$

$$\frac{1}{C_t} = \beta E_t \left[\frac{1 + r_{t+1} - \delta}{C_{t+1}} \right]. \tag{3.10}$$

3.4.3 Equilibrium

Given an initial capital stock, K_0, and a sequence of shocks $\{\{z_{i,t}\}_{i=0}^{\infty}\}_{t=0}^{\infty}$ an equilibrium of this economy is defined as a time path of quantities $\{\{y_{i,t}\}_{i=0}^{\infty}, \{k_{i,t}\}_{i=0}^{\infty}, \{n_{i,t}\}_{i=0}^{\infty}, C_t, K_t, N_t\}_{t=0}^{\infty}$ and a time path of prices $\{w_t, r_t\}_{t=0}^{\infty}$ that satisfy:

1. *Firm optimality:* Taking $\{w_t, r_t\}_{t=0}^{\infty}$ as given, the optimal factor demands for $n_{i,t}$ and $k_{i,t}$ are determined according to equations (3.5) and (3.6).

2. *Household optimality:* Taking $\{w_t, r_t\}_{t=0}^{\infty}$ and K_0 as given, the household's consumption and labor supply satisfy (3.9) and (3.10).

3. *Commodity market clearing:*

$$C_t = \int z_{i,t} k_{i,t}^{\theta} n_{i,t}^{\nu} \, di - \left(K_{t+1} - (1 - \delta) \int k_{i,t} \, di \right)$$

4. *Labor market clearing:*

$$N_t = \int n_{i,t} \, di$$

3.4.4 Calibration

The model period is a quarter. Table 3.3 gives an overview of the standard parameter choices for calibration: Bachmann and Bayer (2011) compute from national accounting data an average annual depreciation rate of 0.094 for Germany. They also estimate the median factor shares of labor and capital in the German manufacturing sector from firm-level micro data. The discount factor generates an annual real interest rate of 2 percent. We fix the Frisch elasticity of labor supply at unity. The disutility parameter of labor, A, is chosen to ensure that the average time spent at work by the representative household is 0.33.

Table 3.3: STANDARD PARAMETER VALUES

Baseline Calibration		
Parameter	Description	Value
δ	Depreciation rate	0.0235
θ	Decreasing returns to capital	0.2075
ν	Decreasing returns to labor	0.5565
β	Discount factor	0.9950
η	Inverse elasticity of labor supply	1.0000
A	Disutility of labor	6.0000

The remaining parameters ρ_1, ρ_2, ρ_3, ϕ and ϵ are calibrated using the IFO-BCS. As a reminder, P^{obj}, the true transition matrix for the idiosyncratic productivity shock process, is given by:

$$P^{obj} = \begin{pmatrix} \rho_1 & \rho_2 & 1-\rho_1-\rho_2 \\ \rho_3 & 1-2\rho_3 & \rho_3 \\ 1-\rho_1-\rho_2 & \rho_2 & \rho_1 \end{pmatrix}.$$

Note that ρ_1, ρ_2 and ρ_3 define transition probabilities conditional on a certain economic state. The IFO-BCS provides such information in a qualitative way with the following question concerning the current business situation:

Q 6 *"Current Situation: We evaluate our business situation with respect to product XY as good, satisfactory, unsatisfactory."*

We start with the calibration of ρ_3. These entries define situations in which the economic states of the firms do not change. Suppose that a firm is in the medium economic state. This state will not change with probability $(1-2\rho_3)$. Therefore, we compute for each quarter the fraction of firms with no upcoming production change, i.e. $REALIZ_{i,t+3}$ is equal to zero, conditional on a normal current business situation. The time series average of these fractions provides an estimate of $(1-2\rho_3)$.[12]

The probability for firms to persist in either the good or bad economic state is given by ρ_1. We compute the fractions of firms that have no decrease in their production level over the next three months, i.e. $REALIZ_{i,t+3}$ is greater or equal to zero, conditional on a good current business situation. Similarly, we compute the fraction of firms that have no increase in their production level over the next three months, i.e. $REALIZ_{i,t+3}$ is less or equal to

[12] As before we use only those firms that pass the outlier and consistency test (Q1 and Q5).

zero, conditional on a bad current business situation. Finally, we take the average of the two time-series averages to get ρ_1.

To calibrate ρ_2 we compute the unconditional quarterly fractions of firms which change their production level over the upcoming quarter, i.e. $REALIZ_{i,t+3}$ is unequal to zero. The time-series average of these fractions provides an empirical moment that has to be matched by the model. Given ρ_1 and ρ_3 we find a value for ρ_2 that yields the same fraction of firms changing their economic state in the model. The first three rows of Table 3.4 summarize the calibration results so far.

Table 3.4: Parameter Values of P^{obj}, ϕ, ϵ

Baseline Calibration			
Parameter	Description	$FE_{i,t}^{narrow}$	$FE_{i,t}^{broad}$
ρ_1	Parameter Transition Matrix	0.8626	0.8479
ρ_2	Parameter Transition Matrix	0.1374	0.1521
ρ_3	Parameter Transition Matrix	0.2072	0.2675
ϵ	Parameter of Technology State	0.0975	0.1979
ϕ	Expectational Bias Parameter	0.1122	0.1073

The calibration of ϵ and ϕ requires the simulation of the entire model and has to be done jointly. In a first step we compute separately for the "realist" and "non realist" firms the means of the absolute value of their expectation errors. The average absolute expectation error for "realist" firms identifies ϵ, the same statistic for "non-realists" identifies ϕ. For a given choice of the forecast error type in the data, each guess for ϵ and ϕ allows us to compute the model average absolute forecast errors for "realist" and "non-realist" firms. We pick ϵ and ϕ such that the model numbers correspond to their data counterparts. The calibrated values of ϕ and ϵ are shown in the last two rows of Table 3.4.

To gauge how these Markov chains behave in terms of standard AR(1) modelling we simulate 100 times the Markov chain defined by P^{obj} and ϵ with 20,000 time series observations each. Then we estimate AR(1)-regressions on each of these time series. Table 3.5 displays the average of the 100 AR(1)-coefficients and standard errors of the regressions.

Table 3.5: AR(1)-Properties of the Idiosyncratic Shock Processes

	AR(1)-coefficient	Standard Error of Regression
$FE_{i,t}^{narrow}$	0.8626	0.0428
$FE_{i,t}^{broad}$	0.8477	0.0926

Notes: This table shows in the first column the average of the estimated 100 AR(1)-coefficients for the simulated Markov chains explained in the text for $FE_{i,t}^{narrow}$ and $FE_{i,t}^{broad}$. The second column displays the average of the standard errors of these regressions. Each Markov chain is simulated over 20,000 observations.

3.5 Welfare Calculations

To gauge - in a first pass - the economic significance of the observed expectational biases in the IFO-BCS, we compare the lifetime utility of the representative agent in the steady state of the actual economy with biased expectations, denoted by $Welfare^{act}$, with her lifetime utility in the following hypothetical scenario: suppose at some point in time t_0 all optimist and pessimist firms become "realists" and use P^{obj} to form their expectations. Welfare is determined by the discounted utility function of the representative household:

$$Welfare = \sum_{t=0}^{\infty} \beta^t \left(\log C_t - \frac{A}{1+\eta} N_t^{1+\eta} \right). \tag{3.1}$$

At t_0 this economy starts out at the steady state capital stock of the economy with expectational biases and then transitions towards a new steady state. We can compute the welfare of the representative household along this transition path, denoted by $Welfare^{hypo}$, according to equation (3.1). Then we determine the welfare loss as the percent deviation of $Welfare^{act}$ and $Welfare^{hypo}$. We also compute the consumption equivalent (in percent of the steady state consumption of the actual economy with expectational biases). Formally, we find a $\overline{C}$, such that:

$$\sum_{t=0}^{\infty} \beta^t \left(\log(C_t^{act} + \overline{C}) - \frac{A}{1+\eta} \left(N_t^{act} \right)^{1+\eta} \right) = Welfare^{hypo} \tag{3.2}$$

The results are presented in Table 3.6. It is clear that the welfare losses from expectational biases of firms in this simple model economy are small. The welfare losses range from 0.01 percent to at most 0.1 percent. The welfare costs under the broad definition of the production expectation error are higher than those under the narrow definition. The main reason for this is that there are more optimist and pessimist firms under the broad definition. But even then the welfare loss in terms of consumption amounts to only roughly

Table 3.6: WELFARE LOSSES ASSOCIATED WITH BIASED EXPECTATIONS

Type of Expectation Error	**Welfare Loss in %**	**Consumption Equivalent in %**
$FE_{i,t}^{narrow}$	0.0071	0.0054
$FE_{i,t}^{broad}$	0.1124	0.0800

Notes: This table provides estimates of the welfare costs. The welfare loss is computed as the percent deviation of $Welfare^{act}$ and $Welfare^{hypo}$. The third column shows the welfare loss expressed in terms of consumption. To compute this number we divide $\overline{C}$ by the steady state consumption level of the actual economy with expectational biases.

0.1 percent, a number comparable to conventional estimates of the costs of business cycle fluctuations (see Lucas, 1987, 2003).

Figure 3.4 presents the transition paths of major macroeconomic aggregates between the steady state with expectational biases and the hypothetical steady state with only realist firms for the broad definition of the production expectation errors.[13] Notice first that the hypothetical steady state with only realist firms features lower output, consumption, investment, employment and real wages, albeit higher average capital and labor productivity. It is clear that resources, in particular capital, are more efficiently allocated in the hypothetical steady state. In contrast, in the steady state with expectational biases optimist firms dominate and the economy has too much capital accumulated and works too much. This is where the transition analysis as opposed to mere steady state comparison is crucial. After the elimination of all expectational biases at firms the economy enjoys a boom in consumption and leisure on impact which ultimately leads to the welfare improvements documented in Table 3.6.

[13]In Appendix A5 we also show the transition paths for the narrow definition of quantitative expectation errors.

Figure 3.4: Transition Paths for the Case of $FE_{i,t}^{broad}$

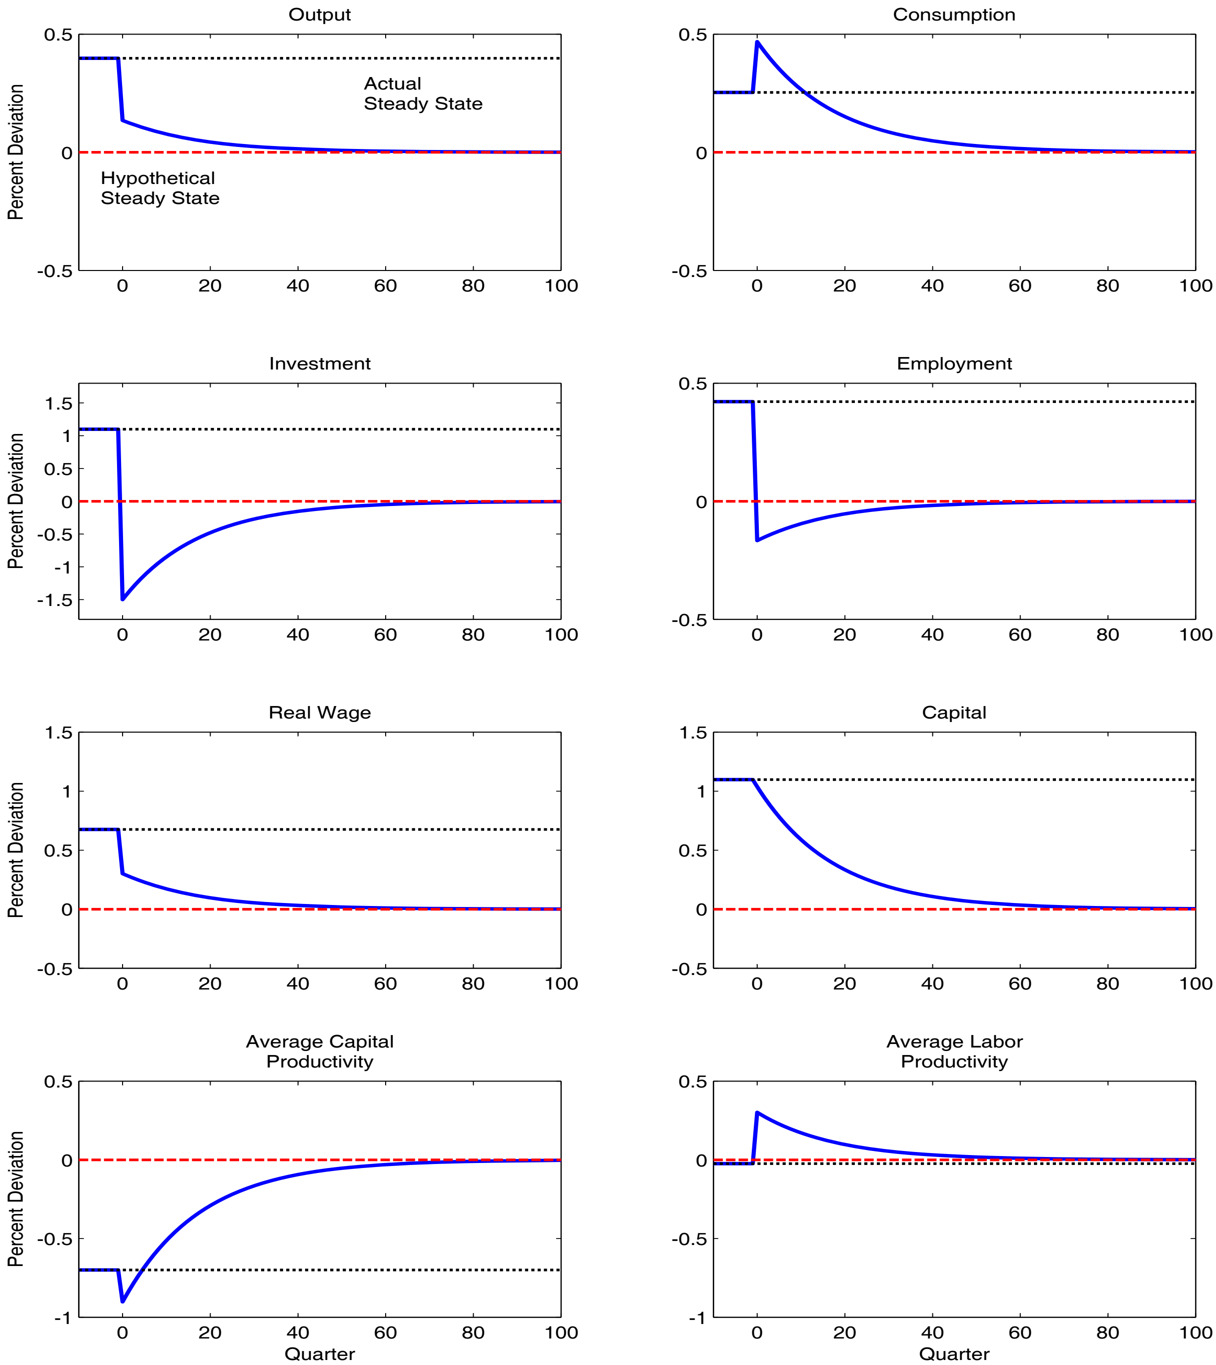

Notes: This figure shows the transition paths of selected variables for the case of $FE_{i,t}^{broad}$. These dynamic responses are expressed as percentage deviations from the steady state of the hypothetical economy.

3.6 Robustness Checks

This section provides a series of robustness checks to our baseline welfare calculations. All robustness checks change one feature of the calibration at a time, relative to the baseline scenario. First, we reconsider the definition of optimistic and pessimistic firms. In the baseline case we defined expectational biases by exploiting firm histories of expectation errors, i.e. we defined expectational biased firms statistically as firms whose average production expectation error was significantly different from zero, given their history of production expectation errors. An alternative is to define optimists as having an average expectation error below the 10th percentile of the average production error distribution and a pessimist as having an average expectation error above the 90th percentile.[14] As an alternative we also use a 25/75th percentile threshold. The second and third panel in Table 3.7 show that in this case welfare losses are somewhat higher, but their overall magnitude is comparable to our baseline estimates.

The next robustness check concerns the calibration of ρ_1, ρ_2 and ρ_3. In the baseline scenario we calibrated those transition probabilities by using the qualitative production index $REALIZ_{i,t}$. In particular, we were interested in the signs of this index for the respective calibration exercises. To see whether our results are sensitive to this choice, we replace it by $sign(\Delta \log u_{i,t})$, which is derived from the quantitative production changes. This modification in our calibration strategy delivers higher welfare costs for both types of expectation errors. Nonetheless, these numbers are of the same order of magnitude as our baseline results.

We also check robustness with respect to the number of observations that we require a firm to have in order to compute an average production forecast error. Instead of 32 observations we use a threshold of 16 observations or four years of quarterly quantitative production errors. This gives us a larger cross section of firms at the cost of having more firms with shorter histories in our sample. In the case of $FE_{i,t}^{narrow}$ the welfare losses get somewhat higher, but they are still in line with our baseline estimates.

In the next robustness check we do not clean our sample from inconsistent statements regarding quantitative and qualitative production changes. After all these inconsistent statements might indicate some form of irrationality. Thus, the removal of these observations could bias the welfare cost estimates downward. If we repeat our analysis without eliminating these observations, we, unsurprisingly, observe slight increases in all welfare losses. But their order of magnitude remains unchanged.

[14] We provide the calibrated values of ρ_1, ρ_2, ρ_3, ϵ and ϕ for these robustness checks in Appendix A6. For some cases ϕ takes on corner solutions that cannot match the data moments perfectly.

Table 3.7: Robustness Checks - Welfare Losses

Type of Expectation Error	**Welfare Loss in %**	**Consumption Equivalent in %**
Baseline Results		
$FE_{i,t}^{narrow}$	0.0071	0.0054
$FE_{i,t}^{broad}$	0.1124	0.0800
10% and 90% Percentile		
$FE_{i,t}^{narrow}$	0.0562	0.0429
$FE_{i,t}^{broad}$	0.1696	0.1199
25% and 75% Percentile		
$FE_{i,t}^{narrow}$	0.0938	0.0719
$FE_{i,t}^{broad}$	0.1557	0.1109
Quantitative Production Changes: $Sign(\Delta \mathbf{log} u_{i,t})$		
$FE_{i,t}^{narrow}$	0.0089	0.0068
$FE_{i,t}^{broad}$	0.2113	0.1556
16 Observations of Expectation Errors		
$FE_{i,t}^{narrow}$	0.0269	0.0201
$FE_{i,t}^{broad}$	0.1076	0.0762
Including Inconsistent Statements		
$FE_{i,t}^{narrow}$	0.0185	0.0138
$FE_{i,t}^{broad}$	0.1260	0.0898
Economy with ϕ_{opt}, ϕ_{pess}, ϵ_u **and** ϵ_l		
$FE_{i,t}^{narrow}$	0.0033	0.0026
$FE_{i,t}^{broad}$	0.1170	0.0938

Notes: see notes to Table 3.6. This table shows the welfare costs estimates of expectational biases for the baseline case and robustness checks.

In our final robustness check we relax the two symmetry assumptions built into the stochastic environment of the firms. We allow for pessimistic and optimistic firms to have different expectational biases. Also, ϵ, the parameter that governs the average absolute expectation error, is now allowed to be different for upward and downward shocks.[15] We replace the symmetric technology state vector $[\epsilon, 0, -\epsilon]$ by an asymmetric one $[\epsilon_u, 0, -\epsilon_l]$. As a result we have to calibrate with ϕ_{opt}, ϕ_{pess}, ϵ_u and ϵ_l four instead of two parameters. We again use the means of the absolute quantitative expectation errors for the calibration and extend the analysis by considering particular subgroups. We calibrate ϵ_u using the absolute values of the positive expectation errors of the realist firms. To compute ϵ_l we use the absolute values of the negative expectation errors for the same type of firms. To determine

[15]The parameter of the transition matrix P_{opt} are the same as in the baseline case.

ϕ_{opt} and ϕ_{pess} we differentiate between expectation errors of optimistic and pessimistic firms. Then we pick up those values for ϵ_u, ϵ_l, ϕ_{opt} and ϕ_{pess} such that the corresponding numbers in the model correspond to the observed means in the data. Table 3.8 summarizes the calibrated values of ϕ_{opt}, ϕ_{pess}, ϵ_u and ϵ_l.

Table 3.8: Calibration of the Asymmetric Case

Calibration of the Asymmetric Case			
Parameter	Description	$FE_{i,t}^{narrow}$	$FE_{i,t}^{broad}$
ϵ_u	Parameter of high Technology State	0.0712	0.0767
ϵ_l	Parameter of low Technology State	0.0691	0.1776
ϕ_{opt}	Expectational Bias Parameter Optimists	0.0786	0.1500
ϕ_{pess}	Expectational Bias Parameter Pessimists	0.1848	0.2675

Notes: This table shows the calibrated parameter values for ϵ_u, ϵ_l, ϕ_{opt} and ϕ_{pess}.

With these parameter values we obtain the following welfare losses summarized in the last panel of Table 3.7. It turns out that these results do not considerably differ from our baseline scenario. The economy experiences at most a welfare loss in consumption of about 0.1 percent.

3.7 Conclusion

This paper, using the micro data from the German IFO Business Climate Survey and the IFO Investment survey, constructs panel data sets of firms' quantitative expectation errors with respect to their production and investment. With these data sets, we can gauge whether firms errors are systematic and thus biased towards optimism or pessimism. We find some degree of biased expectations, but for a large majority of firms realistic expectations in the sense that zero average expectation errors cannot be rejected.

We then ask in the simplest possible neoclassical heterogeneous firm model, where expectation errors play a role, what the welfare implications of these expectational biases are. Obviously, expectational biases lead to factor misallocations in an economy, and a welfare analysis will allow us to gauge the economic significance of such a misallocation. We find that even when expectational errors are very broadly defined, the welfare costs of these misallocations are generally small, at the order of magnitude of conventional estimates for the welfare costs of business cycles.

We do, however, note that our model is somewhat simplistic and expect future research to compute welfare losses in more realistic environments. In this sense, we view the second half of the paper only as a first pass. We speculate that in economies with physical adjustment frictions to capital, financial frictions or endogenous growth elements will lead to larger welfare losses from expectational biases. Future research can then make use of our distributional results for firms' average production and investment errors to calibrate such models.

Acknowledgements

I am indebted to Rüdiger Bachmann who is co-author of Chapter 3. We are grateful to Doris Hauke, Andre Kunkel, Heike Mittelmeier, Wolfgang Ruppert, Christian Seiler, Sigrid Stallhofer and Annette Weichselberger from the IFO Institute for helping us with the data and introducing us to the institutional backgrounds.

Bibliography

Anderson, O., R. K. Bauer, and E. Fels (1954): "On the accuracy of short-term entrepreneurial expectations," in *Proceedings of the Business and Economic Statistics Section*, pp. 124–147. American Statistical Association.

Bachmann, R., and C. Bayer (2011): "Uncertainty Business Cycles - Really?," Working Paper 16862, National Bureau of Economic Research.

Becker, S., and K. Wohlrabe (2008): "Micro Data at the Ifo Institute for Economic Research - The 'Ifo Business Survey' Usage and Access," *Schmollers Jahrbuch*, 128, 307–319.

Bovi, M. (2009): "Economic versus Psychological Forecasting. Evidence from Consumer Confidence Survey," *Journal of Economic Psychology*, 30(4), 563–574.

Brown, P., A. Clarke, J. C. Y. How, and K. Lim (2000): "The Accuracy of Management Dividend Forecasts in Australia," *Pacific-Basin Finance Journal*, 8, 309–331.

Brunnermeier, M. K., and J. Parker (2005): "Optimal Expectations," *American Economic Review*, 95(4), 1092–1118.

Davis, S. J., and J. C. Haltiwanger (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(3), 819–863.

de Leeuw, F., and M. J. McKelvey (1981): "Price expectations of business firms," *Brookings Papers on Economic Activity*, 1981(1), 299–314.

Doms, M., and T. Dunne (1998): "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1(2), 409–429.

Firth, M., and A. Smith (1992): "The Accuracy of Profits Forecasts in Initial Public Offering Prospectuses," *Accounting and Business Research*, 22(87), 239–247.

Hassan, T. A., and T. M. Mertens (2011): "The Social Cost of Near-Rational Investment," Working Paper 17027, National Bureau of Economic Research.

Jaimovich, N., and S. Rebelo (2007): "Behavioral Theories of the Business Cycle," *Journal of the European Economic Association*, 5(2–3), 361–368.

Lucas, R. E. (1987): *Models of Business Cycle*. Oxford University Press.

——— (2003): "Macroeconomic Priorities," *American Economic Review*, 93(1), 1–14.

Lui, S., J. Mitchell, and M. Weale (2008): "Qualitative Business Surveys: Signal or Noise?," Discussion Paper 323, NIESR.

Malmedier, U., and G. Tate (2005): "CEO Overconfidence and Corporate Investment," *Journal of Finance*, 60(6), 2661–2700.

——— (2008): "Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction," *Journal of Financial Economics*, 89(1), 20–43.

McDonald, C. (1973): "An Empirical Examination of the Reliability of Published Predictions of Future Earnings," *The Accounting Review*, 48(3), 502–510.

Müller, H. C. (2010): "Firms' Forecast Errors Regarding their own Future Key Figures: The Disappearance of the Overoptimism Bias," *mimeo*.

Nerlove, M. (1983): "Expectations, Plans, and Realizations in Theory and Practice," *Econometrica*, 51(5), 1251–1279.

Souleles, N. S. (2004): "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 28(3), 39–72.

Tompkinson, P., and M. S. Common (1983): "Evidence on the rationality of expectations in the British manufacturing sector," *Applied Economics*, 15, 425–436.

Zimmermann, K. F., and S. Kawasaki (1986): "Testing the rationality of price expectations for manufacturing firms," *Applied Economics*, 18, 1335–1347.

Appendix

A1 IFO Business Climate Survey (IFO-BCS)

Original German IFO-BCS Questions

Q 1 *"Sonderfragen: Die **Ausnutzung** unserer **Anlagen** zur Herstellung von XY (betriebsübliche Vollauslastung=100%) beträgt **gegenwärtig** bis zu ...%. "*

30	40	50	60	70	75	80	85	90	95	100	mehr als 100 % und zwar

Q 2 *"Erwartungen für die nächsten 3 Monate: **Beschäftigte** (nur inländische Betriebe) - Die Zahl der mit der Herstellung von XY beschäftigten Arbeitnehmer wird: zunehmen, etwa gleichbleiben, abnehmen."*

Q 3 *"Sonderfragen: Unter Berücksichtigung unseres gegenwärtigen Auftragsbestandes und des von uns in den nächsten 12 Monaten erwarteten Auftragseingangs halten wir unsere derzeitige **technische Kapazität** für XY für: mehr als ausreichend, ausreichend, nicht ausreichend. "*

Q 4 *"Erwartungen für die nächsten 3 Monate: Unsere inländische **Produktionstätigkeit** – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY wird voraussichtlich: steigen, etwa gleich bleiben, abnehmen."*

Q 5 *"Tendenzen im vorangegangenen Monat: Unsere inländische **Produktionstätigkeit** – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY ist: gestiegen, etwa gleich geblieben, gesunken."*

Q 6 *"Aktuelle Situation: Wir beurteilen unsere **Geschäftslage** für XY als: gut, befriedigend, schlecht."*

A2 Derivation of Quantitative Expectation Errors under General Production Expectations

We begin by defining the following firm-specific activity variable $REALIZ_{i,t}$ as the sum of the Increase-instances minus the sum of the Decrease-instances in question Q5 over the last three months going backward from t. $REALIZ_{i,t}$ can have seven possible values that live in the interval [-3,3]. A qualitative production expectation error is then computed as follows:

Table A1: Possible Qualitative Expectation Errors

		$Experror_{i,t}$
Expected $Increase_{t-3}$	$REALIZ_{i,t}>0$	0
Expected $Increase_{t-3}$	$REALIZ_{i,t} \leq 0$	$(REALIZ_{i,t}-1)$
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}>0$	$REALIZ_{i,t}$
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}=0$	0
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}<0$	$REALIZ_{i,t}$
Expected $Decrease_{t-3}$	$REALIZ_{i,t}<0$	0
Expected $Decrease_{t-3}$	$REALIZ_{i,t} \geq 0$	$(REALIZ_{i,t}+1)$

Notes: Rows refer to the qualitative expectations in month $t-3$ (Q4).

$Experror_{i,t}$ ranges from $[-4,4]$, where, for instance, -4 indicates a really negative forecast error: the company expected production to increase over the next three months, yet every single subsequent month production actually declined.

Next we compute for all firms with a given value of $REALIZ_{i,t}$ in time t the average $\Delta \log u_{i,t}$, i.e. the cross-sectional average change in capacity utilization. Again, to ensure that we can treat utilization changes as production changes only those firms are considered that state three months before that their future employment levels remain the same and that their technical production capacities are sufficient. We compute this mapping between $REALIZ_{i,t}$ and average production changes for each point in time.[16] Figure A1 illustrates the timing of survey questions that are used to compute this mapping. Figure A2 depicts the results. For instance, it shows that firms with $REALIZ_{i,t}=1$ in the first quarter of 1980 maps into an average change in production of approximately 5 percent.

We subject this mapping to a couple of plausibility tests.[17] For firms with constant production expectations we have a simple plausibility test for our procedure. Suppose we know

[16] We also considered a firm size specific and an industry specific mapping, without much changes to our results.

[17] The pooled Spearman correlation coefficient between $REALIZ_{i,t}$ and $\Delta \log u_{i,t}$ for firms with constant employment and technical capacity expectations is 0.72. The pooled Kendall's tau between $sign(REALIZ_{i,t})$ and $sign(\Delta \log u_{i,t})$ is 0.67.

Figure A1: Mapping between Qualitative and Quantitative Production Changes

Prerequisite: Firm i passes the outlier and inconsistency test (Q1 and Q5)

Firm i states in period $t-3$: Constant employment expectations (Q2) *and* sufficient technical production capacity (Q3)

Firm i states in the periods $t-2$, $t-1$ and t its production change (Q5), i.e. $REALIZ_{i,t}$

$$REALIZ_{i,t} = j \longrightarrow \frac{1}{N}\sum_{i=1}^{N}(\Delta \log u_{i,t} | REALIZ_{i,t} = j), \text{ where } j = -3, ..., 3$$

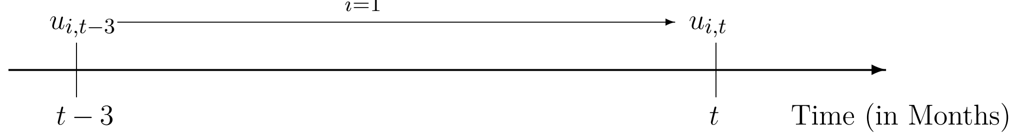

Figure A2: Mapping of $REALIZ_t$ into Quantitative Production Changes

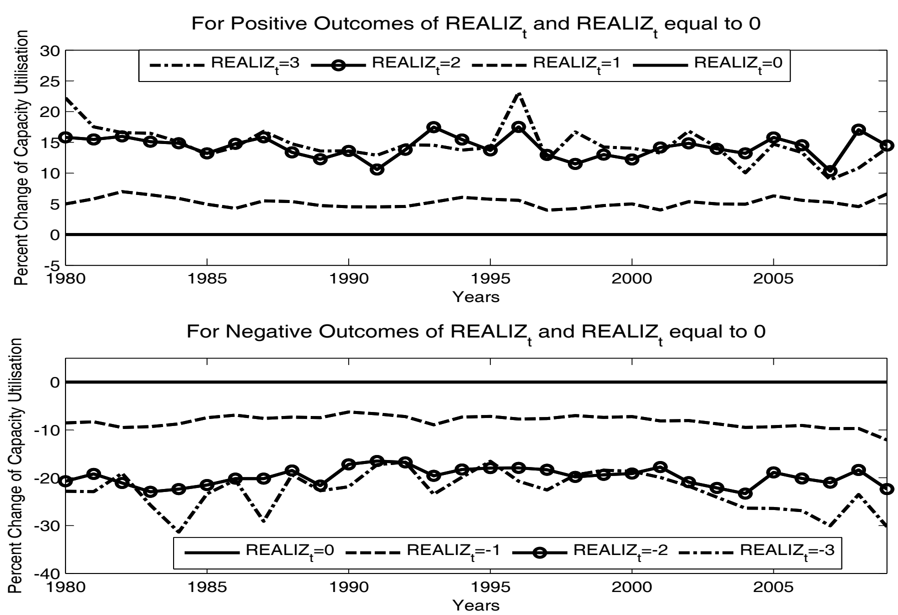

Notes: This figure shows for each possible value of $REALIZ_t$ the average $\Delta \log u_{i,t}$ of all firms with a given value for $REALIZ_t$. For better readability, the quarterly values of $\Delta \log u_{i,t}$ have been averaged to annual numbers. The upper panel shows the results for positive outcomes of $REALIZ_t$ and $REALIZ_t$ equal to zero. The lower panel does the results for negative outcomes of $REALIZ_t$ and $REALIZ_t$ equal to zero. The cross-sectional average $\Delta \log u_{i,t}$ of $REALIZ_t = 0$ has been subtracted from all time series shown in the figure.

with $FE\_BCS_{i,t}^{narrow}$ the "true" forecast errors under constant production expectations. We can also use the mapping strategy for the firms with constant qualitative production expectations, in which case the average percentage change in capacity utilization for a given $REALIZ_{i,t}$ would be another estimate for their production expectation error. We can then compare the "true" forecast errors with the forecast errors determined in the mapping procedure.

We do this by computing a quarterly time series of the cross-sectional means and standard deviations for these two types of expectation errors. The time series correlation coefficients between the two series are high: 0.97 for the average expectation errors and 0.84 for their standard deviation. This means that *over time* the mapping strategy captures the first and second moment of the expectation error distribution rather well. However, *on average* the "true" expectation errors are much more disperse than those based on the mapping strategy. Since the time series behavior of the cross-sectional standard deviation is very similar, we view this as a scaling problem, resulting from the discretization of expectation errors in the mapping strategy. As a consequence, we rescale the quantitative production change values that $REALIZ_{i,t}$ is mapped into by a constant.[18]

[18]This constant is calculated by dividing the time average of the cross-sectional standard deviations of $FE\_BCS_{i,t}^{narrow}$ by the corresponding value derived from the mapping procedure on the same subset of observations. Its value is 1.7.

A3 Firm-Specific Average Production Expectation Errors (IFO-BCS)

Figure A3: Histograms of the Firm-Specific Average Production Expectation Errors (IFO-BCS)

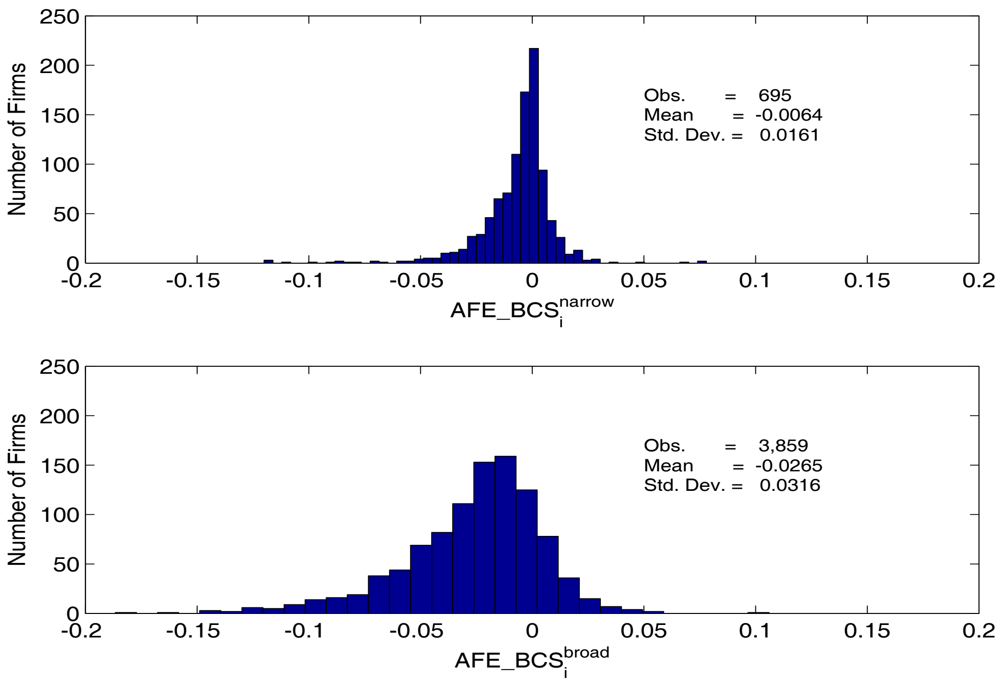

Notes: This figure shows the histograms of the distributions of the firm-specific means of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$.

A4 Firm-Specific Average Investment Expectation Errors (IFO-IS)

Figure A4: Histograms of the Firm-Specific Average Investment Expectation Errors (IFO-IS)

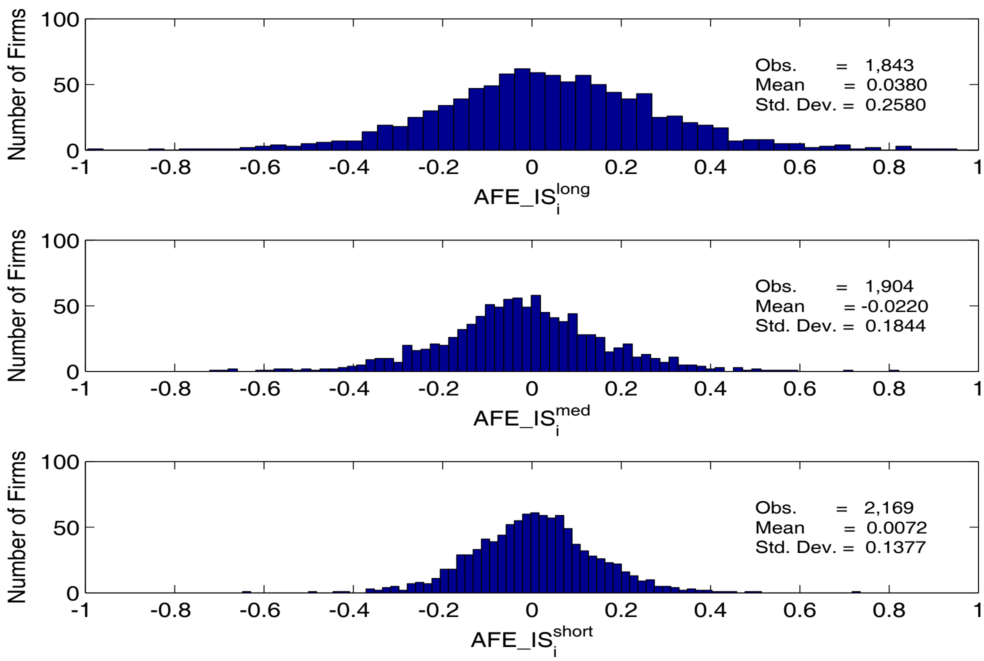

Notes: This figure shows the histograms of the distributions of the firm-specific means of $FE\_IS_{i,t}^{long}$, $FE\_IS_{i,t}^{med}$ and $FE\_IS_{i,t}^{short}$.

A5 Transition Paths for the Case of $FE_{i,t}^{narrow}$

Figure A5: TRANSITION PATHS FOR THE CASE OF $FE_{i,t}^{narrow}$

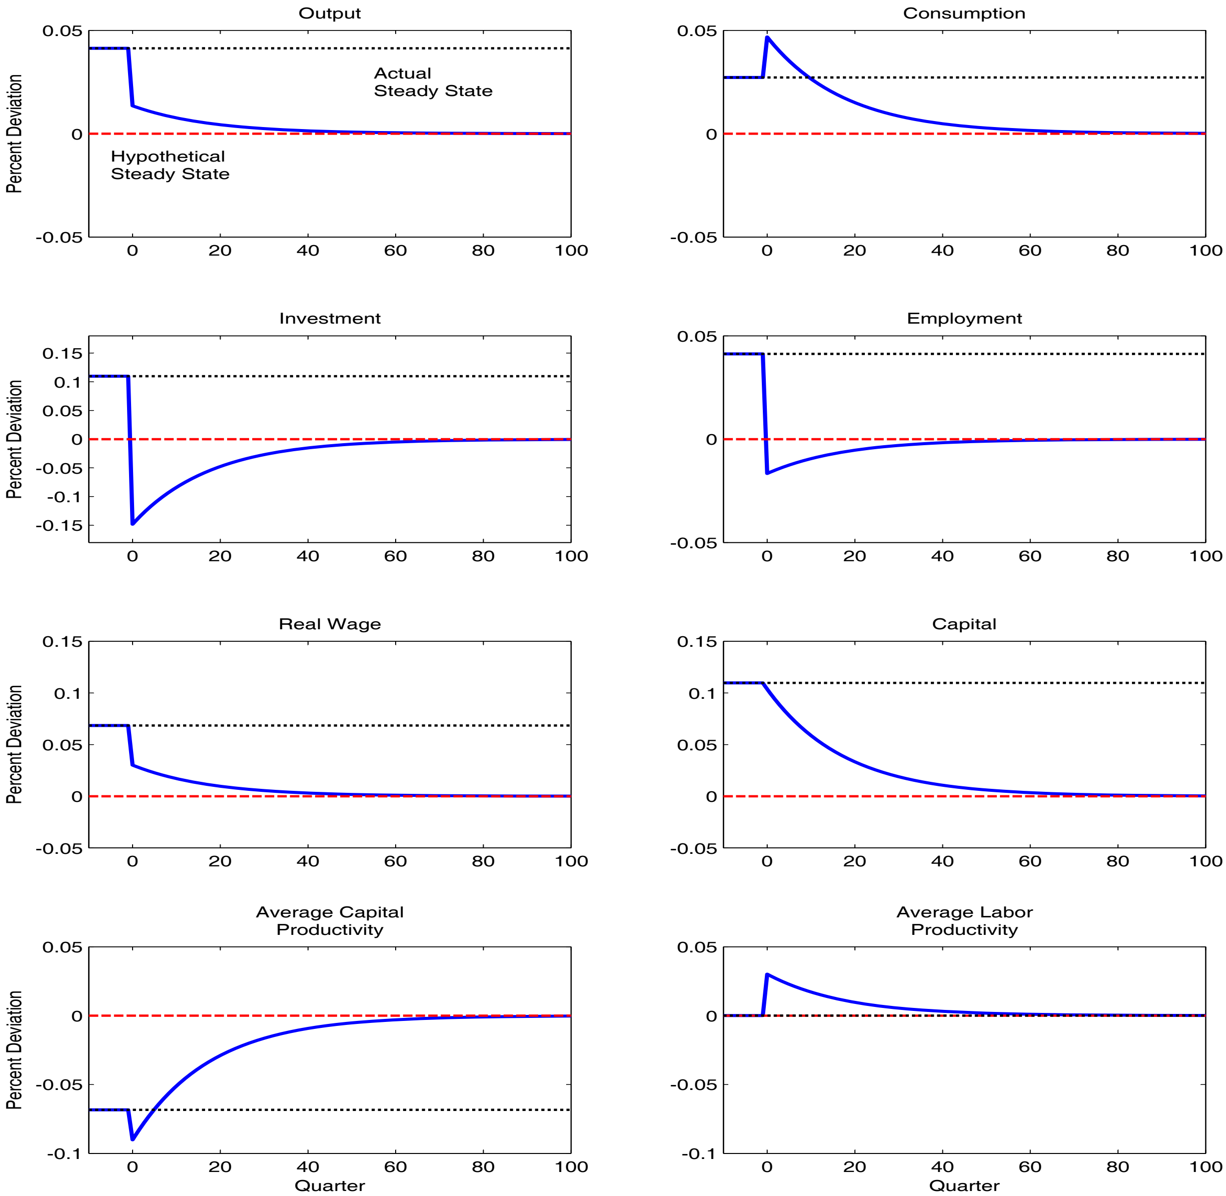

Notes: This figure shows the transition paths of selected variables for the case of $FE_{i,t}^{narrow}$. These dynamic responses are expressed as percentage deviations from the steady state of the hypothetical economy.

A6 Robustness Checks - Calibration

Table A2: Calibration Robustness Checks

	Parameter Transition Matrix ρ_1	Parameter Transition Matrix ρ_2	Parameter Transition Matrix ρ_3	Parameter Technology State ϵ	Expectational Bias Parameter ϕ	AR(1)-coefficient	S.E. of Regress.
Baseline Results							
$FE_{i,t}^{narrow}$	0.8626	0.1374	0.2072	0.0975	0.1122	0.8626	0.0428
$FE_{i,t}^{broad}$	0.8479	0.1521	0.2675	0.1979	0.1073	0.8477	0.0926
Robustness Check - 10th and 90th Percentile							
$FE_{i,t}^{narrow}$	0.8626	0.1374	0.2072	0.0771	0.2071	0.8630	0.0338
$FE_{i,t}^{broad}$	0.8479	0.1521	0.2675	0.2073	0.1370	0.8477	0.0970
Robustness Check - 25th and 75th Percentile							
$FE_{i,t}^{narrow}$	0.8626	0.1374	0.2072	0.0630	0.2071	0.8627	0.0276
$FE_{i,t}^{broad}$	0.8479	0.1521	0.2675	0.1988	0.0867	0.8477	0.0930
Robustness Check - Quantitative Production Changes: $Sign(\Delta \mathbf{log} u_{i,t})$							
$FE_{i,t}^{narrow}$	0.8506	0.1404	0.2980	0.0830	0.1455	0.8411	0.0404
$FE_{i,t}^{broad}$	0.8013	0.1987	0.3987	0.1544	0.1790	0.8008	0.0827
Robustness Check - At Least 16 Observations							
$FE_{i,t}^{narrow}$	0.8536	0.1463	0.2437	0.1202	0.1794	0.8607	0.0545
$FE_{i,t}^{broad}$	0.8476	0.1524	0.2762	0.2037	0.1088	0.8498	0.0912
Robustness Check - Including Inconsistent Statements							
$FE_{i,t}^{narrow}$	0.8606	0.1394	0.2300	0.1223	0.2299	0.8533	0.0550
$FE_{i,t}^{broad}$	0.8496	0.1504	0.2647	0.1959	0.1122	0.8477	0.0957

Notes: This table provides the calibrated values of ρ_1, ρ_2, ρ_3, ϵ and ϕ for the robustness checks. The next to last column displays the average of the estimated 100 AR(1)-coefficients for the simulated Markov chains resulting from these parameter choices. The last column shows the average of the standard errors of these regressions. Each Markov chain is simulated over 20,000 observations.

Chapter 4

Heterogeneous Expectation Errors of Firms:
Evidence From The IFO Business Climate Survey

Abstract

Are firms' expectations rational? I use micro data from the IFO Business Climate Survey on firms' production expectations and compare them to realization data from the same survey. I construct a unique quarterly panel data set of firms' quantitative expectation errors for a time horizon of 30 years. In a conservative estimate of expectation errors, I find that about two-thirds of the firms in my sample have rational expectations, i.e. their expectations are unbiased and the firms use all available information efficiently. When the expectation error estimate is defined less conservatively, this number reduces substantially to 33.2 percent. Thus, there is evidence that heterogeneous firms form their expectations in heterogeneous ways, i.e. firms differ in regard to their beliefs and in their ability to process information.

4.1 Introduction

Are firms' expectations rational? This is an important question since expectation formation plays a crucial role in macroeconomic models. Recent theoretical literature is critical of the traditional rational expectation assumption. In this paper, I answer this question by exploiting unique micro data from the IFO Business Climate Survey (IFO-BCS).

Ever since the seminal works of Muth (1961) and Lucas (1976) almost all business cycle research has been conducted under the rational expectations assumption (see e.g. Gali, 2008). One key feature of the rational expectations hypothesis is that economic agents are able to acquire and process *all* available information in a model-consistent way. Recently, macroeconomic models using the rational expectations assumption have been questioned as they fail to produce the sluggish adjustment behavior of macroeconomic variables such as inflation and output that is usually found in empirical data. To incorporate sluggish adjustment behavior, many studies (see e.g. Christiano et al., 2005, and Smets and Wouters, 2007) rely on special kinds of adjustment costs and price indexation rules. These model assumptions, however, often lack proper microfoundations. Sims (2003) proposes an alternative way of solving this problem. He shows that models can produce sluggish adjustment behavior if they drop the traditional rational expectations assumption. The observed heterogeneity in firms' and households' expectations with respect to aggregate variables such as inflation also weakens the rational expectations assumption. Weale and Pesaran (2006) point out that expectation heterogeneity arises due to differences in subjective probability densities (belief disparity) and differences in individual-specific information sets (information disparity).

Concerning belief disparity, Brunnermeier and Parker (2005) show that it can be optimal for an economic agent to bias his expectations concerning future outcomes upward to increase expected life time utility. On the empirical side, Müller (2010), using quantitative expectations about plants' sales and employment development from the annual IAB Establishment Panel in Germany, and Bachmann and Elstner (2011), exploiting the IFO Business Climate Survey and IFO Investment Survey, show that a sizeable proportion of firms can be classified as either optimists or pessimists.[1]

There are two streams of literature relevant to the topic of information disparity: the first points out that even though economic agents are able to efficiently acquire and process all available information, they do not continuously update their information sets. Mankiw and Reis (2002) analyze the implications for inflation of this sticky information approach. Reis (2006) provides a model in which it is suboptimal for a single firm to acquire and process all

[1]There is an accounting literature (e.g. McDonald, 1973, Firth and Smith, 1992, Brown et al., 2000), finding that managers tend to bias their public earnings and dividend forecasts upward, presumably as a strategy to attract investors.

available information due to information costs.[2] Consequently, firms update their information sets only at intervals. Cai (2010) shows that large firms update their information sets more quickly than do smaller firms, with the consequence that macroeconomic shocks are transmitted with more persistence throughout the economy. The second body of literature dealing with information disparity (see e.g. Veldkamp and Wolfers, 2007) points out that each individual information set consists of public and private information. Economic agents do not acquire all information as they entail a cost. The main difference compared to the sticky information approach is that people are able to continuously update their information sets but they do not gather all information. Specifically, firms and households decide whether public or private information is more beneficial to them.[3]

This study is motivated by the idea that firms may differ with respect to their beliefs and their ability to process information. I am specifically interested in three things: first, I want to discover whether firms have biased expectations (distorted beliefs) and, if so, if there are certain characteristics specific to these firms. Second, I am interested in the proportion of firms that omit private information from their expectation formation processes. Third, I aim to provide an estimate of the fraction of firms that have "truly" rational expectations. In pursuing answers to these questions, it is crucial to analyze individual response data.

One advantage of business surveys is that they ask firms directly about their expectations. However, empirical literature on firms' expectation formation at the micro level is scarce: Anderson et al. (1954) conduct a very early study on qualitative expectation errors using the first few installments of the IFO Business Climate Survey. König et al. (1981) test a number of different models of expectation formation and also use models that jointly determine production and price expectations. Nerlove (1983) uses German (from IFO) and French data to analyze firms' expectations about idiosyncratic firm variables (such as prices, demand, etc.). Flaig and Zimmermann (1983) provide evidence that the production expectations of firms do not capture all available information. Zimmermann and Kawasaki (1986), using IFO price expectation and realization data, test whether firms are rational about how the market prices of their own commodities will develop. Zimmermann (1986) employs log-linear probability models to test the rationality of German survey expectations. All these studies have in common that they reject the strict rational expectations assumption.[4]

[2]According to Reis (2006) these costs include payments to outside consultants and/or the effort expended in supervising the production and monitoring of sales figures.

[3]Veldkamp and Wolfers (2007) provide a model in which most firms prefer public information since it is cheaper. Therefore, it can be rational for these firms not to include idiosyncratic information in their expectation formation process. Heid and Larch (2011) stress the importance of social linkages across firms in the expectation formation process.

[4]There are, of course, studies that use other survey data. A survey can be found in Weale and Pesaran (2006).

The chief reason for such slow progress of research into firms' expectation formation has to do with the qualitative nature of survey data.[5] Typically, firms report qualitative answers by choosing one of the following: "we expect our production to increase, decrease, stay the same over the next three months." While useful, in that more firms are inclined to participate in a survey when the information demands are low, qualitative information has limits, particularly when forecasting errors need to be aggregated over time in order to measure the long-run average forecasting errors of firms and possible biases therein. How does one aggregate a qualitative forecasting error of $+1$ (up) today and -1 (down) tomorrow?[6] Therefore, many studies rely on aggregate survey responses (for business survey data see e.g. Tompkinson and Common, 1983; for consumer survey data see e.g. Bovi, 2009). However, as highlighted by Souleles (2004) and Zimmermann and Kawasaki (1986), such an analysis can be misleading due to aggregation problems and because different economic agents have different information sets.

I overcome this problem in my analysis by using micro data from the manufacturing part of the IFO-BCS from 1980 onward. I then combine the qualitative three-month ahead production outlook from the monthly survey with the quantitative change in percentage capacity utilization as proposed by Bachmann and Elstner (2011). This procedure allows me to construct a panel of quarterly production expectation errors for a 30-year period. With this panel data set I can analyze firm histories of expectation errors, which is a unique advantage compared to previous studies and investigations based on consumer confidence surveys (e.g. Souleles, 2004).

To analyze distorted belief in firms' expectations, I compute long-run averages of expectation errors for those firms that have at least eight years of observations. I find that 25 percent of firms *consistently* overpredict their one-quarter ahead upcoming production by at least 1.2 percent in a conservative estimate or by at least 4.2 percent in an upper bound estimate. To investigate whether these firms really have a bias in their expectations, I test for *each* firm whether its average expectation error is significantly different from zero. In my conservative estimate only 1 percent of all firms have distorted beliefs which is good news for the rational expectations hypothesis. However, extending the analysis to a broader measure of quantitative expectation errors reveals that more than 18 percent of firms have distorted beliefs. To see whether firms efficiently exploit their information sets I check for

[5]There also are quantitative business surveys. For example, de Leeuw and McKelvey (1981, 1984) use a BEA survey of business expenditures on plant and equipment in the United States to study quantitative annual expectation data about aggregate prices and find that firms do not have rational expectations.

[6]This is a problem, even when qualitative expectation data predict quantitative ex-post data rather well, which is what Lui et al. (2008) find, using survey data from the 'Confederacy of British Industry' business survey and ex-post administrative data.

correlation between their firm-specific expectation errors and their information sets. These computations show that a considerable proportion of firms do not process their information sets efficiently. In my conservative estimate of expectation errors about two-thirds of the firms in my sample have rational expectations, i.e. their expectations are unbiased and the firms use all available information efficiently. When the expectation error estimate is defined less conservatively, this number reduces substantially to 33.2 percent. Thus, there is evidence that heterogeneous firms form their expectations in heterogeneous ways.

The remainder of the paper is structured as follows. The next section describes the IFO-BCS and the construction of the production expectation errors from it. Note that Section 4.2 is based on Bachmann and Elstner (2011). Readers already familiar with that work will lose nothing by skipping this section and proceeding directly to Section 4.3. Section 4.3 deals with the business cycle properties of my production expectation errors. Section 4.4 describes the empirical analysis at the firm level. Section 4.5 sets out and discusses the results. Section 4.6 provides a series of robustness checks. Section 4.7 concludes.

4.2 Evidence from the IFO Business Climate Survey

4.2.1 The IFO Business Climate Survey

The IFO Business Climate index is a much-followed leading indicator for economic activity in Germany. It is based on a firm survey which has been conducted since 1949 and, therefore, is one of the oldest and broadest monthly business confidence surveys available (see Becker and Wohlrabe, 2008, for details). Due to longitudinal consistency problems in other sectors and the availability of micro data in a processable form I limit my analysis to the manufacturing sector from 1980 until the present. From 1991 on, the sample includes East–German firms.

One of the IFO-BCS's main advantages is the high number of survey participants. The average number of respondents at the beginning of my sample is approximately 5,000; towards the end it is about 2,500.[7] Participation in the survey is voluntary and confidential. Thus, there is little incentive for firms to provide overoptimistic forecasts as a signal to investors. There is some fraction of firms that are only one time participants. However, conditional on staying two months in the survey, most firms continue on and this allows me to construct an unbalanced panel data set of expectation errors. For my narrow, very conservative definition of expectation errors the final baseline panel consists of 695 firms with at least 32 quarterly observations each; for a broader definition I follow 3,859 firms for again 8 years at least.

[7] The IFO-BCS is a survey at the product level, so that these numbers do not exactly correspond to firms.

4.2.2 Construction of Quantitative Production Expectation Errors

To construct firms' quantitative production expectation errors I would ideally need the following quantitative information about production expectations and realizations: "By how much do you expect your production to grow over the next quarter? By how much did your production grow in the preceding quarter?" To the best of my knowledge there is no firm survey that asks these questions for a long time horizon and repeatedly at underyearly frequencies. However, the quantitative quarterly supplement of the IFO survey allows me to construct - under certain assumptions - quantitative production expectations and quantitative production realizations. I am thus able to construct a panel of quarterly production expectation errors for the last thirty years.

Specifically, I use the following supplementary question about capacity utilization to compute production changes:[8]

Q 1 *"Supplementary Question: The utilisation of our production equipment for producing XY (customary full utilization = 100) currently amounts to..%."*

30	40	50	60	70	75	80	85	90	95	100	more than 100 % namely

I start from the following production relationship of an individual firm i:

$$y_{i,t}^{act} = u_{i,t} y_{i,t}^{pot}, \tag{4.1}$$

where $y_{i,t}^{act}$ denotes the firm's actual output, $y_{i,t}^{pot}$ its potential output level and $u_{i,t}$ the level of capacity utilization. Only $u_{i,t}$ is directly observable in the IFO-BCS. Taking the natural logarithm and the three-month difference, I get:[9]

$$\Delta \log y_{i,t}^{act} = \Delta \log u_{i,t} + \Delta \log y_{i,t}^{pot}. \tag{4.2}$$

Under the assumption that potential output remains constant, $\Delta \log y_{i,t}^{pot} = 0$, percentage changes in actual output can be recovered from percentage changes in capacity utilization.

[8]Here I provide a translation, for the German original see Appendix A1.

[9]Time intervals are months. For me to construct an expectation error in t, I need an observation for capacity utilization in t and $t-3$.

To implement this idea I restrict the analysis to firms of which I can reasonably expect that they did not change their production capacity in the preceding quarter, making use of the following two questions in the IFO-BCS:

Q 2 *"Expectations for the next three months: Employment related to the production of XY in domestic production unit(s) will probably increase, roughly stay the same, decrease."*

Q 3 *"Supplementary Question: We evaluate our technical production capacity with reference to the backlog of orders on books and to orders expected in the next twelve months as more than sufficient, sufficient, insufficient."*

Given hiring frictions in the labor market I view a firm's expectation, stated in $t-3$, that its employment level will remain the same in the next three months as highly indicative that its productive capacity did not change between $t-3$ and t. Similarly, given capital adjustment frictions a firm's statement, again in $t-3$, that its technical production capacity is sufficient for the future incoming orders suggests that this firm has no reason to change its production capacity in the near future. To be conservative I require a firm satisfy both criteria in $t-3$ for me to assume that its production capacity has not changed between $t-3$ and t. In this case, I use the quarterly percentage change in capacity utilization in t as a proxy for the quarterly percentage change in production in t. The existence of non-convex or kinked adjustment costs for capital and labor adjustment as well as time to build (see Davis and Haltiwanger, 1992, as well as Doms and Dunne, 1998) make this a reasonable assumption.

To derive production expectation errors I also need information on firms' production expectations. This allows me to compute production *surprises* out of mere production *changes*. In the IFO-BCS firms report only qualitative production expectations:

Q 4 *"Expectations for the next three months: Our domestic production activities with respect to product XY will (without taking into account differences in the length of months or seasonal fluctuations) increase, roughly stay the same, decrease."*

Qualitative expectations have a built-in asymmetry in the sense that the middle category also constitutes a quantitative expectation, zero change, whereas the increase and decrease category conveys no quantitative information. I therefore proceed in two steps. First, I consider only firms whose answer to Q4 is that their production level, $y_{i,t}^{act}$, will not change in the next three months. Under the assumption that $y_{i,t}^{pot}$ remains constant over this time period, all $\Delta \log u_{i,t}$ are automatically expectation errors. In a second step I extend my analysis to arbitrary qualitative production expectations. This will give me a broader picture of expectation errors, albeit with the added cost of more assumptions.

I also clean my sample from firm-quarter observations with extreme capacity utilization outliers, i.e. those that exceed 150%, and from firm-quarter observations with inconsistent statements. To determine the latter I consider the following monthly qualitative IFO-BCS question concerning actual production changes in the months t, $t-1$, $t-2$:

Q 5 *"Trends in the last month: Our domestic production activities with respect to product XY have (without taking into account differences in the length of months or seasonal fluctuations) increased, roughly stayed the same, decreased."*

I drop all observations as inconsistent in which firms report a strictly positive (negative) change in $\Delta \log u_{i,t}$ and no positive (negative) change in Q5 in the last 3 months. For firms that report $\Delta \log u_{i,t} = 0$, I proceed as follows: Unless firms in Q5 either answer three times in a row that production did not change, or they have at least one "Increase" *and* one "Decrease" in their three answers, I drop them as inconsistent.

4.2.2.1 Quantitative Expectation Errors under Constant Production Expectations

If the production capacity can be assumed not to have changed in the preceding quarter, and if no change in production was expected three months prior, a change in capacity utilization, $\Delta \log u_{i,t}$, is also a production expectation error of firm i in month t. I first restrict my analysis to the subset of firm-quarter observations that satisfy these assumptions. For this case, Figure 4.1 illustrates the move from capacity utilization changes to production expectation errors.

Figure 4.1: Link between Capacity Utilization and Production Expectation Errors

Prerequisite: Firm i passes the outlier and inconsistency test (Q1 and Q5)

Firm i has an observation for a change in capacity utilization $\Delta \log u_{i,t}$ (Q1)

No change in potential output $\Delta \log y_{i,t}^{pot}$

Firm i answers "Constant employment expectations" (Q2)

and "Technical production capacity is sufficient" (Q3) in $t-3$

$\Delta \log u_{i,t}$ is a production change $\Delta \log y_{i,t}^{act}$

No expected production change

Qualitative production expectations in $t-3$ are constant (Q4)

$\Delta \log u_{i,t}$ is a quantitative production expectation error

Notes: The time dimension t is measured in months.

Figure 4.2 illustrates the exact timing of the questions in the Ifo-BCS that I use to compute production expectation errors. As a first pass I consider only firms which state in period $t-3$ that their production level, employment level and technical production capacity will remain the same in the next three months. Then I compute $\Delta \log u_{i,t}$ three months later in t. These $\Delta \log u_{i,t}$ constitute my narrow definition of production expectation errors. I denote them by $FE\_BCS_{i,t}^{narrow}$.

Figure 4.2: Derivation of Production Expectation Errors under Constant Production Expectations - Timing

Prerequisite: Firm i passes the outlier and inconsistency test (Q1 and Q5)

Firm i states in period $t-3$:
Constant employment expectations (Q2)
Sufficient technical capacity (Q3)
Constant production expectations (Q4)

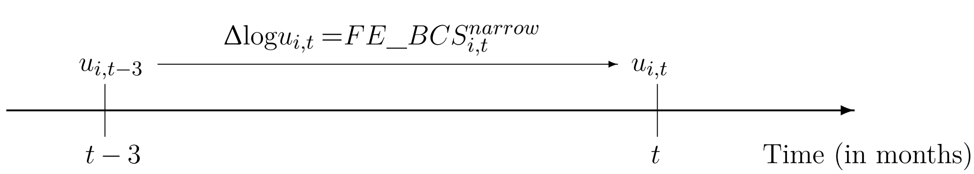

4.2.2.2 Quantitative Expectation Errors under General Production Expectations

The derivation of quantitative production expectation errors for firms with increasing or decreasing qualitative production expectations in Q4 requires additional assumptions. I take the following three steps. First, I define qualitative indices of production changes and expectations errors using questions Q4 and Q5. Then, changes in capacity utilization are mapped into qualitative production changes. In a final step, I convert these quantitative production changes into quantitative production expectation errors. The expectation errors for firms with constant production expectations remain the same as in the previous section.

I begin by defining the following firm-specific activity variable $REALIZ_{i,t}$ as the sum of the Increase-instances minus the sum of the Decrease-instances in question Q5 over the last three months going backward from t. $REALIZ_{i,t}$ can have seven possible values that live in the interval [-3,3]. A qualitative production expectation error is then computed as follows:

Table 4.1: Possible Qualitative Expectation Errors

		$Experror_{i,t}$
Expected $Increase_{t-3}$	$REALIZ_{i,t}>0$	0
Expected $Increase_{t-3}$	$REALIZ_{i,t} \leq 0$	$(REALIZ_{i,t}-1)$
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}>0$	$REALIZ_{i,t}$
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}=0$	0
Expected $Unchanged_{t-3}$	$REALIZ_{i,t}<0$	$REALIZ_{i,t}$
Expected $Decrease_{t-3}$	$REALIZ_{i,t}<0$	0
Expected $Decrease_{t-3}$	$REALIZ_{i,t} \geq 0$	$(REALIZ_{i,t}+1)$

Notes: Rows refer to the qualitative expectations in month $t-3$ (Q4).

$Experror_{i,t}$ ranges from $[-4,4]$, where, for instance, -4 indicates a really negative forecast error: the company expected production to increase over the next three months, yet every single subsequent month production actually declined.

Next I compute for all firms with a given value of $REALIZ_{i,t}$ in time t the average $\Delta \log u_{i,t}$, i.e. the cross-sectional average change in capacity utilization. Again, to ensure that I can treat utilization changes as production changes only those firms are considered that state three months before that their future employment levels remain the same and that their technical production capacities are sufficient. I compute this mapping between $REALIZ_{i,t}$ and average production changes for each point in time.[10] Figure 4.3 illustrates the timing of survey questions that are used to compute this mapping. Figure A1 in Appendix A2 depicts

[10] I also considered a firm size specific and an industry specific mapping, without much changes to my results.

the results. For instance, it shows that firms with $REALIZ_{i,t} = 1$ in the first quarter of 1980 maps into an average change in production of approximately 5 percent.

Figure 4.3: Mapping between Qualitative and Quantitative Production Changes

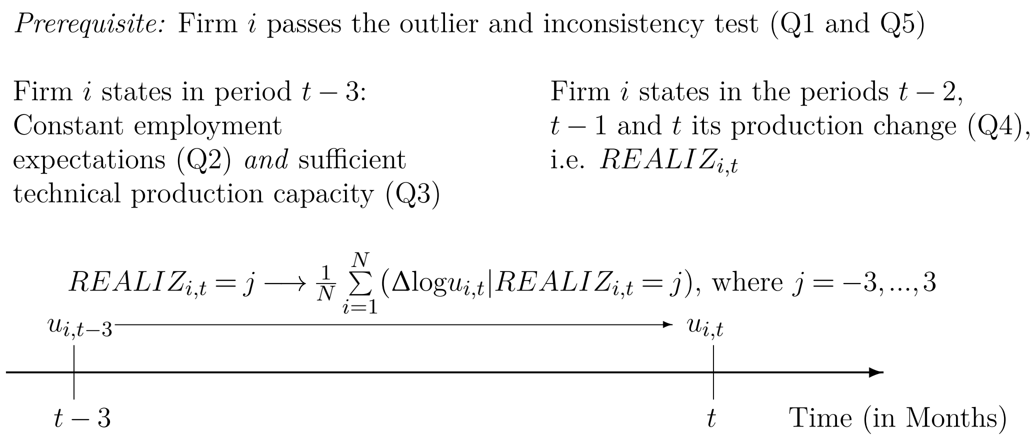

I subject this mapping to a couple of plausibility tests.[11] For firms with constant production expectations I have a simple plausibility test for my procedure. Suppose I know with $FE\_BCS_{i,t}^{narrow}$ the "true" forecast errors under constant production expectations. I can also use the mapping strategy for the firms with constant qualitative production expectations, in which case the average percentage change in capacity utilization for a given $REALIZ_{i,t}$ would be another estimate for their production expectation error. I can then compare the "true" forecast errors with the forecast errors determined in the mapping procedure.

I do this by computing a quarterly time series of the cross-sectional means and standard deviations for these two types of expectation errors. The time series correlation coefficients between the two series are high: 0.97 for the average expectation errors and 0.84 for their standard deviation. This means that *over time* the mapping strategy captures the first and second moment of the expectation error distribution rather well. However, *on average* the "true" expectation errors are much more disperse than those based on the mapping strategy. Since the time series behavior of the cross-sectional standard deviation is very similar, I view this as a scaling problem, resulting from the discretization of expectation errors in the mapping strategy. As a consequence, I rescale the quantitative production change values that $REALIZ_{i,t}$ is mapped into by a constant.[12]

[11] The pooled Spearman correlation coefficient between $REALIZ_{i,t}$ and $\Delta \log u_{i,t}$ for firms with constant employment and technical capacity expectations is 0.72. The pooled Kendall's tau between $sign(REALIZ_{i,t})$ and $sign(\Delta \log u_{i,t})$ is 0.67.

[12] This constant is calculated by dividing the time average of the cross-sectional standard deviations of

In the last step the qualitative production expectation errors are converted into quantitative production expectation errors. This is simply done by assigning to each value of $Experror_{i,t}$ the quantitative equivalent of $REALIZ_{i,t}$. The extreme cases of $Experror_{i,t}$, +4 and -4, are added by extrapolation.[13] I denote my final measure of the quantitative production expectation errors under general production expectations by $FE\_BCS_{i,t}^{broad}$.

4.3 Business Cycle Properties of Expectation Errors

To discover whether firms' expectations are biased and to lend support to my measures of quantitative expectation production errors, it will be useful to consider the business cycle properties of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$. To this end, I compute quarterly time series of the cross-sectional means and standard deviations for both measures and compare them with the quarter-on-quarter percent changes in manufacturing production.

The upper two panels of Figure 4.4 show the quarterly cross-sectional means of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{narrow}$ together with the quarter-on-quarter percent changes in manufacturing production. The lower panels show the quarterly cross-sectional standard deviations.[14] At first glance, the cross-sectional means of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ seem strongly procyclical. The very high contemporaneous correlation coefficients between these expectation error time series and manufacturing production, which are shown in Table 4.2, provide further evidence of strong comovement. However, the cross-sectional means of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ are, on average, negative. This suggests that a large proportion of firms tends toward too optimistic and, therefore, biased expectations, i.e. they have negative expectation errors. To prove this conjecture, I perform two tests that are outlined in Souleles (2004).

First, I test whether $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ have significant time effects, which is easily accomplished by regressing both expectation error time series on year dummies. If firms are too optimistic, the time effects should be mostly negative and significantly different from zero.[15] Figure 4.5 depicts the point estimates together with the 95 percent confidence intervals. To construct the latter, I use robust standard errors that are adjusted for cross-correlations of the expectation errors within a quarter. F tests for both

$FE\_BCS_{i,t}^{narrow}$ by the corresponding value derived from the mapping procedure on the same subset of observations. Its value is 1.7.

[13] The extrapolation procedure tries to capture the change in the differences of the quantitative equivalents of $REALIZ_{i,t}$. It concerns only 860 of the 319,172 one-quarter-ahead production expectation errors. That corresponds to 0.27 % of all observations. Neglecting these extreme values in our upcoming quantitative analysis would not alter our results.

[14] All time series are seasonally adjusted using the Census X-12-ARIMA seasonal adjustment method.

[15] I obtain similar results for quarterly time dummies.

$FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ suggest that significant time effects are present. Figure 4.5 also shows that these time effects have been *significantly* negative for almost all years. Even during boom periods such as 2006 and 2007, a great many firms tend to overestimate their upcoming production.

In the second test, I determine the mean μ of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$, respectively, by regressing all their observations on a constant. For both cases, I obtain significant negative estimates, further confirming the initial conjecture that a large proportion of firms tends towards too optimistic expectations. However, as highlighted by Souleles (2004) and Zimmermann and Kawasaki (1986), the analysis at the macro level can be misleading due to aggregation problems and because different economic agents have different information sets. Therefore, it is necessary to analyze the expectation errors directly at the firm level in order to detect expectational biases.

As additional support for my quantitative production expectation errors I present results for the cross-sectional standard deviations of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$. These time series are countercyclical and have a lag of roughly two quarters. These negative correlations are in line with the uncertainty literature (see e.g. Bachmann and Bayer, 2010, Bloom, 2009, and Bloom et al., 2010). This result makes me confident that the proxies I use for qualitative production errors are appropriate. In a next step, I shed more light on the expectational biases of firms by analyzing firm-specific expectation errors at the firm level.

Table 4.2: Cyclical Properties of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$

Variable in $t+j$	Manufacturing production in t								
	-4	-3	-2	-1	0	1	2	3	4
Cross-sectional means of:	Leading Property of Variable					Lagging Property of Variable			
$FE\_BCS_{i,t+j}^{narrow}$	-0.12	-0.02	0.10	0.48	0.74	0.53	0.35	0.14	-0.11
$FE\_BCS_{i,t+j}^{broad}$	-0.08	0.17	0.30	0.42	0.61	0.38	0.12	-0.03	-0.14
Cross-sectional standard deviations of:									
$FE\_BCS_{i,t+j}^{narrow}$	0.25	0.27	0.12	-0.04	-0.22	-0.32	-0.37	-0.31	-0.24
$FE\_BCS_{i,t+j}^{broad}$	0.23	0.32	0.13	0.00	-0.08	-0.25	-0.39	-0.38	-0.25

Notes: This table provides the unconditional correlations between the quarter-on-quarter percent changes in manufacturing production in period t and a variable in period $t+j$. All time series are seasonally adjusted.

Figure 4.4: Comparison of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ with Manufacturing Production

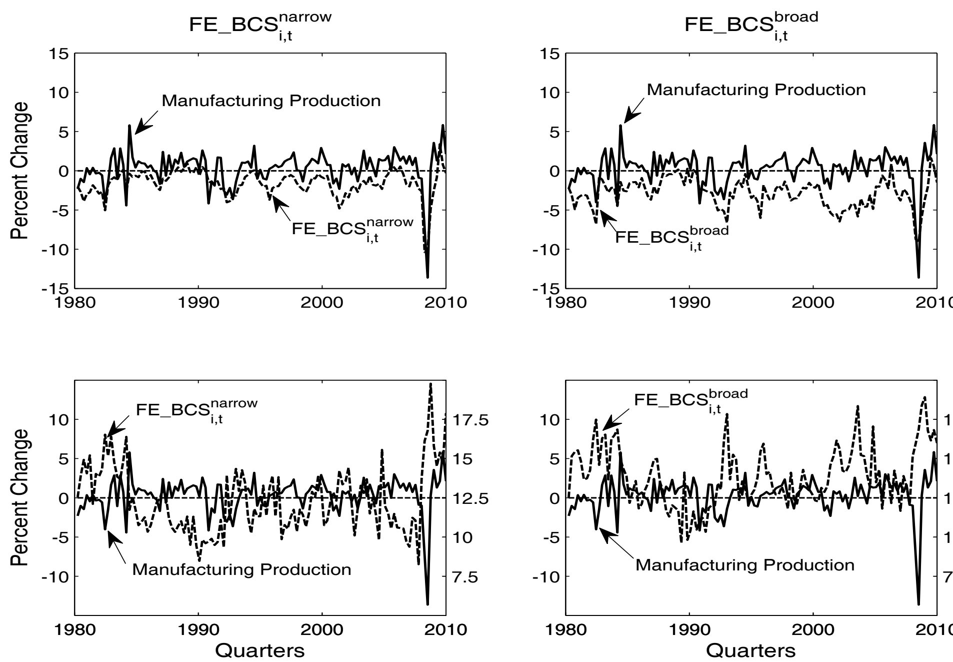

Notes: The upper panels show the quarterly cross-sectional means of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ together with quarter-on-quarter percent changes in manufacturing production from the German Federal Statistical Agency. The lower panels show the quarterly cross-sectional standard deviations of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ (left scale) together with the quarter-on-quarter percent changes in manufacturing production (right scale). All time series are seasonally adjusted.

Figure 4.5: Biasedness of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ on the Macro Level

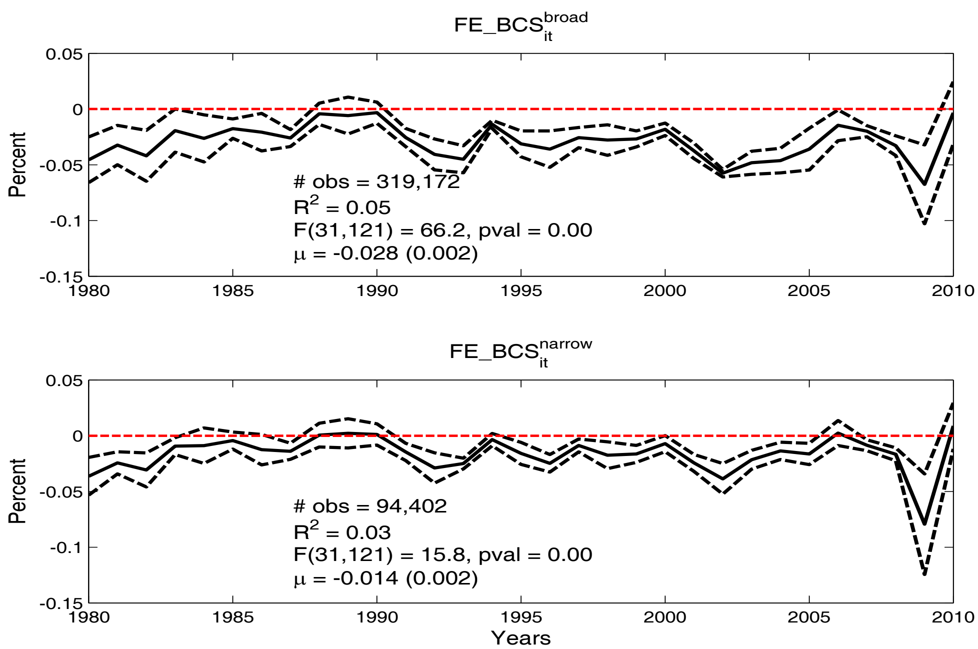

Notes: The upper panel shows the annual time effects of $FE\_BCS_{i,t}^{broad}$. The lower panel does the same for $FE\_BCS_{i,t}^{narrow}$. The results are derived by a linear regression of either $FE\_BCS_{i,t}^{narrow}$ or $FE\_BCS_{i,t}^{broad}$ on year dummies (no constant included in the regression). The solid lines represent the point estimates of the year dummies and the dashed lines depict 95 percent confidence intervals. To test the joint significance of the year dummies, an F test is employed. For each category of expectation error I determine the mean μ by regressing respectively all observations on a constant (standard error of the point estimate in parentheses). I use robust standard errors that are adjusted for cross-correlations of the expectation errors within a quarter.

4.4 Firm-Level Heterogeneity in Expectation Errors

4.4.1 Statistics at the Firm-Level

In this section I provide summary statistics concerning the average firm-specific production expectation errors. I restrict my sample to firms that have at least 32 observations or eight years of expectation errors: $FE\_BCS_{i,t}^{narrow}$ or $FE\_BCS_{i,t}^{broad}$. The average firm-specific expectation errors are denoted by $AFE\_BCS_{i}^{narrow}$ and $AFE\_BCS_{i}^{broad}$. Table 4.3 displays the distributions of firms' average expectation errors; Figure 4.6 provides a graphical illustration. Note that positive values of $AFE\_BCS_{i}^{narrow}$ or $AFE\_BCS_{i}^{broad}$ indicate that a firm was on average too pessimistic in the sense that its predicted production changes were on average lower than its actual production changes. The cross-sectional means are based on numbers different from those reported in Figure 4.5, since there I computed the average of all firms with expectation errors, not just for those with at least 32 observations. Especially in the case of $AFE\_BCS_{i}^{broad}$ the distribution is skewed toward negative values. This observation is supported by the fact that at least 25 percent of all firms have long-run averages of expectation errors that are too optimistic by 4 percent or more.

Table 4.3: Firm-Specific Average Production Expectation Errors (IFO-BCS)

Statistics	$AFE\_BCS_{i}^{narrow}$	$AFE\_BCS_{i}^{broad}$
Obs	695	3859
Mean	-0.0064	-0.0265
Std.Dev.	0.0161	0.0316
Percentiles		
5th	-0.0339	-0.0859
10th	-0.0242	-0.0672
25th	-0.0119	-0.0418
50th	-0.0030	-0.0212
75th	0.0013	-0.0060
90th	0.0068	0.0063
95th	0.0119	0.0140

Notes: This table provides a summary of the distributions of $AFE\_BCS_{i}^{narrow}$ and $AFE\_BCS_{i}^{broad}$.

Figure 4.6: Firm-Specific Average Production Expectation Errors (IFO-BCS)

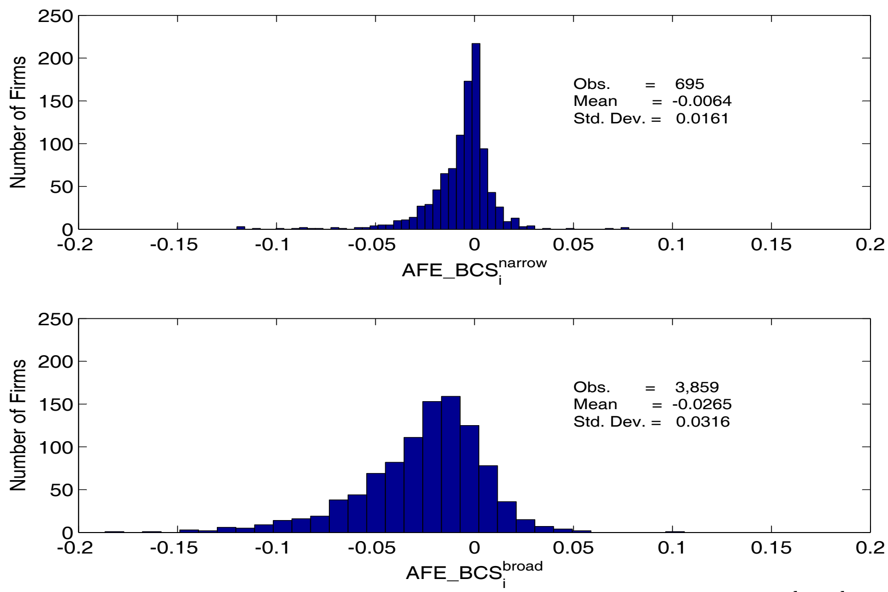

Notes: This figure displays the distributions of $AFE\_BCS_i^{narrow}$ and $AFE\_BCS_i^{broad}$.

4.4.2 Unbiasedness and Efficiency of Firms' Expectations

According to Muth (1961) firms have rational expectations if their expectations are unbiased and they use all available information efficiently. Firms' expectations are unbiased if their subjective probabilities with respect to future economic states are not distorted. This implies that their long-run average expectation errors are not significantly different from zero. However, even if a firm passes the test for unbiased expectations, it does not mean that this firm has rational expectations; efficient information processing is also necessary. Efficiency means that there is no correlation between firm-specific expectation errors and the firm's information set. In the following, I conduct three distinct tests to check for unbiasedness and efficiency of firms' expectations.

Other work follows the classical approach of Mincer and Zarnowitz (1969) to test for unbiased expectations, using the following regression equation:

$$x_t = b_0 + b_1 x_t^e + u_t,$$

where x_t denotes the realized value of some variable, say, firm production in month t, and x_t^e denotes the expected value of this variable, say, expected firm production for month t. In this setup, expectations are assumed to be unbiased, if the intercept coefficient b_0 is zero *and* the slope coefficient b_1 is one.

Due to a lack of quantitative expectation data, I use a different, but very similar, approach in testing for biased expectations: for each firm I determine $AFE\_BCS_{i,t}^{narrow}$ and $AFE\_BCS_{i,t}^{broad}$, respectively, by regressing all its observations of $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$ on a constant. Then I use a two-sided t test to assess whether the firm-specific average expectation error is significantly different from zero. If this is the case, I conclude that this firm has distorted beliefs. For statistical reasons, the level of significance is defined to be 1 percent. Under the assumption that all firms can be treated independently from each other and in a world with perfect rational expectations, the test would wrongly find that 1 percent of all firms have biased expectations. I want to ensure that my results are not susceptible to this error.

The test is carried out on the restricted samples of $FE\_BCS_{i,t}^{narrow}$ or $FE\_BCS_{i,t}^{broad}$, i.e. I consider only those firms with at least 32 observations of the respective category of expectation error. The upper panel of Table 4.4 displays the results. For $FE\_BCS_{i,t}^{narrow}$, I identify distorted beliefs for only 9 out of 695 firms. This corresponds to 1.3 percent of all firms and provides very little evidence of distorted beliefs at the firm level. This picture, however, changes dramatically when looking at the results of the broader measure $FE\_BCS_{i,t}^{broad}$. In this case, the number of firms with biased expectations is considerably larger and almost 20 percent of all considered firms have distorted beliefs.

Table 4.4: Firms with Rational Expectations - Results

	Category of Expectation Error: $FE_{i,t}^{narrow}$		$FE_{i,t}^{broad}$	
	Number of Firms	Proportion of Firms	Number of Firms	Proportion of Firms
Test of Unbiasedness				
Unbiased expectations	686	98.7%	3,137	81.3%
Biased expectations	9	1.3%	722	18.7%
Tests of Efficiency				
(1) Autocorrelation in Expectation Errors				
No autocorrelation	650	93.5%	3,609	93.5%
Autocorrelation	45	6.5%	250	6.5%
(2) Correlation of Expectation Errors with Available Information Set				
No correlation	505	72.7%	1,474	38.2%
Correlation	190	27.3%	2,385	61.8%
Firms with Rational Expectations				
Rational expectations	475	68.3%	1,282	33.2%
No rational expectations	220	31.7%	2,577	66.8%

Next, I use two tests to check for efficiency. The goal of the first is to discover whether the expectation errors of a single firm are uncorrelated over time. I thus estimate autoregressive time series models for $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$, respectively, with two lags for each firm. The decision to use a lag length of two is due to incomplete firm histories of expectation errors. Especially in the case of $FE\_BCS_{i,t}^{narrow}$ the number of observations decreases dramatically the higher the number of lags. I use an F test to find out whether the expectation errors of a single firm are autocorrelated.

The second test analyzes for each firm whether its expectation errors are uncorrelated with its available information set at the time of its expectation report. I use the answers to four survey questions to capture firm-specific information sets. The first two questions concern expected production (Q4) and actual production compared to the last month (Q5). Furthermore, answers to the following two questions are used to control for information regarding the firm's business situation:

Q 6 *"Expectations for the next six months: Our business situation with respect to XY will in a cyclical view: improve, remain about the same, develop unfavourably."*

Q 7 *"Current situation: We evaluate our business situation with respect to product XY as: good, satisfactory, unsatisfactory."*

The answers to all four questions are qualitative and have three possible outcomes. While the answer categories are ordered, I split each of them into two dummy variables, which is a standard procedure. I run regressions with $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$, respectively, as the dependent variable. The set of independent variables consists of a constant and the firm's answers to the four questions at the time of its expectation report. Again, I use an F test to decide whether the explanatory variables are jointly significant. If this is the case, I conclude that the firm does not process its available information efficiently. I again impose a significance level of 1 percent.

The middle panel of Table 4.4 shows the test results. For both categories of expectation errors, the proportion of firms that feature autocorrelated expectation errors is 6.5 percent. Thus, it can be concluded that the majority of firms exploit their information efficiently. However, the same conclusion cannot be made when looking at the results of the second test. In the case of $FE\_BCS_{i,t}^{broad}$ more than 60 percent of all firms show a significant correlation between their expectation errors and their information sets. Thus, many firms do not form expectations in an efficient and, therefore, rational way. This result is supported by the figures in the last panel of Table 4.4, which show the proportions of firms that have rational expectations. I define firm expectations as strictly rational if they pass the tests for unbiasedness and efficiency. I observe that in the case of $FE\_BCS_{i,t}^{narrow}$ about two-thirds of the firms in my sample have rational expectations. However, in the case of $FE\_BCS_{i,t}^{broad}$ this number decreases substantially to 31.7 percent. Thus, there is evidence that heterogeneous firms form their expectations in heterogeneous ways, i.e. firms differ with respect to their beliefs and their ability to process information.

4.5 Systematic Relationships in Expectation Errors

4.5.1 Model Specifications

The aim of this section is to discover whether there are systematic relationships between firm-specific characteristics and distorted beliefs or non-rationality of firms. The following estimations are agnostic and provide stylized facts, e.g. whether smaller firms are more likely to have distorted beliefs.

I first analyze the relationship between firm-specific characteristics and distorted beliefs. In the first model specification, the independent variable y_i denotes a dummy variable that assigns to a single firm i the value +1 if it has no distorted beliefs and 0 otherwise. y_i is constructed for both $FE\_BCS_{i,t}^{narrow}$ and $FE\_BCS_{i,t}^{broad}$. I assume that y_i^* is a latent quantitative variable that relates to the corresponding qualitative variable y_i and takes the following form:

$$y_i^* = \beta_0 + \beta_1 DEXPORT_i + \beta_2 DSIZE_i + \beta_3 DIND_i + \beta_4 DSTATE_i + u_i, \tag{4.1}$$

where $DSTATE_i$ defines a set of federal state dummies, $DSIZE_i$ denotes firm size dummies, $DIND_i$ summarizes a set of dummies defining the sectoral classification of the single firm and $DEXPORT_i$ is a dummy variable that takes the value of +1 for an exporting firm; 0 otherwise.[16] Concerning firm size, I know whether the number of employees in firm production is less than 50, between 50 and 199, between 200 and 499, between 500 and 999, or equal to or more than 1,000. From this, five dummy variables are constructed. The dummy variables summarized by $DIND_i$ allocate each firm to one of the following 14 manufacturing subsectors: food and tobacco; textiles and textile products; leather; cork and wood products except furniture; pulp, paper, publishing and printing; refined petroleum products; chemicals and chemical products; rubber and plastic products; other non-metallic mineral

[16] I use the monthly IFO-BCS question concerning expected export trade. If a single firm has stated in more than half its answers that it does not export, that firm receives the value 0.

products; basic and fabricated metal products; electrical and optical equipment; transport equipment; furniture and jewelry and machinery and equipment. The last set of dummy variables, represented by $DSTATE_i$, indicates the German federal state in which the firm operates.

In the second model specification, I focus on correlations between firm-specific characteristics and non-rationality. I define firms with rational expectations as those firms that have passed both the test for unbiasedness and the tests for efficiency. My independent variable is again a dummy variable that assigns to a single firm i the value +1 if it has rational expectations and 0 otherwise. The other assumptions concerning the estimation equation are the same as in Equation (4.1).

For both model specifications I estimate probit models with robust standard errors. "Basic and fabricated metal products", the federal state "Bavaria" and firm size of 200 to 499 employees are the reference categories for both model specifications.[17]

4.5.2 Results

The left-hand side of Table 4.5 provides the estimation results for the relationships between firm-specific characteristics and distorted beliefs. For the sake of brevity I do not show the point estimates of the federal state dummies, $DSTATE_i$. Several explanatory dummies are not considered in the regressions since they do not contribute to the explanation of the dependent variable. This problem is particularly evident in the case of $FE\_BCS_{i,t}^{narrow}$, for which only 9 out of 695 firms feature distorted beliefs. Hence, the information content of this regression is limited and, therefore, in the following I discuss only the results for $FE\_BCS_{i,t}^{broad}$.

To begin with, the χ^2 tests reported in the last rows of Table 4.5 indicate that the explanatory variables have a joint impact on the outcome of whether a firm has distorted beliefs. I also report separate results of joint significance for firm size, industrial sector and federal state dummies. For each of these groups I find good evidence that they can explain what kinds of firms are more likely to feature distorted beliefs. Even though there are significant correlations between firms' expectational biases and firm-specific characteristics, their total impact is rather small, as can be inferred from the very low value for the "Pseudo R^2".

Looking at the individual coefficients reveals that smaller firms are more likely to have distorted beliefs than are larger firms. A firm with less than 50 employees has a 7 percent lower probability of having no distorted beliefs than the reference firm with 200 to 499 employees.

[17] The use of usual robust standard errors is problematic since I face a "generated regressand" problem. However, it is beyond the scope of this paper to employ a suitable bootstrap method to account for this issue. Thus, inferences based on these standard errors should be viewed with caution.

Table 4.5: Results

		Firms with undistorted beliefs:				Firm with rational expectations:			
		(1) $FE\_BCS_{i,t}^{narrow}$		(2) $FE\_BCS_{i,t}^{broad}$		(1) $FE\_BCS_{i,t}^{narrow}$		(2) $FE\_BCS_{i,t}^{broad}$	
		Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.
DEXPORT		0.014	0.021	0.006	0.017	-0.013	0.055	0.047	0.021**
DSIZE	<50 employees	0.025	0.011**	-0.068	0.022***	-0.073	0.070	-0.023	0.024
	$50-199$ employees	0.018	0.015	-0.032	0.018*	-0.043	0.048	-0.020	0.021
	$500-999$ employees	0.002	0.022	0.040	0.023*	0.010	0.064	-0.016	0.029
	$\geq 1,000$ employees			0.066	0.023***	0.040	0.075	-0.024	0.031
DIND	Food and tobacco			-0.076	0.037**	-0.005	0.089	-0.001	0.039
	Textiles and textile products			-0.069	0.032**	0.084	0.069	0.059	0.035
	Leather			0.032	0.050	0.095	0.116	-0.031	0.061
	Cork and wood products except furniture	-0.058	0.064	-0.076	0.042*	0.059	0.090	0.026	0.046
	Furniture and jewelery			-0.048	0.037	-0.104	0.124	0.027	0.041
	Pulp, paper, publishing and printing	-0.019	0.027	-0.060	0.028**	-0.028	0.070	0.007	0.030
	Electrical and optical equipment			-0.042	0.029	0.138	0.064**	0.022	0.032
	Chemical products			-0.060	0.055	-0.072	0.129	0.002	0.057
	Rubber and plastic products	-0.019	0.040	-0.099	0.037***	-0.152	0.097	-0.030	0.036
	Other non-metallic mineral products	0.003	0.023	-0.063	0.036*	0.118	0.074	0.034	0.039
	Machinery and equipment			0.010	0.025	-0.006	0.067	0.010	0.030
	Transportation equipment			0.038	0.046	0.069	0.121	0.076	0.057
	Oil industry			-0.116	0.166			0.121	0.165
DSTATE		yes		yes		yes		yes	
Observations		281		3859		692		3859	
log likelihood		-34.15		-1805.25		-417.12		-2434.78	
Pseudo R^2		0.14		0.03		0.04		0.01	
Joint Significance of (χ^2 test):									
All Variables [pval]		0.00***		0.00***		0.33		0.34	
DSIZE [pval]		0.17		0.00***		0.63		0.86	
DIND [pval]		0.52		0.02**		0.12		0.75	
DSTATE [pval]		0.13		0.04**		0.65		0.30	

Notes: This table provides the estimation results of the model specifications described in the text. Both model specifications use probit models for the estimations and all coefficients are expressed as marginal effects. Some variables (missing entries) are not able to explain the variation of the dependent variable and thus are not considered in the regression. “Basic and fabricated metal products", the federal state “Bavaria” and firm size of 200 to 499 employees are considered the reference categories for both model specifications. Robust standard errors in parenthesis. ***: p<0.01; **: p<0.05; *: p<0.1.

In contrast, a firm with more than 1,000 employees has a 7 percent higher probability of having no distorted beliefs compared to the reference firm. This result reflects the idea that larger firms put more resources into analyzing their current and upcoming business environment than do smaller firms. Concerning industrial sectors I also find significant relationships, but it is difficult to interpret them. In summary, the first model specification provides evidence that there are significant correlations between distorted beliefs and firm-specific characteristics. It appears that it is especially differences in firm size that explain these correlations.

The right-hand side of Table 4.5 shows the results for the relationships between firm-specific characteristics and non-rationality. Briefly, there are almost no significant and economically meaningful relationships for either $FE\_BCS_{i,t}^{narrow}$ or $FE\_BCS_{i,t}^{broad}$.

4.6 Robustness Checks

To this point in the analysis, I find that for $FE\_BCS_{i,t}^{broad}$ a substantial proportion of firms have distorted beliefs and for both types of expectations errors at least 30 percent of all firms have no rational expectations. In the following I present three distinct robustness checks. As a first check, I retain inconsistent statements in my sample. It could be argued that these inconsistent statements indicate some form of irrationality. Thus, removal of these observations could exert a downward bias on the proportion of irrational firms and firms with distorted beliefs. After repeating the whole analysis I indeed find in the case of $FE\_BCS_{i,t}^{narrow}$ that the proportion of firms with irrational expectations increases. The results are shown on the left-hand side of Table 4.6. The figures for $FE\_BCS_{i,t}^{broad}$ remain roughly the same. In total, I find for both types of expectation errors that considerably less than half of all firms have no rational expectations.

The next robustness check concerns the number of observations that I require a firm to have in order to compute an average production forecast error. Instead of 32 observations, I now use a threshold of 16 observations or four years of quarterly quantitative production errors. This gives me a larger cross section of firms, but at the cost of including more firms with shorter histories. For this test, for $FE\_BCS_{i,t}^{narrow}$ the number of firms increases from 695 to 2,194 and for $FE\_BCS_{i,t}^{broad}$ I now consider 5,980 firms instead of 3,859. However, increasing the number of firms does not affect my results. The proportions of firms with distorted and irrational expectations remain largely unchanged.

Table 4.6: Robustness Checks - Results

	Including Inconsistent Statements				**At least 16 Observations**				**Qualitative Expectation Error**			
	Category of Expectation Error:				Category of Expectation Error:				Category of Expectation Error:			
	$FE_{i,t}^{narrow}$		$FE_{i,t}^{broad}$		$FE_{i,t}^{narrow}$		$FE_{i,t}^{broad}$		Narrow Definition		Broad Definition	
	Number of Firms	Proportion of Firms	Number of Firms	Proportion of Firms	Number of Firms	Proportion of Firms	Number of Firms	Proportion of Firms	Number of Firms	Proportion of Firms	Number of Firms	Proportion of Firms
Test of Unbiasedness												
Unbiased expectations	1,704	99.7%	3,622	79.2%	2,170	98.9%	5,090	85.1%	2,308	82.4%	3,253	84.3%
Biased expectations	5	0.3%	954	20.8%	24	1.1%	890	14.9%	494	17.6%	606	15.7%
Tests of Efficiency												
(1) Autocorrelation in Expectation Errors												
No autocorrelation	1,252	73.3%	4,284	93.6%	2,114	96.4%	5,585	93.4%	2,408	85.9%	3,588	93.0
Autocorrelation	457	26.7%	292	6.4%	80	3.6%	395	6.6%	394	14.1%	271	7.0
(2) Correlation of Expectation Errors with Available Information Set												
No correlation	989	57.9%	1,882	41.1%	1,554	70.8%	2,486	41.6%	1,516	54.1%	1,228	31.8%
Correlation	720	42.1%	2,694	58.9%	640	29.2%	3,494	58.4%	1,286	45.9%	2,631	68.2%
Firms with Rational Expectations												
Rational expectations	744	43.5%	1,605	35.1%	1,498	68.3%	2,197	36.7%	1,216	43.4%	1,108	28.7%
No rational expectations	965	56.5%	2,971	64.9%	696	31.7%	3,783	63.3%	1,586	56.6%	2,751	71.3%

As a last check, I review the concept of quantitative expectation errors. It may be that this is the wrong approach entirely and that the qualitative expectation errors defined in Table 4.1 should be used instead. Therefore, I conduct the same analysis with $Experror_{i,t}$ and a new variable $Experror_{i,t}^{nar}$. The latter variable is constructed for qualitative expectation errors with initially reported constant production expectations and allows a comparison with the results of $FE\_BCS_{i,t}^{narrow}$. Appendix A3 contains summary statistics concerning the business cycle properties of these qualitative expectation errors. However, even for these measures, it is evident that many firms do not form their expectations in a rational way. Again, I find that more than half of all the firms in my sample have no rational expectations for either $Experror_{i,t}^{nar}$ or $Experror_{i,t}$.

4.7 Conclusion

Are firms' expectations rational? For a large number of firms the answer is "no". To arrive at this conclusion, I used business survey data and provided a comprehensive analysis of the rationality of firms' expectations. To overcome the problem of qualitative business survey data, I employed a supplementary question about capacity utilization to construct measures of quantitative production expectation errors.

I show that firms' production expectations tend to be biased both at the macro and at the micro level. At the macro level, I find that the cross-sectional means of expectation errors are significantly negative at most points in time. This suggests that a large proportion of firms tend to have too optimistic and, therefore, biased expectations. At the firm level, I find in my conservative estimate that only 1 percent of all firms have distorted beliefs. However, using a broader measure of quantitative expectation errors reveals that more than 18 percent of firms have distorted beliefs. Nevertheless, this analysis alone is not enough to prove that a firm has rational expectations; it is necessary to also check whether firms efficiently exploit their information sets. I find that a considerable proportion of firms does not. Specifically, my conservative estimate of expectation errors demonstrates that about two-thirds of all the firms in my sample have rational expectations, i.e. their expectations are unbiased and they use all available information efficiently. For the broader definition, this number decreases substantially to 33.2 percent. Thus, there is evidence that heterogeneous firms form their expectations in heterogeneous ways, i.e. a large proportion of firms differ with respect to their beliefs and their ability to process information. In addition, I provide evidence that undistorted beliefs are correlated with firm-specific characteristics such as firm size or affiliation with a specific industrial sector. Smaller firms are especially prone to holding distorted beliefs. No systematic relationships are found between firm-specific characteristics and non-rationality.

The results provide evidence that a large fraction of firms does not form expectations in a rational way. Thus, further research into firms' expectation formation processes would seem a fruitful line of inquiry. This topic is also highly relevant for the transmission of macroeconomic shocks as shown in Cai (2010) and Bachmann and Sims (2011). Furthermore, allowing for some deviation from the rational expectation assumption usually incorporated in macroeconomic models has implications for the optimal conduct of economic policy (Ball et al., 2005).

Acknowledgements

I am grateful to Rüdiger Bachmann, Kai Carstensen, Christian Grimme and Michael Kleemann for discussions.

Bibliography

Anderson, O., R. K. Bauer, and E. Fels (1954): "On the accuracy of short-term entrepreneurial expectations," in *Proceedings of the Business and Economic Statistics Section*, pp. 124–147. American Statistical Association.

Bachmann, R., and C. Bayer (2011): "Uncertainty Business Cycles - Really?," Working Paper 16862, National Bureau of Economic Research.

Bachmann, R., and S. Elstner (2011): "Firms' Optimism and Pessimism: Evidence from the IFO Survey," *mimeo*.

Bachmann, R., and E. R. Sims (2011): "Confidence and the Transmission of Government Spending Shocks," Working Paper 17063, National Bureau of Economic Research.

Ball, L., N. G. Mankiw, and R. Reis (2005): "Monetary Policy for Inattentive Economies," *Journal of Monetary Economics*, 52(4), 703–725.

Becker, S., and K. Wohlrabe (2008): "Micro Data at the Ifo Institute for Economic Research - The 'Ifo Business Survey' Usage and Access," *Schmollers Jahrbuch*, 128, 307–319.

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

Bloom, N., M. Floetotto, and N. Jaimovich (2010): "Really Uncertain Business Cycles," *mimeo*.

Bovi, M. (2009): "Economic versus Psychological Forecasting. Evidence from Consumer Confidence Survey," *Journal of Economic Psychology*, 30(4), 563–574.

Brown, P., A. Clarke, J. C. Y. How, and K. Lim (2000): "The Accuracy of Management Dividend Forecasts in Australia," *Pacific-Basin Finance Journal*, 8, 309–331.

Brunnermeier, M. K., and J. Parker (2005): "Optimal Expectations," *American Economic Review*, 95(4), 1092–1118.

Cai, J. A. (2010): "Information Heterogeneity by Firm Size and Business Cycles," *mimeo*.

Christiano, L. J., M. Eichenbaum, and C. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), 1–45.

Davis, S. J., and J. C. Haltiwanger (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(3), 819–863.

de Leeuw, F., and M. J. McKelvey (1981): "Price expectations of business firms," *Brookings Papers on Economic Activity*, 1981(1), 299–314.

——— (1984): "Price Expectations of Business Firms: Bias in the Short and Long Run," *American Economic Review*, 74(1), 99–110.

Doms, M., and T. Dunne (1998): "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1(2), 409–429.

Firth, M., and A. Smith (1992): "The Accuracy of Profits Forecasts in Initial Public Offering Prospectuses," *Accounting and Business Research*, 22(87), 239–247.

Flaig, G., and K. F. Zimmermann (1983): "Misspecification and seasonal adjustment of qualitative panel data," *CIRET-Sudien*, 33, 63–95.

Gali, J. (2008): *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.

Heid, B., and M. Larch (2011): "Animal spirits in expectation formation: Evidence from firm-level survey data," *mimeo*.

König, H., M. Nerlove, and G. Oudiz (1981): "On the formation of price expectations: an analysis of business test data by log-linear probability models," *European Economic Review*, 92(5), 1521–1534.

Lucas, R. E. (1976): "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.

Lui, S., J. Mitchell, and M. Weale (2008): "Qualitative Business Surveys: Signal or Noise?," Discussion Paper 323, NIESR.

Mankiw, N. G., and R. Reis (2002): "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117(4), 1295–1328.

McDonald, C. (1973): "An Empirical Examination of the Reliability of Published Predictions of Future Earnings," *The Accounting Review*, 48(3), 502–510.

Mincer, J. A., and V. Zarnowitz (1969): "The Evaluation of Economic Forecasts," in *Economic Forecasts and Expectations*, ed. by J. A. Mincer, pp. 3–46, New York. Columbia University Press.

Müller, H. C. (2010): "Firms' Forecast Errors Regarding their own Future Key Figures: The Disappearance of the Overoptimism Bias," *mimeo*.

Muth, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29(3), 1–23.

Nerlove, M. (1983): "Expectations, Plans, and Realizations in Theory and Practice," *Econometrica*, 51(5), 1251–1279.

Reis, R. (2006): "Inattentive Producers," *Review of Economic Studies*, 73(3), 793–821.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50(3), 665–690.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach.," *American Economic Review*, 97(3), 586–606.

Souleles, N. S. (2004): "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 28(3), 39–72.

Tompkinson, P., and M. S. Common (1983): "Evidence on the rationality of expectations in the British manufacturing sector," *Applied Economics*, 15, 425–436.

Veldkamp, L., and J. Wolfers (2007): "Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement," *Journal of Monetary Economics*, 54(S), 37–55.

Weale, M., and M. H. Pesaran (2006): "Survey Expectations," in *Handbook of Economic Forecasting*, ed. by G. Granger, and A. Timmermann, pp. 715–776.

Zimmermann, K. F. (1986): "On Rationality of Business Expectations: A Micro Analysis of Qualitative Responses," *Empirical Economics*, 11, 23–40.

Zimmermann, K. F., and S. Kawasaki (1986): "Testing the rationality of price expectations for manufacturing firms," *Applied Economics*, 18, 1335–1347.

Appendix

A1 IFO Business Climate Survey (IFO-BCS)

Original German IFO-BCS Questions

Q 1 *"Sonderfragen: Die **Ausnutzung** unserer **Anlagen** zur Herstellung von XY (betriebsübliche Vollauslastung=100%) beträgt **gegenwärtig** bis zu ...%. "*

30	40	50	60	70	75	80	85	90	95	100	mehr als 100 % und zwar

Q 2 *"Erwartungen für die nächsten 3 Monate: **Beschäftigte** (nur inländische Betriebe) - Die Zahl der mit der Herstellung von XY beschäftigten Arbeitnehmer wird: zunehmen, etwa gleichbleiben, abnehmen."*

Q 3 *"Sonderfragen: Unter Berücksichtigung unseres gegenwärtigen Auftragsbestandes und des von uns in den nächsten 12 Monaten erwarteten Auftragseingangs halten wir unsere derzeitige **technische Kapazität** für XY für: mehr als ausreichend, ausreichend, nicht ausreichend. "*

Q 4 *"Erwartungen für die nächsten 3 Monate: Unsere inländische **Produktionstätigkeit** – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY wird voraussichtlich: steigen, etwa gleich bleiben, abnehmen."*

Q 5 *"Tendenzen im vorangegangenen Monat: Unsere inländische **Produktionstätigkeit** – ohne Berücksichtigung unterschiedlicher Monatslängen und saisonaler Schwankungen – bezüglich XY ist: gestiegen, etwa gleich geblieben, gesunken."*

Q 6 *"Unsere **Geschäftslage** für XY wird in konjunktureller Hinsicht: eher günstiger, etwa gleich bleiben, eher ungünstiger."*

Q 7 *"Aktuelle Situation: Wir beurteilen unsere **Geschäftslage** für XY als: gut, befriedigend, schlecht."*

A2 Mapping of $REALIZ_t$ into Quantitative Production Changes

Figure A1: Mapping of $REALIZ_t$ into Quantitative Production Changes

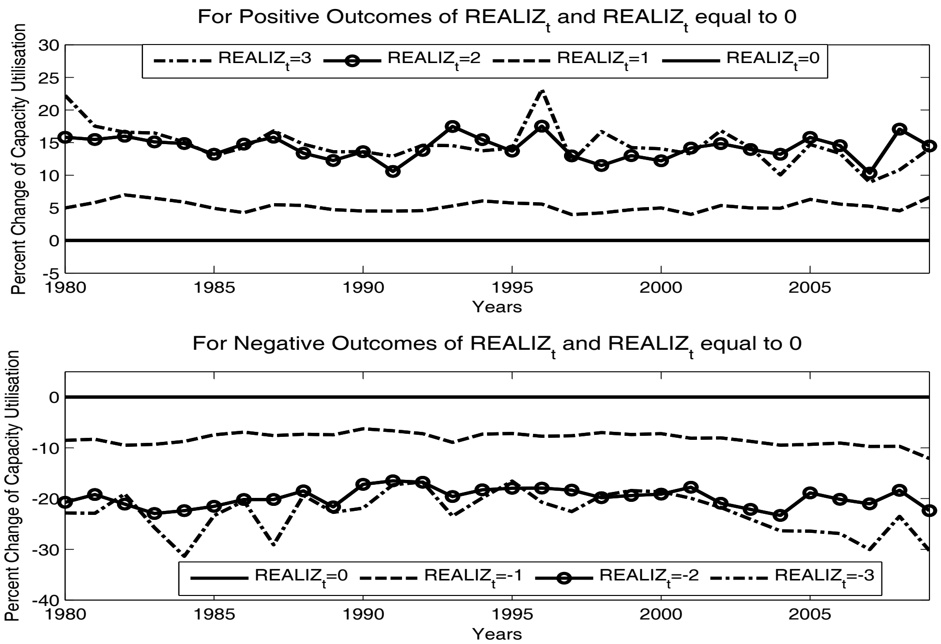

Notes: This figure shows for each possible value of $REALIZ_t$ the average $\Delta\log u_{i,t}$ of all firms with a given value for $REALIZ_t$. To improve readability, the quarterly values of $\Delta\log u_{i,t}$ have been averaged to obtain annual numbers. The upper panel shows the results for positive outcomes of $REALIZ_t$ and $REALIZ_t$ equal to zero. The lower panel displays the results for negative outcomes of $REALIZ_t$ and $REALIZ_t$ equal to zero. The original average $\Delta\log u_{i,t}$ of $REALIZ_t = 0$ has been subtracted from all time series shown in the figure.

A3 Robustness Check - Qualitative Expectation Errors (IFO-BCS)

Figure A2: Comparison of $Experror_{i,t}^{nar}$ and $Experror_{i,t}$ with Manufacturing Production

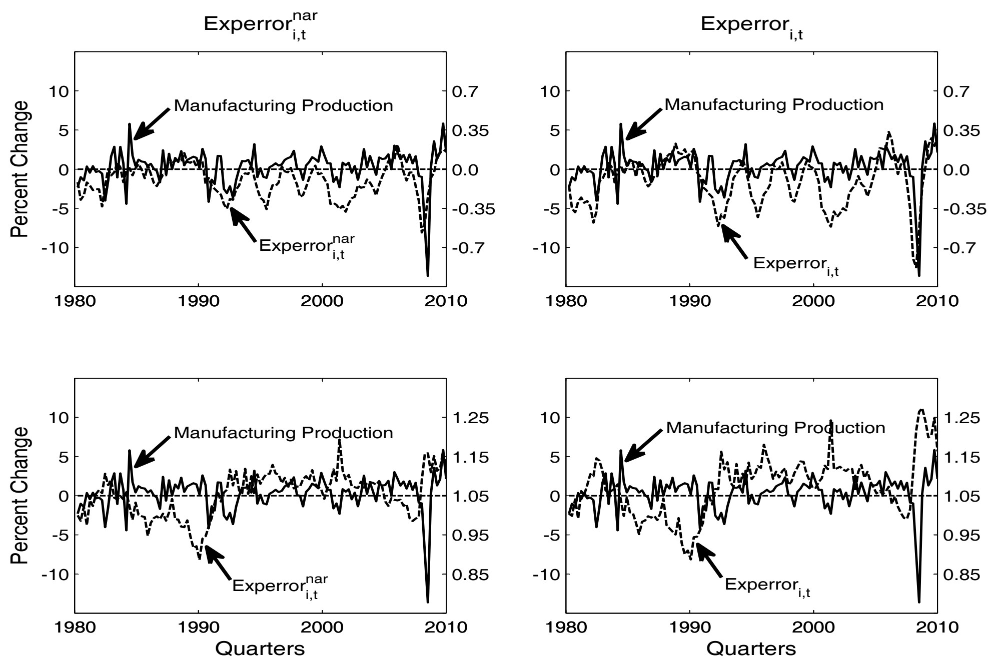

Notes: The upper panels show the quarterly cross-sectional means of $Experror_{i,t}^{nar}$ and $Experror_{i,t}$ (left scale) together with the quarter-on-quarter percent changes in manufacturing production (right scale). The lower panels show the quarterly cross-sectional standard deviations of $Experror_{i,t}^{nar}$ and $Experror_{i,t}$ (left scale) together with the quarter-on-quarter percent changes in manufacturing production (right scale). All time series are seasonally adjusted.

Table A1: Cyclical Properties of $Experror_{i,t}^{nar}$ and $Experror_{i,t}$

Variable in $t+j$	Manufacturing production in t -4	-3	-2	-1	0	1	2	3	4
Cross-sectional means of:	Leading Property of Variable					Lagging Property of Variable			
$Experror_{i,t+j}^{nar}$	0.05	0.18	0.44	0.62	0.50	0.34	0.20	0.01	-0.12
$Experror_{i,t+j}$	-0.02	0.12	0.40	0.63	0.60	0.44	0.24	0.05	-0.08
Cross-sectional standard deviations of:									
$Experror_{i,t+j}^{nar}$	0.22	0.24	0.05	-0.06	-0.20	-0.22	-0.25	-0.22	-0.20
$Experror_{i,t+j}$	0.26	0.23	0.09	-0.03	-0.22	-0.30	-0.32	-0.26	-0.24

Notes: see notes to Table 4.2.

Figure A3: Biasedness of $Experror_{i,t}^{nar}$ and $Experror_{i,t}$ on the Macro Level

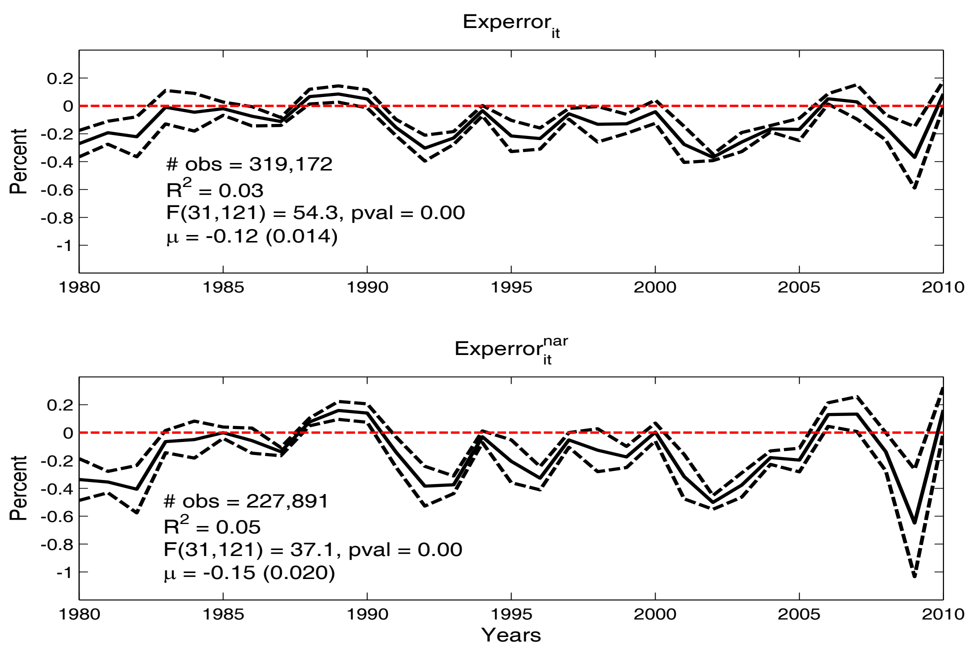

Notes: see notes to Figure 4.5. The upper panel shows the annual time effects of $Experror_{i,t}$. The lower panel does the same for $Experror_{i,t}^{nar}$.

Eidesstattliche Versicherung

Ich versichere hiermit eidesstattlich, dass ich die vorliegende Arbeit selbständig und ohne fremde Hilfe verfasst habe. Die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sowie mir gegebene Anregungen sind als solche kenntlich gemacht. Die Arbeit wurde bisher keiner anderen Prüfungsbehörde vorgelegt und auch noch nicht veröffentlicht.

München, 21. September 2011

Steffen Elstner

Curriculum Vitae

04/2008 – 09/2011	Ph.D. Student at the University of Munich,
08/2007 – 05/2008	Advanced Studies Program in International Economic Policy Research, Kiel Institute for World Economics
10/2002 – 09/2007	Diploma in Economics, University of Magdeburg
04/2002	Abitur, Werner-von-Siemens Gymnasium, Magdeburg
16.08.1982	Born in Magdeburg, Germany

München, 21. September 2011