THE BEHAVIOR OF THE FORECAST ERRORS: INSIGHTS FROM FIRM LEVEL SURVEY DATA

by

ALEXANDROS BOTSIS

A thesis submitted to the University of Birmingham for the degree of

DOCTOR OF PHILOSOPHY

Department of Economics

Birmingham Business School

College of Social Sciences

University of Birmingham

January 2021

Synopsis

Expectations are the quintessence of modern economic theory. Indeed, economic agents base their decisions on their forecasts of the future, which distinguishes the economic science from the natural sciences. Depending on how agents form their expectations, economic models predict vastly different outcomes, which has policy making implications. For instance, money supply cannot affect output when the full information rational expectations hypothesis holds. Therefore, there is a need for a better understanding of the properties of the agents' expectations. To improve our understanding of the agents' expectations, in this thesis, I focus on their forecast errors. I study an innovative combination of data comprising of two datasets. First, the survey data that records the expectations of the firms in Greece's Manufacturing sector. Second, I match the survey responses with the respondents' financial statements. Studying the forecast errors of the firms, I offer some novel insights into the full information rational expectations hypothesis. Namely, I show that only major forecast errors (the least accurate ones) show systematic patterns that lead to the rejection of the rational expectations. Minor forecast errors show no systematic patterns. In order to arrive at this conclusion develop a quantification model that allows me to compute quantified forecast errors on sales growth using the survey-based expectations and the realized sales growth of the financial statements. This quantification is pivotal in distinguishing between major and minor forecast errors is impossible. The contribution of my quantification model is that it delivers quantitative annual forecasts at the firm-level. Second, I adapt and implement a recent threshold estimator that endogenously estimates where the behavior of the forecast errors changes. Finally, in the last part, I provide a new test for the full information rational expectations which measures the extent of the information inefficiencies in survey data. This test is broadly applicable to all survey-based forecast errors.

To my partner Luke,

to my family,

to my friends.

Acknowledgements

For this work I am most grateful to my partner Luke, to my family and to my friends for their unconditional love and support. Especially my friends Bárbara and Rosalia who welcomed me when I first moved to the United Kingdom to start my PhD journey.

I want to express my deepest appreciation to my supervisors Christoph Görtz and Kausik Mitra for their patience and support through all the stages of this work. They have always been available, they have taught me the art of perseverance, the attention to detail, and they have had an immense impact on my development. Words cannot express how thankful I am. Next, I would like to thank Prof. Plutarchos Sakellaris without whom I could never acquire the data that I have and who has been most helpful with his feedback and suggestions. My work for this thesis and the collaboration with Christoph Görtz and Plutarchos Sakellaris have resulted in a working paper, Botsis et al. (2020). Some passages have been quoted verbatim from this working paper.

I am also particularly grateful to Prof. Alexandros Sarris who guided me through my early steps as a researcher. He taught me that economists should pay great attention to detail — perhaps even more than physicians — because one miscalculation in economic policy can irreversibly impact millions of lives.

Additionally, I owe a big 'Thank you!' to Prof. Patricia Chelley-Steeley for always encouraging me to prioritize my PhD work and for intriguing me to explore the field of Behavioral Finance.

Special thanks to all the internal seminar participants for their insightful comments. Finally, I am grateful to the Department of Economics and to the ESRC for their financial assistance, without which I would not be able to finish this journey.

This has been a fun albeit rugged journey.

Contents

List of Tables

Chapter 1

Introduction

The expectations of the economic agents is the quintessence of modern economic theory. By using this theory to inform their actions, policy makers also rely on assumptions about the expectations of the agents their policies affect. This is because all agents make forward-looking decisions. Households optimally chose consumption and savings based on their expected future income. Firms decide how much to invest and how many people to hire based on their forecasts for their future demand and revenue. They also decide about their prices using their forecasts of their future marginal costs. Investors estimate the fair price of equity shares by forecasting the future cash flows these shares will generate.

Evidently, expectations are important in all aspects of economic theory. But, now comes the tricky part. Different mechanisms of expectation formation can lead to very different outcomes of economic policy. For example, the full information rational expectations of Lucas (1972) results into money neutrality, that is monetary policy cannot affect what firms chose to produce. Intuitively, the full information rational expectations hypothesis assumes that agents make forecasts about the future having *correctly* used *all* the available information at the time of the forecast. If this is true, then the errors

of these forecasts will be completely random and will not display any systematic behavior. Contrary to the rational expectations, when firms delay updating their information — like those in Mankiw and Reis (2002) — also delay the adjustment of prices and increase their output in response to an expansionary monetary policy. This outcome of the monetary policy is completely different to what the rational expectations would imply.

In addition to different outcomes of policy making, the deviation from the full information rational expectations hypothesis can also lead to different optimal behaviors. Intuitively, if agents show persistent bias in their forecasts, then their optimal choices will be inferior. For example, if households or firms over-estimate their future income they are also likely to over-estimate their ability to pay back their debts and will ultimately be forced into default. In another example, if investors over-estimate the future dividend streams of a stock they will be over-valuing that stock and they will most likely end up with losses. Under full-information rational expectations, agents are assumed not to demonstrate any persistent biases, so any unexpected losses or gains are purely random and on average independent to their forecasting mechanism.

What is more, under the assumption of full information rational expectations (as well as in standard econometric theory), agents should use all available data when making forecasts. Mitra (2005) shows that this is not the case when agents are not fully informed about how the variable they want to forecast evolves (the data generating process). When agents are not aware of the data generating process, they use a simple rule of thumb to make a forecast, they use the simple arithmetic mean of the lagged values. In case of persistent autocorrelated variables, agents will optimally forecast using naive expectations — i.e. next period's expected value equals the last observable value. In fact, the mean squared forecast error resulting from naive forecasts is lower than the mean

squared forecast error when agents use the arithmetic mean of all past observable values. All the aforementioned examples clearly illustrate the importance of understanding the properties of the expectations.

A robust finding in the literature is that the expectations do not follow the full information rational expectation paradigm. Lovell (1986) takes a decisive step in challenging the primacy of the full information rational expectations. He provides evidence drawn from existing research findings that the full information rational expectations hypothesis is overwhelmingly rejected in the empirical research. More recently, Gennaioli et al. (2016) study firm-level expectations, and in particular the Chief Financial Officers' (CFOs) forecasts about their own firm's earnings, and find evidence that the rational expectations hypothesis is violated. In fact, CFOs make forecast errors that are predictable by the information that was available to them when they were making their forecasts (ex-ante). Tanaka et al. (2020) also find evidence of serial correlation in the firms' forecasts about GDP growth. This is again a violation of the rational expectations hypothesis. Both the predictability and the auto-correlation indicate a systematic and non-random behavior of the forecast errors, which contradicts the full information rational expectations hypothesis.

Further empirical evidence on the violation of the full information rational expectation hypothesis can be found abundantly. Coibion and Gorodnichenko (2012) find that both consumers and professional forecasters deviate from the full information rational expectations. Coibion and Gorodnichenko (2015) also find that forecast errors on inflation made by professional forecasters, non-financial businesses, banks and consumers are all predictable by their past forecast revisions. Bachmann and Elstner (2015) study firm-level survey data that shows that a significant portion of the firms make biased forecasts that are persistently over-optimistic or over-pessimistic. Massenot and Pettinicchi

(2018) also find that firm-level forecast errors are predictable by past realizations.

How can we explain the deviation from the full information rational expectations hypothesis? To answer this question, economists have focused on the information the agents use to make their forecasts and on how the agents use this information. One can identify two strands of theoretical models that explain these stylized deviations from the ideal. The first strand is the behavioral biases that occur when agents use heuristics in order to make a forecast. Bordalo et al. (2018) show that when agents use diagnostic expectations to forecast the credit spreads, their forecast errors become predictable by past realizations. If they suffer from diagnostic expectations, when agents receive new information that increases the probability of possible future outcomes, they will overestimate the increase of that probability. In other words, in light of new information a future event that becomes more likely, agents tend to mis-perceive it as even more likely, they 'overweight' it. Gabaix (2014) also shows how cognitive limitations can result into biased assessments regarding the relationship among variables. In Gabaix (2014), agents' cognitive limitations allows them to include in their information set only a subset of all the possible variables when they form a forecast. In that way, their forecasts are not fully informed. Finally, Mitra (2005) shows that when agents are not fully aware of the functional form of the data generating process of the variable they attempt to forecast, it is optimal to use less data than what full information rational expectations would imply.

The second strand of theoretical models that can track the origins of the deviation from the full information rational expectations hypothesis is rational inattention. In rational inattention the presence of information costs either prevent the agents from updating their information more frequently as in Mankiw and Reis (2002) or limit the precision of that information as in Sims (2003). In rational inattention, agents face

costs in improving the precision of the information they receive and rationally chose to sacrifice precision in order to economize on the costs. The difference between the behavioral models and the rational inattention lies in the nature of the imperfection. In the behavioral models, the cognitive limitations of agents forces them to use simplified versions of their information set, such as heuristics, or ignoring a part of it. In the rational inattention, there are costs in acquiring and processing information that make inattention the optimal, the rational, choice. Most importantly, most rational inattention models lay the micro-foundations that result in limited attention, whereas the majority of the behavioral models assume ad-hoc the agent behavior. One exception in the Behavioural literature is Gabaix (2014)'s sparsity matrix which is micro-founded on a signal-extraction mechanism resembling with rational inattention to a certain extent.

The rational inattention models that Sims (2003) introduced are based on the information theory (the details of information theory can be found in, for example, Cover and Thomas (2006)). In these models, agents use noisy signals to reduce uncertainty about the variables they intend to forecast. In the jargon of the field, the reduction in uncertainty is called '(Shannon's) mutual information' — see Sims (2003) and Mackowiak et al. (2018). The costs the agents face in collecting information is a positive function of the reduction in uncertainty. That is, better and more accurate signals resulting in more reduction in uncertainty are also more costly. Therefore, agents optimally choose to buy the best possible signal given the constraint of its cost. Clearly, by using less or lower quality information, their forecasts deviate for the full information rational expectations and their forecast errors start displaying systematic patterns. With full information rational expectations, agents accurately observe and process all the available information at the time of the forecast, and their forecast errors do not display any systematic pattern.

It has become evident that the forecast errors the firms, and any agent, make are key in showing us whether the forecasts were rational and fully informed. By studying the forecast errors the firms make, I will provide some new insights on the properties of the expectations of the firms. I study a novel combination of data comprising of firm-level survey responses matched with the respondents' financial statements. The survey data is of firm-month frequency and the financial statements are of firm-year frequency. I focus on the Manufacturing sector of Greece to maintain comparability with existing research and because the Manufacturing survey is the largest dataset available. Financial statements are available from 1998 until 2015, which also determines my time horizon for the initial chapters; Chapters 2, 3 and 4. In Chapter 5, I will entirely focus on the survey, which allows me to extend my horizon to cover a period from 1998 until 2019.

I study the surveys because they constitute an important tool that allows us to observe the expectations of the agents and their forecast errors. The economic agents that usually respond to surveys recording expectations are firms, households and professional economists. Surveys directly record the agents' expectations about the future development of aggregate or micro-level variables. The survey that I will focus on asks firms about the future development of their own variables. Even though forecasts about firm-specific variables are also collected from professional forecasters, asking firms about their own variables has one important advantage. The forecasts of the firms are conditional on the information the firm has, as opposed to the publicly available information on which professional forecasters rely. After all, it is the firms themselves that make hiring and investment decisions that drive the aggregate economic activity. The anonymity of the survey that I study guarantees that firms do not have any incentive of strategic responses.

The main reason why firm-level survey expectations are collected is that they also

offer a way to assess the overall business climate in a timely and frequent manner. As such, they are collected every month and give us a quick look into how optimistic or pessimistic firms are about their near future. Being collected every month comes, however, at a cost; firm's responses are qualitative. In order for the surveys to collect responses every month, they have to be less demanding on what they ask firms to do. It is easier for the respondents to indicate each month what is the direction of change of their future sales, for example, than to indicate the exact growth rate of that change.

Given that I want to study the firm-level forecast errors in depth, the qualitative nature of the survey responses is a limitation. As I will demonstrate later, the largest forecast errors (in absolute value) behave entirely differently from the smaller ones with respect to the violation of the rational expectations hypothesis. Clearly, in order to arrive at this finding one has to analyze quantitative forecast errors that allow for a distinction between the least and the most accurate ones. Qualitative data are bounded by construction and cannot distinguish between the largest and the smallest forecast errors. For example, qualitative forecast errors cannot distinguish between firms expecting a positive growth of 10% and a 1% followed by a realized contraction of -1%. Both these forecast errors have the same qualitative value that only indicates that the two forecasts completely miss-predicted the direction of change. It is therefore essential to measure forecast errors in a continuous scale rather than in a qualitative one.

The importance of quantitative forecasts and forecast errors in examining the full information rational expectations is also stressed by Pesaran and Weale (2006). Pesaran and Weale (2006) suggest that in order to examine the full information rational expectations hypothesis in survey data, researchers need to obtain quantitative forecasts first. To examine these properties in the forecast errors in my data, I also quantify the survey based forecasts. I can achieve this by matching the survey data with the financial

statements of the surveyed firms, their balance sheets and their income statements. This combination of data is quite novel and allows me to assess the quality of the information in the survey responses in addition to quantifying the qualitative survey forecasts on sales growth. After having computed the quantitative firm-level forecasts I use the realized sales growth from the financial statements to compute the quantitative forecast errors that I study.

With the estimated quantitative forecast errors at my disposal, I investigate whether they display any systematic behavior, which indicates that the full information rational expectations hypothesis does not hold. I find that the forecast errors that the firms make are both predictable and autocorrelated. This is a standard finding in the literature — for example, Bachmann and Elstner (2015), Gennaioli et al. (2016) and Tanaka et al. (2020). I contribute by showing that not all of the forecast errors show this systematic behavior. In fact, I provide robust evidence that only major forecast errors display this behavior that violates the rational expectation hypothesis. I define major forecast errors to be the least accurate ones (the largest ones in absolute value), while minor forecast errors are the relatively more accurate ones. Clearly, the availability of quantified forecast errors is essential for this exercise, because the qualitative survey data does not allow for the distinction between the minor and the major forecast errors.

Before obtaining my results, I present the data in detail and I examine its quality in Chapter 2. This is a necessary step, because I am using a novel dataset. In particular, I firstly describe how the data is collected and recorded, and I show that it is a fairly representative sample. Second, I document that the responses are well informed. I show that on average there are no contradictions among the responses that each firm gives in each month. I also give evidence that the survey responses are strongly correlated with their quantitative counterparts in the survey data. With these findings I affirm the

quality of the information in my data, before using it for the analysis that follows.

Next, in Chapter 3 I develop my quantification model with which I obtain quantitative forecast errors on sales growth. I also explain in detail that my quantification model contributes to the literature by providing quantified forecasts (and forecast errors) of firm-year frequency whereas all existing methodologies deliver time-series. Intuitively, the quantification model I develop consists of two non-linear equations. The first equation shows how the survey based forecasts can be aggregated by a non-linear function that gives us a quantitative forecast for each firm in each year. To implement this model and compute the quantitative forecasts I need to obtain estimates of the parameters of this first equation. The second equation, I show that I can use it to estimate these parameters. In the process of obtaining the second equation, I also show how to control for the unobserved firm-specific characteristics. As a result, with the estimated parameters and the survey data, I then compute the quantified forecasts. The difference of these quantified forecasts from the realizations in the financial statements gives me the quantified forecast errors.

To validate my estimates of the quantitative forecasts and forecast errors, I additionally develop a series of exercises in Chapter 3. Due to the unavailability of observable quantitative forecasts, I cannot directly examine how close the estimated values are to the actual ones. The unavailability of quantitative forecasts is the reason why I need the quantification in the first place. This is standard in the literature of quantified expectations, which is why there are no pre-existing ways to assess the quantified forecasts. However, I design a series of exercises to validate as much as possible the quantified forecasts that I obtain. First, I show that my quantified forecasts are fully consistent in terms of the overall annual direction of change with the corresponding qualitative survey-based expectations. Additionally, my methodology also substantially outperforms models that

could be potential alternatives for obtaining quantified predictions. Second, I perform a Monte Carlo exercise that provides a benchmark based on simulated data. I find that the forecast errors imputed from my methodology in the simulated data are highly accurate when compared with the observable — in the artificial data — forecast errors. These two exercises demonstrate that I can rely on the quantified forecast errors to perform my analysis.

Having the quantified forecast errors on sales growth, I examine their behavior in depth in Chapter 4. I firstly document that the quantified forecast errors violate the full information rational expectations hypothesis by displaying a systematic behavior which is well documented in the literature. I find that forecast errors are predictable by past realizations of sales growth and that they are autocorrelated. I contribute to the literature by providing robust evidence that only the major forecast errors (the least accurate ones) display this behavior. In contrast, the minor forecast errors are neither predictable nor autocorrelated. I arrive at this finding by using an estimator that identifies endogenously where that shift in the behavior occurs; the Dynamic Panel Threshold Estimator of Seo and Shin (2016). To implement this estimator I firstly adapt it to better accommodate the missing data of my unbalanced panel. The estimates I obtain document that the forecast errors that show this systematic behavior are those that lie either at the lower or at the upper 26% of their empirical pool distribution. Even though this might sound trivial, I demonstrate that the minor forecast errors are still quite far from being zero. In other words, minor forecast errors are not small enough to make their systematic behavior disappear 'under the radar', they simply do not show any systematic behavior.

To rationalize my empirical findings, in the last part of Chapter 4, I develop a simple and analytical model of rational inattention. At the time of the forecast, firms have to

rely on noisy signals to make their forecast on future sales growth. Knowing that they use signals, firms limit their attention to them. Firms can acquire more accurate signals (and pay full attention to them) at a cost. I show that firms optimally choose to pay less than full attention because being more careful bears higher costs — this is the very essence of rational inattention. An increase in the cost of attention pushes firms to make major forecast errors. Without a positive attention cost, firms accurately observe the lagged value, do not make major forecast errors and form rational and fully informed forecasts. Therefore, only in the presence of major forecast errors that originate from rational inattention, agents violate the full information rational expectations hypothesis. This is consistent with the empirical findings I obtain in the first parts of Chapter 4.

Finally, in Chapter 5, I study the firm-level survey-based forecast errors on production. In this chapter I do not use the financial statements data, which enables me to study a larger sample covering a horizon from 1998 to 2019 with firm-monthly survey data. I devise a decomposition of the variance of these forecast errors into a forecastable and an unforecastable component. For the unforecastable component, I argue that it is a measure of micro-uncertainty. The contribution of this new measure is that it is of firm-month frequency and can be obtained from survey-based forecast errors without additional data requirements. To validate my proposed measure, I engage in two exercises. First, I show that it behaves in the same manner as other existing and widely used uncertainty measures. Second, I show that larger micro-uncertainty increases the probability of a subsequent forecast errors.

The most important contribution in Chapter 5 comes, however, from the forecastable component of the forecast error variance. Using the forecastable component I develop a novel test for the full information rational expectations hypothesis. Importantly, the test I propose determines whether the rejection comes from inefficient use of information, and

it measures the magnitude of this inefficiency. Additionally, it is generally applicable to survey-based forecast errors, qualitative and quantitative, without additional data requirements. The test stems from the variance of the forecast errors, which eliminates the need for quantified forecast errors to test the full information rational expectations hypothesis.

Overall, I study an innovative combination of data comprising of two datasets. First, the survey data that directly records the expectations of the firms. Second, I match the survey responses with the respondents' financial statements. Studying the forecast errors of the firms, I offer some novel insights into the full information rational expectations hypothesis. Namely, I show that only major forecast errors show systematic patterns that lead to the rejection of the rational expectations. Minor forecast errors show no systematic behavior. In order to arrive at this conclusion I firstly develop a quantification model that allows me to compute quantified forecast errors on sales growth using the survey-based expectations and the realized sales growth of the financial statements. Without this quantification, the distinction between major and minor forecast errors is impossible. The main contribution of my quantification model is that it delivers quantitative annual forecasts at the firm-level. Second, I adapt and implement a recent threshold estimator that endogenously estimates where the behavior of the forecast errors changes. Finally, in the last part of my thesis, I provide a new test for the full information rational expectations which measures the extent of the information inefficiencies. This test is generally applicable to all survey-based forecast errors without any additional data requirements.

Chapter 2

Data Description and its Quality of Information

Survey data has been the focus of economic research since decades. In addition to their use for now-casting by policy makers, economists have used survey data to study the expectation formation. Some of the earliest attempts were made by Nerlove (1983), De Leeuw and McKelvey (1984) and Pesaran (1985). They all find biases in the expectations that constitute a violation of the rational expectations hypothesis. The interest on survey data has been recently revived by Bachmann et al. (2013) who show that uncertainty in Germany is negatively associated with the economic activity. Bachmann and Elstner (2015) also find evidence of persistent biases in firms' expectations for the German economy. Coibion et al. (2018) conduct a survey in New Zealand business sector to study the expectations of firms about inflation and find evidence of violations of the rational expectations.

In addition to studying the expectations formation mechanism, survey data has also been a fruitful research tool in numerous ways. Enders et al. (2019a) show that firms' expectations about their future output affects their decisions about production and prices

at the time of the forecast. Enders et al. (2019b) use German data from the IFO Business Survey to study the effect of monetary policy announcements to the expectations of the firms. Bloom et al. (2019) use survey responses to study the effects of the Brexit vote on the UK economy and documented a long-term increase in uncertainty and a decline in investment and productivity. Bachmann and Zorn (2020) based on the Ifo survey study the effect of aggregate demand and technology shocks on the investment dynamics.

My data is a novel combination of two databases. The first one is survey responses of Greek firms, administered every month by 'IOBE' (Foundation for Economic and Industrial Research). The survey responses are used to construct the business climate index in compliance with the guidelines of Directorate-General for Economic and Financial Affairs (DGECFIN (2016)). Similar surveys are administered by the Confederation of British Industry for the United Kingdom and by the IFO Institute for Germany. IOBE surveys four 'broad-sectors': manufacturing; construction; retail trade; and services sector. Finally, all survey data concern the recent developments (retrospective or realization variables) and the expectations of firm-level variables. Surveys do not collect forecasts about aggregate macroeconomic or sector-level conditions. For my thesis, I focus on the manufacturing sector to maintain comparability with existing literature findings. More importantly, manufacturing is the largest of these broad sectors as it includes 38% of survey observations and 36% of observations in the financial statements data.

The second part of my data includes the balance sheet and income statements. I obtained it from ICAP S.A., a private consultancy firm, who obtains and digitizes them from official publicly available sources. ICAP S.A. also provides the same data to Bureau van Dijk for Orbis. The financial statements are compiled by certified auditors (chartered accountants) and are used for reporting to tax authorities, to investors and all interested parties. They are only available in digital form from 1998 to 2015, therefore this will be

the sample period of my analysis in this Chapter and in Chapters 3 and 4.[1]

The reason why I chose Greece is, most importantly, because of data availability: under a strict confidentiality agreement, I was given access to the de-unonymized survey data, and that is how I was able to combine it with the financial statements. That is to say, Greece was the only economy – to my knowledge – that this combination of data was available. Lui et al. (2011) used a similar approach matching business climate survey data for the UK with the quantitative data for the surveyed firms collected by the Office of National Statistics. My data-base, though, comprises a much larger sample in terms of number of firms and years. In addition, I use the officially published and audited financial statements, whereas the quantitative data from the Office of National Statistics are survey based, and might be vulnerable to higher measurement error and misreporting risk. Tanaka et al. (2020) also combine annual survey responses with the financial statements of the respondents for Japanese manufacturing firms. My main advantage to Tanaka et al. (2020) is that I obtain monthly frequency forecasts about firm's own variables (e.g. production and sales) and not for aggregate macro-economic variables as in Tanaka et al. (2020). Enders et al. (2019a) combine monthly firm-level survey data from Germany's Ifo Institute with the firms' annual balance sheets. My data is different by focusing on a smaller economy like Greece.

The combination of data that I have is what allows me to quantify the qualitative survey responses in Chapter 3. In turn, the availability of quantitative forecasts is essential in studying in depth the continuous annual forecasts in Chapter 4. Qualitative survey responses are by construction bounded and cannot allow for distinguishing between major and minor forecast errors. As I demonstrate in Chapter 4, there are paramount differences between the behavior of major forecast errors and that of minor.

[1]In Chapter 5 I extend the horizon as I am not constrained by the availability of the financial statements. There, I provide further evidence about the quality of the information in the extended sample horizon.

However, before working with this novel dataset, it is important to describe its details and assess its information content. I do so in this chapter. First, I describe the details of the survey data in Section 2.1 and some general sample characteristics. Second, in Section 2.2, I examine the representativeness of the data and its quality. Overall, I show that the data is fairly representative of the manufacturing sector in Greece and that responses have a rich information content. In particular I document that on average survey responses are consistent with each other and not contradictory. Also, I document that the survey responses are consistent with the financial statements of the firms. The latter is important as it demonstrates not only the validity of the responses but also the fact that the respondents have complete knowledge of the activities of the firms. This is because the financial statements of the firms are finalized during the following year well after each firm responds to the surveys. Overall, with these exercises I add credibility to my findings of the following chapters.

2.1 Survey Data Details

The firm-level survey responses are collected every month by 'IOBE' (Foundation for Economic and Industrial Research), from a sample of a larger firm directory that covers more than 1/2 of the output of the total economy and includes a large fraction of the total firm population. The sample is chosen initially to represent the distribution of firm sizes, in terms of gross sales, in each 2-digit sector.[2] Then, every 4-5 years it is replenished by removing those who never reply and those who have stopped replying and replacing them with new ones following the same sampling principles. The firms that have been responsive, though, continue to be sampled. According to IOBE researchers,

[2]Because the firm directory might be more biased towards larger firms, random sampling from it would not produce a sample representative of the economy.

the response rate is $< 20\%$. This is however the case with almost all surveys that aim at identifying business (and consumer) trends as contemporaneously as possible.

IOBE mails all the surveys between the 22th and the 25th of each month – surveys refer to the following month. Responses that arrive well past the month they initially refer to are dropped as it is unclear which month the firms have responded for. Much less than 10% of the sample responds electronically via e-mail. Finally, IOBE researchers ask firm managers or a person who has complete knowledge of the entire activity of the surveyed firm to respond to the survey. As soon as firms respond, they mail it back to IOBE in a prepaid (by IOBE) postal envelope. In Section 2.2 that follows, I provide evidence that, on average, the respondents have complete knowledge of the firm operations.

Although the surveys are conducted monthly, there is one exemption. During August the majority of the firms are closed as the majority of the managers and the employees take their annual leave, so surveys are not sent out at all. Additionally, IOBE uses imputation methods to produce data for August and for some monthly non-responses. Lui et al. (2011) report that for the UK business climate survey, the CBI (who administer the survey) also implements imputation techniques for missing data; they document that, on average, this does not jeopardize the credibility of the survey data. I also provide evidence on the credibility of my data in Section 2.2 that follows. In all parts of my analysis I ignore the firm-month observations that are imputed. In Appendix A, I document all the details of the cleaning steps I take.

2.2 The Quality of Survey Responses

I match firms' financial statements data with the corresponding survey responses using the firm's unique tax identifier. As I describe in Appendix A, there are 1,093 firms in

the cleaned survey data. I matched 73.1% of these firms or 76.7% of the firm-month observations with their financial statements. In Appendix B I also document the detailed steps I take to clean the financial statement data. The final sample for which I have both survey responses matched with their financial statements comprises of 799 firms in the manufacturing sector with 25,764 monthly responses from the survey. In this section I document that my final matched sample is representative for the manufacturing sector and I also evaluate the quality of the survey responses.

Given that my data is an unbalanced panel, there will be some selection bias. Controlling for selection bias is crucial in order to be able to make inference about the whole population of firms and the whole economy based on the sample. Assuming that the selection bias stems from the characteristics of each firm and are time invariant, then estimators that control for the unobserved firm heterogeneity (panel fixed effects) also eliminate the selection bias (see Van Beveren (2012)). In addition to using estimators that control for the panel fixed effects in all my estimates, I provide in this chapter information about the characteristics of the sample and about its representativeness.

In Table 2.1 I provide the key characteristics of my sample. My sample comprises of firms of various sizes; very small firms as well as large firms with more than 4,000 employees and annual sales turnover of over six billion Euros. On average, firms respond in six out of the 11 months in which surveys are sent out in each year.

2.2.1 Representativeness

I study the representativeness of my sample three ways. First, I examine the business sentiment index. I report the time-series correlation of 95% between the business sentiment index calculated from my sample and the official IOBE index. The monthly

Table 2.1: Sample Characteristics.

	Min.	Max.	Mean	Median	St. Dev.
Firm-Year Characteristics					
# of Employees	1	3,811	162	75	278
Real Sales (in thousands, 2005 Euros)	6	6,710,000	29,100	7,202	179,000
Survey Responses per Annum	3	11	6	6	3
Firm Level Characteristics					
Age at First Appearance in Sample	0	110	25	24	17
Time-Series Length in Sample (Years)	1	18	5	4	4

sentiment index for the manufacturing sector is computed as

$$\frac{QS_{im} + QS^e_{im} - INV_{im}}{3}.$$

The survey variables used to compute the business sentiment index are:

INV_{im}, question E.1: *'The level of your final goods inventories is...'* with the possible responses being above/at/below normal levels and coded as $+1/0/-1$, respectively.

QS^e_{im}, question D.1: *'During the next 3 months you expect that your total production will...'*.

QS_{im}, question A.1: *'During the 3 previous months, your total production did...'*.

For the last two questions, QS^e_{im} and Q_{im} , the possible responses are rise/no change/fall, coded as $+1/0/-1$, respectively.

Second, I examine the manufacturing real output growth. I report a 64% time-series correlation between: (i) the average real growth rate of output of my matched sample of manufacturing firms and (ii) of the output growth of the manufacturing sector from the national accounts. That is, my matched sample roughly captures 2/3 of the

Manufacturing output growth.[3] Because my data is an unbalanced panel, a few very large firms entering and exiting the sample creating huge spikes and troughs in the annual total sum of the sample output. Therefore, I cannot use this to assess the representativeness and I have to use the average of the growth rates.

Third, to further examine the representativeness of my final sample I study the share of each 2-digit sector in the total manufacturing sector sales. I compare their contributions based on my sample with the ones from the official Eurostat data. Table 2.2 exemplifies these statistics for two years — 2009 and 2012 — and I observe that most of the shares based on my dataset are close to the ones reported by Eurostat with few exceptions of over- and under- representativeness.

2.2.2 Quality of Survey Responses

In this section I aim at examining the information content of my data. First, I will determine whether the surveyed firms provide consistent or conflicting answers in each month m. In the spirit of Coibion et al. (2015), I follow a regression-based approach in order to document the consistency of the survey responses. I refer to this consistency as the 'internal consistency'. Second, apart from the internal consistency of the responses, I also have the unique opportunity to verify the survey responses using external sources, the financial statements. To my knowledge, no such validation has been conducted so far in the literature concerning firm surveys focusing on firms' own variables. My data combination is what gives me this opportunity. I will call this consistency 'external'.

I begin by examining the internal survey consistency of the manufacturing survey. I will do two exercises to establish consistency that will jointly cover around 2/3 of the

[3]Output from the financial statements is the sum of sales plus the contemporaneous first difference of the final goods inventories. I deflated the firm-year output of the financial statements using the ratio of the nominal over real (chain linked volumes) gross value added at the NACE 2-digit level. I use the simple arithmetic mean of the firm-year observations to obtain the average growth rate of my sample. The manufacturing growth rate of real output from Eurostat for Greece is from Table nama_10_a64.

Table 2.2: Share of NACE 2-digit industry sales in the total manufacturing sales in years 2009 and 2012.

	2009		2012	
NACE Code	Sample Data	Eurostat Data	Sample Data	Eurostat Data
10	13.35%	20.23%	16.01%	19.74%
11	10.11%	3.94%	6.03%	2.98%
12	2.67%	1.01%	1.60%	0.74%
13	1.99%	1.93%	1.94%	1.26%
14	0.58%	3.16%	0.29%	1.84%
15	0.74%	0.50%	0.20%	0.21%
16	0.95%	1.50%	0.06%	0.82%
17	1.63%	2.02%	0.89%	1.76%
18	0.84%	1.63%	0.30%	1.06%
19	19.71%	21.77%	45.19%	36.54%
20	5.58%	4.44%	4.22%	3.48%
21	10.70%	2.63%	6.08%	1.80%
22	2.42%	3.24%	2.17%	3.04%
23	6.95%	5.90%	1.92%	2.78%
24	7.23%	7.49%	1.19%	8.62%
25	6.93%	7.60%	7.29%	5.31%
26	2.60%	0.68%	1.00%	0.68%
27	0.68%	2.49%	0.57%	2.68%
28	2.14%	2.39%	1.57%	1.68%
29	0.53%	0.51%	0.25%	0.27%
30	0.39%	1.12%	0.78%	0.36%
31	0.62%	1.76%	0.21%	0.91%
32	0.38%	0.96%	0.23%	0.58%
33	0.28%	1.11%	0.00%	0.87%

For my sample, total manufacturing sales is the sum of sales of all firms in a particular year. The shares reported show the sum of sales in a 2-digit sector over total manufacturing sales in my sample for a particular year. The shares in the 'Eurostat' columns are the corresponding ratios based on Eurostat sales data based on Table sbs_sc_sca_r2 for Greece.

survey questions.

First, based on economic intuition, if a firm expects higher sufficiency of production capacity, it is more likely to: report higher than normal inventory level; expect a drop in the demand or sales; expect it will have to decrease employment; have lower capacity utilization that would allow it to increase production if need be. To confirm that my economic intuition holds in my data I estimate the following linear equation:

$$D3_{im} = \beta_0 + \boldsymbol{\beta}\Big[INV_{im}, XS^e_{im}, L^e_{im}, U_{im}\Big]' + \psi_i + \psi_y + \eta_{im}, \tag{2.1}$$

where the vector $\boldsymbol{\beta} = [\beta_1, \beta_2, \beta_3, \beta_4]$, ψ_i and ψ_y control for firm and year fixed effects respectively, and η_{im} is the idiosyncratic error. The variables $D3_{im}$, INV_{im}, XS^e_{im}, L^e_{im}, and U_{im} denote current production capacity, inventory level, sales, the number of employees, and capital utilization of firm i in month m and are derived from survey questions. The precise questions are as follows:

INV_{im}, question E.1: *'The level of your final goods inventories is: above normal/normal/below normal'.*

$D3_{im}$, question E.2: *'Given the outstanding orders you have at the moment and the possible evolution of demand during the next months, the current production capacity is more than sufficient/sufficient/insufficient'.*

XS^e_{im}, question D.2: *During the next 3 months, you expect your total sales to increase/remain unchanged/decreased.*

L^e_{im}, question D.3: *'During the next 3 months, you expect your number of employees to increase/remain unchanged/decrease'.*

In these questions, a numerical value -1 refers to reduction or lower than normal level or insufficient production capacity as appropriate; $+1$ refers to an increase or

higher than normal level or more than sufficient as appropriate; and 0 refers to no change or normal level or sufficient capacity as appropriate.

U_{im}, question E.3: *'During the ongoing period, what is your percentage (%) utilization of your production capacity?* Firms respond to this question with a quantitative answer.

I estimate equation (2.1) twice: (i) I eliminate ψ_i using standard fixed effects tools, the Least Squares Dummy Variable; (ii) I use NACE sector dummies for ψ_i. In Table 2.3 and in Panel A, I report the estimation of equation (2.1). I observe that the signs of the variables under examination are aligned with the economic intuition, and they are all statistically significant at 1%. The value of the coefficient of determination, R^2, is somewhat low, which indicates that there are other factors that result in more than sufficient production capacity. However, with this exercise I am only interested in verifying the direction of the relationship.

In the second exercise, I focus on production and capacity utilization. Our economic intuition suggests that an increase in the production reported in the surveys could be positively associated with an increase in capacity utilization. I check this by estimating the following linear equation:

$$QS_{im} = \beta_0 + \beta_1[U_{i,m-1} - U_{i,m-3}] + \psi_i + \psi_y + \eta_{im}, \tag{2.2}$$

where β_0 and β_1 are parameters to be estimated, U_{im} corresponds to the survey question asking about the percentage of capacity utilization for firm i in month m, QS_{im} indicates the change in past production, ψ_i and ψ_y control for firm and year fixed effects respectively, and η_{im} is the idiosyncratic error. QS_{im} corresponds to survey question A.1: *'During the 3 previous months, your total production has increased/remained un-*

Table 2.3: Consistency of survey responses across questions

PANEL A: Dependent Var. $D3_{im}$			**PANEL B:** Dependent Var. QS_{im}		
INV_{im}	0.137***	0.140***	$U_{i,m-1} - U_{i,m-3}$	0.00506***	0.00508***
XS^e_{im}	-0.0485***	-0.0489***			
L^e_{im}	-0.179***	-0.185***			
U_{im}	-0.00418***	-0.00413***			
Constant	0.336***	0.242***	Constant	0.277***	0.363***
RE/FE	FE	RE	RE/FE	FE	RE
NACE FE	NO	YES	NACE FE	NO	YES
Observations	22,168	22,168	Observations	9,411	9,411
Overall R^2	0.243	0.262	Overall R^2	0.0537	0.0767
Number of firms	791	791	Number of firms	627	627

Estimations with NACE FE were made with Random Effects GLS in pool data (RE). All variables (apart from NACE 2-digit code) are survey questions. NACE fixed effects are taken at the 2-digit level. Fixed year effects are omitted to simplify representation but are included in the estimation. $D3_{im}$ is the sufficiency of production capacity; U_{im} is the percentage capacity utilization; QS_{im} is the recent change of production; INV_{im} is the level of inventories; XS^e_{im} is a forecast about sales; L^e_{im} is a forecast about the number of employees. Complete details about the exact wording of the questions are in the text of this section. *** denotes significance at the 1% level.

changed/decreased'. Responses to QS_{im} are coded with $+1/\ 0\ /\ -1$.

As previously, I estimate equation (2.2) twice: (i) I eliminate ψ_i using standard fixed effects tools; (ii) I substitute NACE sector dummies for ψ_i. In Panel B of Table 2.3 I report the estimation of equation (2.2). These agree with my economic intuition: a reported increase in production is positively and significantly correlated with a three-month increase in % capacity utilization (from $m-3$ to $m-1$).

Having substantiated the internal consistency of survey responses, I now turn to their consistency with the financial statements. Naturally, I expect that the annual growth of real gross sales, x_{iy}, in the income statement, must be positively correlated with the survey question concerning the evolution of sales, XS_{im}. This is equivalent to the linear equation

$$XS_{im} = \beta_0 + \beta_1 x_{iy} + \psi_i + \psi_y + \eta_{im}, \tag{2.3}$$

where XS_{im} indicates the change in past sales and x_{iy} the real growth rate of sales; ψ_i and

Table 2.4: Consistency of survey responses with variables in financial statements

Dependent Variable XS_{im}		
x_{iy}	0.221***	0.227***
Constant	0.155***	0.223***
Observations	24,261	24,261
Number of Firms	785	785
Overall R^2	0.0670	0.0801
RE/FE	FE	RE
NACE FE	NO	YES

Estimations with NACE FE were made with Random Effects pool OLS (RE). NACE fixed effects are taken at the 2-digit level. Fixed year effects are omitted to simplify representation but are included in the estimation. x_{iy} is gross sales growth from financial statements. Significance at the 1% level is indicated by ***.

ψ_y control for firm and year fixed effects respectively; η_{im} is the idiosyncratic error. As previously, I estimate equation (2.3) in two ways: (i) with standard fixed effects tools; (ii) with NACE sector dummies. XS_{im} corresponds to survey question A.2: *During the previous 3 months, your total sales, have increased/remained unchanged/decreased.* Responses to XS_{im} are coded with $+1/\ 0\ /\ -1$.

In Table 2.4, I observe that the monthly survey responses about the change in sales are on average positively and significantly (sig. 1%) correlated with the sales growth rate from the financial statements. In other words, survey responses are on average consistent with the financial statements.

Overall, based on the results in Tables 2.3 and 2.4, I can conclude that survey responses are consistent. Responses are both consistent with each other within the questionnaire, but also with the financial statements. That is, the surveyed firms do not respond contradictory and erratically. In addition, the fact that survey responses are positively correlated with the sales growth from the financial statements indicates that

respondents have a complete knowledge of their firm's activities. I can draw this conclusion, because the financial statements are published the year after the respondents fill in the survey.

2.3 Concluding Remarks

In this chapter I describe the details of the data I am using and assess its quality of information. First, I describe how the data is collected and recorded, and I show that it is a fairly representative sample for the Greek manufacturing sector. Second, using linear regressions I document that responses are well informed and show both internal and external consistency. On average, survey responses are consistent with each other without contradictory responses (internal consistency). Additionally, survey responses concerning past sales developments are also consistent, on average, with the sales growth reported in the financial statements (external consistency). With these findings I am affirming the quality of the information in my data, before using it for the analysis that follows.

Chapter 3

Quantifying the Qualitative Survey Data: Maintaining the Panel Data Structure

A I discussed in Chapter 2, the survey-based expectations are collected each month and are qualitative. Qualitative forecasts are inherently bounded and therefore do not allow to distinguish the major from the minor ones ones. As I show in Chapter 4 that follows, this distinction is pivotal for the behavior of the forecast errors. To pursue this analysis I firstly need to overcome the limitation of the qualitative survey data. In doing so, I draw upon the 'regression approach' developed by Pesaran (1987), as extended by Smith and McAleer (1995). My contribution is that I redesign the regression method to maintain the panel data structure. Pesaran (1987) and Smith and McAleer (1995) derive monthly time-series from firm-month survey observations. What I propose is a similar method that uses the firm-month survey expectations on sales growth to generate firm-year forecasts on sales growth. My quantification methodology is broadly applicable as long as there is quantitative data on realizations and higher frequency qualitative

expectations.

The quantification model I develop comprises of two key equations. In the first one, I show that the unobservable firm-year forecasts are a non-linear function of the the observable firm-month forecasts from the surveys. This equation cannot be estimated because the firm-year forecast on the left-hand side is not observed. In the second equation, I show that the unobserved forecast on the left hand side can be replaced with the realized sales growth rates that are observed in the financial statements. Importantly, all the variables in the second equation are observable and the parameters are the same as in the first one. Therefore, I use the second equation to estimate the parameters using non-linear least squares (NLS), and I show how to control for the unobserved firm heterogeneity (firm fixed effects). Then, by using the estimated parameters and the survey-based forecasts in the first equation, I compute the quantitative firm-year forecasts on sales growth. Having the estimated quantitative forecasts and the observed realized sales growth rates, I compute the ex-post forecast error that I will use in my a analysis in Chapter 4.

I provide evidence of external validity and accuracy for my methodology in a number of ways. First, I show that my quantified estimates on sales growth expectations are fully consistent in terms of sign with the corresponding qualitative survey-based expectations. In a horse race, my methodology also substantially outperforms ordered response models which are potential alternatives for obtaining quantified predictions. Second, I provide evidence for the accuracy of my quantification methodology in terms of the magnitude of firm growth forecasts. In practice, this is challenging to do due to the unavailability of data on quantitative firm-level expectations. Note, that the vast majority of surveys contain qualitative questions about future developments. In fact, this lack of quantitative data highlights the need for and value of my quantification methodology. To overcome

this obstacle, I perform a Monte Carlo exercise that provides a benchmark based on simulated data that mimics the observable data that I have. I find that the forecast errors estimated by implementing my methodology in the artificial data are highly accurate when compared with the forecast errors that can now be observed in the underlying artificial data.

In order to examine the rationality in economic agents' expectations using survey data, the quantification of this data is the imperative first step. Pesaran and Weale (2006) emphasize that in order to examine the full information rational expectations hypothesis in survey data, researchers need to obtain quantitative forecasts first. Among others, Pesaran (1985), Souleles (2004), Bachmann and Elstner (2015) all quantify the qualitative survey-based forecasts and forecast errors in order to test the full information rational expectation hypothesis. Given that I also want to study these properties of the expectation formation, I also need to quantify the qualitative survey forecasts. As I explain below in more detail, my contribution with my quantification is that instead of computing a time-series quantified forecasts, my model provides quantified forecasts of firm-year frequency using all available survey data.

Evidently, the quantification model I develop is not the only available method in the literature, but I argue it is the most appropriate. Recently, Bachmann and Elstner (2015) used the % change of capacity utilization, from the IFO Business Climate survey, to extract forecast errors. To do so, they assumed and verified that potential output remains constant. If potential output is constant, the percentage change of the utilization represents the percentage change of output. According to their methodology, for the firms who expected no change in their output of the following three months, the non-zero percentage change of their reported utilization (i.e. output % growth) is a forecast error. This technique has one important limitation, because, it only uses observations

for which the expectation indicates 'no change'. It ignores roughly 1/2 of the firm-month survey observations that expect either a fall or an increase. As a result, it might ignore the largest forecast errors on which I want to focus. For instance, when firms expect a decrease, while the realized outcome is a positive growth rate.

A second quantification method is the one developed by Theil (1952) and Anderson Jr (1952); also known as the 'probability method'. It provides the theoretical grounds for the 'balance statistics' that are widely used for the published business and consumer sentiment indexes. The specific details about this methodology extend beyond the scope of my analysis, but the reader can find them in all Theil (1952), Anderson Jr (1952), Pesaran (1987) and Smith and McAleer (1995). I avoid using the probability method for a key reason. First, remember that the survey forecasts record an expected rise or a fall or no change. If in a given year all firm responses indicate expectation of only the two out of the three possible outcomes, then the probability method implies that the outcome that was not in firm's forecasts has zero probability in that year and for that firm. However, this is a contradiction, because the probability method assumes a continuous probability function which cannot allow zero probability.[1]

Another approach is the one proposed by Souleles (2004). By estimating ordered logit and probit models, Souleles (2004) uses the latent variable of these ordered response models as a quantified forecast. However, this quantification has two important drawbacks. First, the latent variable is conditional on the information used by the econometrician when estimating the ordered response models and not on the information available to the firm. Econometricians cannot know what information the firm used nor the functional form it used to map that information to a forecast. Second, the latent variable has no unit of measurement and is therefore incomparable to the realized growth rates that I use to compute the forecast errors.

[1]Pesaran (1987) has all the technical details about this.

The absence of a unit of measurement in the ordered response models further stresses the suitability of my methodology estimations in quantifying the forecast errors compared to the alternative quantification methodologies. The estimated values that I will obtain are fitted to the sales growth, so their unit of measurement is the growth rate. This makes them directly comparable with the realized growth rates and hence suitable to compute the forecast error. The estimates of the latent variable in the ordered response models do not have units of measurement and as such are not comparable with the observable realizations. If I wanted to obtain forecast errors using the alternative methods, I would have to use quantified realizations too in order for the forecasts and realizations to be comparable. Estimated realizations however cannot obtain very large values and cannot identify large forecast errors which I show play a vital role in the behavior of forecast errors.

In the following section I develop the quantification model from which I will obtain estimates of quantified forecasts. Subsequently, I will give evidence of external validity of my quantification model. There, I also show that the ordered response models underperform my quantification model in terms of data fit.

3.1 Data

As a reminder my data comprises of firm-level survey responses combined with the financial statements of the firms. I focus on the Manufacturing sector of Greece and for the period 1998-2015. In Chapter 2, I have all the details concerning the data collection and my data sources, the representativeness of my data and the quality of its information. I document in that Chapter that my data is fairly representative of the Greek manufacturing sector and responses show consistency both internally and externally.

For the quantification model in this chapter I use the following (translated) questions of the survey

> Question A.2: *During the previous 3 months, your total sales, has increased/remained unchanged/decreased.*
>
> Question D.2: *During the next 3 months, you expect your total sales to increase/remain unchanged/decreased.*

These qualitative survey responses are coded in the data as +1/0/-1 indicating an increase/remain unchanged/decrease, respectively. In the following, I label the variables that include the responses of firm i in month m to questions A.2 XS_{im}, and to question D.2 XS^e_{im}. The qualitative survey variable on current sales developments, XS_{im}, has a direct quantitative counterpart with sales growth, denoted as x_{iy} for firm i in year y, in the financial statements.

I have matched the aforementioned qualitative survey responses with the realized sales growth from the financial statements of the firm. The cleaned and matched dataset includes 799 firms with 25,764 monthly responses from the survey on the above two questions and 4,104 annual balance sheet observations on sales.[2] In Table 2.1 of Chapter 2, I provide an overview of the firms in my sample. I show there that my sample includes very small firms but also large firms with more than 4,000 employees and annual sales turnover of over six billion Euros.

Before moving to the the quantification model, I make an overview of the survey-based sales forecasts in my final (cleaned and matched) sample. Figure 3.1 shows the distribution of monthly responses to survey question D.2 on firms' expected sales for the following three months. These possible responses, increase/no change/decline, are coded

[2]In Appendix C I document all the cleaning steps for the survey responses. Some additional cleaning was necessary in this chapter as it relates to the quantification.

as +1/0/1, respectively. We observe that more than 50% of the observations indicate increase or decline. This further illustrates my point that the quantification developed by Bachmann and Elstner (2015) is not suitable for my data as it would ignore half of the survey forecasts.

Figure 3.2 shows for each year the share of survey responses on sales growth expectations that indicate an increase/unchanged/decrease (shown in green/orange/blue). The share of responses expecting an increase is higher during the years of the strong boom in Greece that ended in 2008. Accordingly, the share of the responses expecting a decline increases during the subsequent contraction starting from 2009. Finally, in Figure 3.3, I report the number of firm-month survey responses on sales expectations per year. The number of responses is somewhat constant across my sample and well above 1,000, with the exception of the last year, 2015. The reason is that the financial statements are digitized only about 2 years after they have been published. At the time I obtained the data not all financial statements had been digitized so I could not match all the survey responses with the financial statements of that year.

3.2 Quantification Model

In order to obtain continuous forecasts, I develop in this chapter the quantification model that allows me to compute firm-year forecasts of sales growth. Firms sacrifice time and resources at the end of each year to forecast the sales growth of the following year as a whole. That is, at the end of $y-1$ they forecast the following year's sales growth which I cannot observe. I define the annual forecast as

$$x^e_{iy} \triangleq \mathbb{E}\big[x_{iy}|\mathscr{F}_{i,y-1}\big],$$

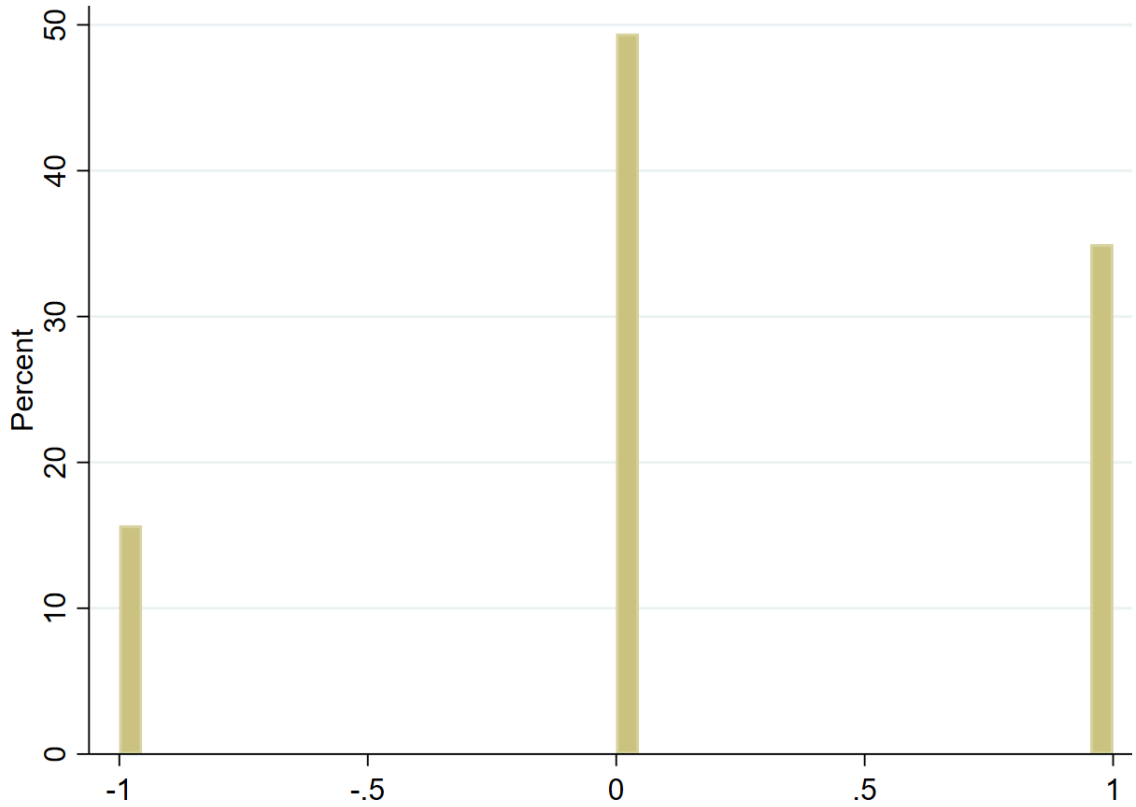

Figure 3.1: **Distribution of Sales Forecasts based on Qualitative Survey Data.** The figure on the left shows the distribution of firm-month sales forecasts based on survey question D.2. The figure on the right shows the distribution of the survey based firm-year sales forecasts when the monthly survey responses are annualized using a yearly weighted average.

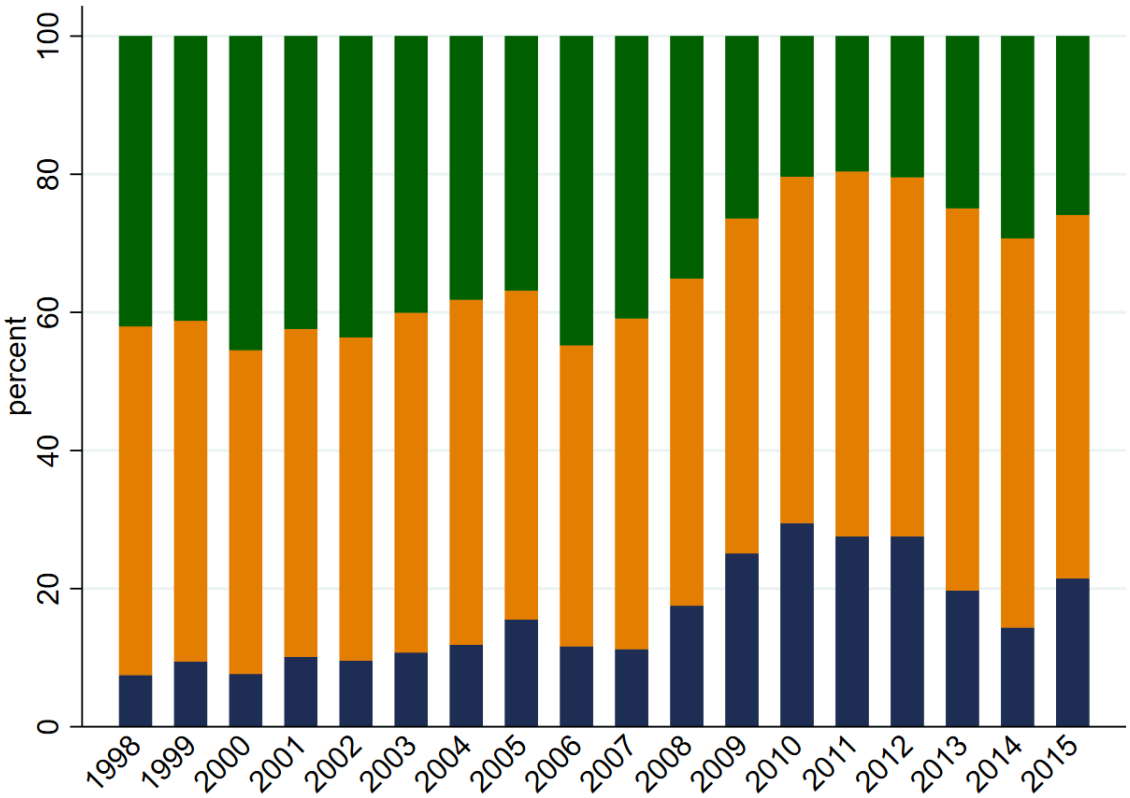

Figure 3.2: **Distribution of Qualitative Survey Responses on Expected Sales Growth over Time (Survey Question D.2).** The figure shows the responses indicating an increase/unchanged/decrease in green/orange/blue as share of total monthly observations per year.

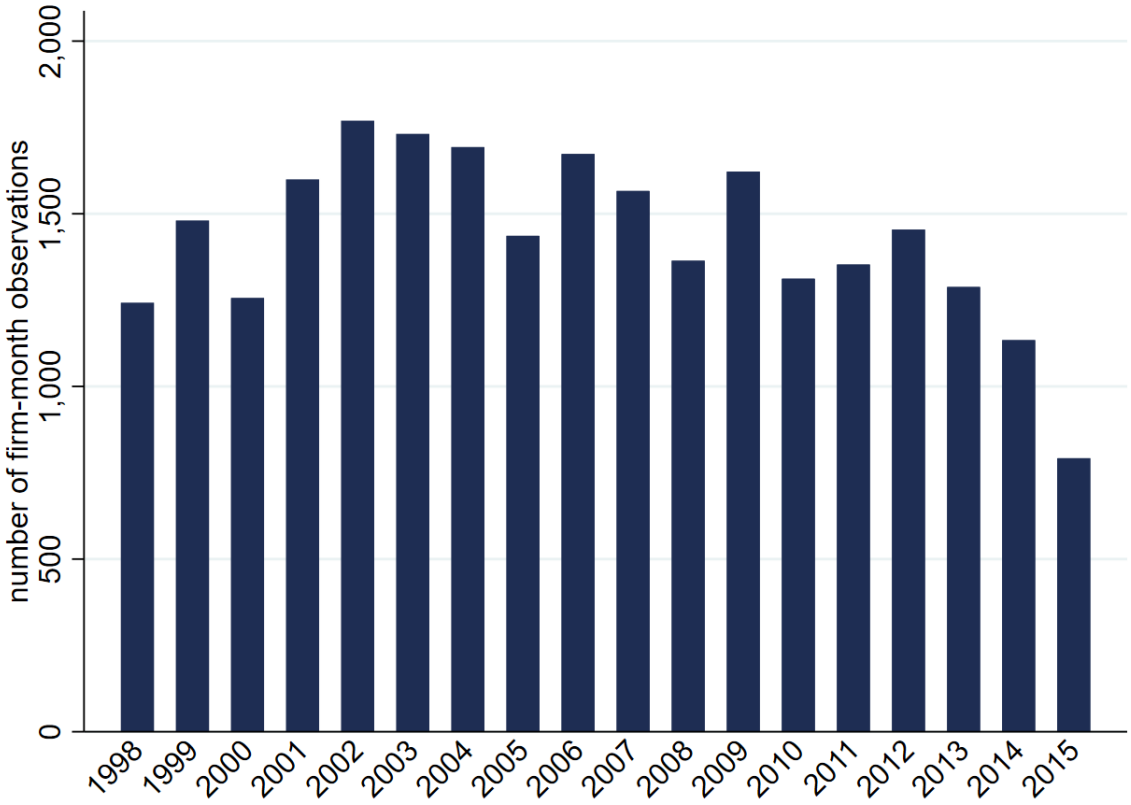

Figure 3.3: **Observations of Qualitative Survey Responses on Expected Sales Growth over Time (Survey Question D.2).** The figure shows the total number of firm-month survey responses per year to question D.2.

with $\mathscr{F}_{i,y-1}$ being the information set available to the firm at the time of the forecast. This quantitative annual forecast is unobservable but with the quantification I will impute its value from the observable monthly qualitative forecasts.

Second, I define firm i's expectation about average sales growth in the next three months as $x^e_{im} \triangleq \mathbb{E}\big[x_{i,\{m,m+1,m+2\}}|\mathscr{F}_{i,m-1}\big]$, where $x_{i,\{m,m+1,m+2\}}$ is the average growth rate of sales for the following three-month period.[3] Note that the monthly forecast is based on information available to the firm at time $m-1$. Naturally, the firm at any month of the following year it has at least as much information as at the end of the preceding year when it posted the annual forecast. Formally, this is given by $\mathscr{F}_{i,y-1} \subseteq \mathscr{F}_{i,m-1\in y} \subseteq \mathscr{F}_{i,m\in y}$, which implies that as the time passes the information of the firm is non-decreasing.

Next, I follow Pesaran (1987) and assume that for each firm the monthly expectations are linearly correlated with their annual counterparts. Intuitively, this means that firms do not fully update their information every single month. Instead, they partially rely on the forecast for the whole year that they did at the end of the preceding year. One reason for that may be costly information collection and processing. I further distinguish between positive and negative monthly forecasts to allow for this linear correlation to be asymmetrical (see Smith and McAleer (1995)). This is my first identifying assumption which is equivalent to

$$x^{e,+}_{im} = \alpha + \gamma_1 x^e_{iy} + \nu^+_{im}, \qquad \text{and} \qquad x^{e,-}_{im} = -\beta + \gamma_2 x^e_{iy} + \nu^-_{im}. \qquad \text{[ID1]} \qquad (3.1)$$

$x^{e,+/-}_{im}$ is the quantitative monthly forecast for the following three-month period as a whole, α, β, γ_1 and γ_2 are parameters, and νs are the error terms. Equations (3.1)

[3] I define x^e_{im} for November (December) to include expectations about the next two (one) months only. This is consistent with my treatment of the survey data for these months (which are weighted with 2/3 (1/3)) which is standard in the literature — for details see Appendix C. This scheme avoids double counting of months.

do not attempt any causal inference, they merely state that the monthly forecasts are linearly correlated with the corresponding annual expected growth rate. The horizon and the timing of the monthly forecasts is consistent with those in the surveys: at each month m, firm i makes a forecast for the average growth rate of following three months, that is for $\{m, m+1, m+2\}$. Note that I cannot observe the quantitative monthly forecasts of equation (3.1), but I algebraically eliminate them at a later stage. I additionally assume that the error terms ν_{im}^{+} and ν_{im}^{-} are independently distributed across firms. Any month-over-month serial correlation in these error terms is not concerning either, because at a later stage I will aggregate them at the firm-year frequency.

Next, I assume that the expected growth rate can be decomposed into a weighted average of its monthly positive and negative forecasts. Zeros have no effect for the quantification. This assumption formally is

$$x_{iy}^{e} = \mathbb{E}_{i,y-1} \sum_{m \in y} W_{im}^{+} x_{im}^{e,+} + \mathbb{E}_{i,y-1} \sum_{m \in y} W_{im}^{-} x_{im}^{e,-}, \tag{3.2}$$

where the operator $\mathbb{E}_{i,y-1}(\cdot)$ is a shorter version of the conditional expectation $\mathbb{E}\big[(\cdot)|\mathscr{F}_{i,y-1}\big]$ that I defined earlier. The weights are defined as $W_{im}^{+} = W_{im}\mathbb{1}_{[XS_{i,m}^{e}=+1]}$ and $W_{im}^{-} = W_{im}\mathbb{1}_{[XS_{i,m}^{e}=-1]}$ and consist of two components. The first component in each weight, W_{im}, accounts for the fact that some months have a higher level of firm sales than others and therefore represent a larger share of the final annual forecast. It is defined as

$$W_{im} \triangleq \frac{w_{im}}{\sum_{m \in y} w_{im}}, \tag{3.3}$$

where w_{im} is the ratio of the seasonally unadjusted over the seasonally adjusted real gross value added. Intuitively, when this ratio is larger that 1, unadjusted gross value added is higher than the seasonally adjusted one, meaning that during this month value added

is above normal levels, and this month is more important than others for the annual outcome. While a purely theoretical decomposition would allow for individual weights for each firm, in my practical implementation here, data availability limits the design of w_{im} to be the same across all firms in the manufacturing sector at quarterly frequency.[4] The second component of the weights is dummy variables that take a value of unity if the expected sales growth rate of monthly frequency is either positive, $\mathbb{1}_{[XS^e_{i,m}=+1]}$, or negative, $\mathbb{1}_{[XS^e_{i,m}=-1]}$.

The distinction in positive and negative monthly forecasts is essential for the quantification. While I do not observe the quantitative expectations of sales growth in equation (3.2) nor in equations (3.1), I can observe the qualitative expected growth of sales, $XS^e_{i,m}$. In the next step, this distinction permits me to use the qualitative survey-based forecasts and algebraically eliminate the unobservable quantitative monthly forecasts $x^{e,+}_{im}$ and $x^{e,-}_{im}$. In particular, the linear relationship in equations (3.1) can be used to eliminate the unobserved variables $x^{e,+}_{im}$ and $x^{e,-}_{im}$ from equation (3.2). Combining equations (3.1) and (3.2) and solving for x^e_{iy} yields (detailed derivations are shown in Appendix D.1)

$$x^e_{iy} = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \qquad \text{with} \qquad \xi_{iy} = \frac{\mathbb{E}_{i,y-1} \sum_{m \in y} \left(W^+_{im} \nu^+_{im} + W^-_{im} \nu^-_{im} \right)}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}, \tag{3.4}$$

where, to ease notation, I have defined

$$P_{iy} \triangleq \sum_{m \in y} W_{im} \mathbb{1}_{[XS^e_{im}=1]}, \qquad \text{and} \qquad N_{iy} \triangleq \sum_{m \in y} W_{im} \mathbb{1}_{[XS^e_{im}=-1]}. \tag{3.5}$$

[4] I use 2-digit seasonally unadjasted and adjusted real gross value added for the manufacturing sector from Eurostat, Table namq_10_a10 for Greece, both in 2005 Chain Linked Volumes. I use value added since information on sales is not available at monthly or quarterly frequency. In Appendix E I show that quantitative sales forecasts are almost identical if I use a simple firm-year arithmetic mean instead instead of a weighted average in equation (3.2).

P_{iy} (N_{iy}) denotes the weighted share of months per year that record a rise (decline) in expected sales of firm i and can be computed from the surveys. The survey question linked to the variable XS^e_{im} covers expectations about the next three months. By construction the aggregation to annual frequency in P_{iy} and N_{iy} is almost identical if the underlying survey variable would refer to expectations about one month only. The only differences would stem from survey responses in November and December. I document in Appendix C that responses in these months are weighted with 2/3 and 1/3, respectively, to reflect the period of expectations that belongs to the current calendar year. This is a common issue in the use of survey data. My results are robust to alternative weightings, as I discuss later.

By observing the survey-based qualitative forecasts I can directly compute P_{iy} and N_{iy}. However, I cannot estimate equation (3.4) since I do not observe quantitative expectations of annual sales growth, x^e_{iy}. In fact, deriving quantitative sales growth expectations was my goal in the first place. Instead, if one has estimates for the parameters and knowledge of the error term — and given that P_{iy} and N_{iy} are observable — they could use equation (3.4) to derive fitted values for x^e_{iy}. Indeed, in the following steps I show how to obtain estimates of the parameters. First, I reduce the unavailability of x^e_{iy} to an omitted variable problem and show that it does not affect the estimates of the parameters. Second, I show how to control for the the unobserved firm heterogeneity of the error term ξ_{iy}.

Omitted Variable Problem

We know that for each firm i, realized sales growth in year y is the sum of expected sales growth in that year plus a forecast error, $x_{iy} = x^e_{iy} + x^{fe}_{iy}$. Using this expression to

replace x^e_{iy} in equation (3.4) I obtain after rearranging

$$x_{iy} = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + x^{fe}_{iy} + \xi_{iy}. \tag{3.6}$$

In principle, this equation can be estimated, as I can observe the realized annual sales growth, x_{iy}, from the financial statements. While the forecast error, x^{fe}_{iy}, is still unobserved, estimating equation (3.6) without this variable is simply an omitted variable problem that adds to the error term. I show in this subsection that the omitted variable does not affect the estimates of the parameters.

To ease the notational burden in this section, I use equation (3.4) and define the conditional expectation of the quantitative sales growth expectation as

$$\tilde{x}^e_{iy} \triangleq \mathbb{E}\big[x^e_{iy} \big| P_{iy}, N_{iy}\big] = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{3.7}$$

I aim to obtain consistent estimates of the parameters in equation (3.6) using Non-Linear Least Squares. To do so, the composite error term $x^{fe}_{iy} + \xi_{iy}$ needs to be mean independent of the non-linear function $\tilde{x}^e_{iy}$ (see Davidson and MacKinnon (2004)). I show here that this is true. Note that the forecast error, x^{fe}_{iy}, is mean independent from the forecast x^e_{iy}. Indeed, $\mathbb{E}\big[x^{fe}_{iy} \big| x^e_{iy}\big] = \mathbb{E}\big[x_{iy} - x^e_{iy} \big| x^e_{iy}\big] = x^e_{iy} - x^e_{iy} = 0$. Note that this does not imply rational expectations, because mean independence from the forecast does not imply mean independence from the information set that was used by that forecast. The latter only implies full information rational expectation. Note that when a firm makes a forecasts, it maps its information to a value that is the forecast.

Since $\mathbb{E}\big[x^{fe}_{iy} \big| x^e_{iy}\big] = 0$ holds, it also implies that $\mathbb{E}\big[x^{fe}_{iy} \big| \tilde{x}^e_{iy}\big] = 0$. I provide a proof of this statement in Appendix D.2. Intuitively, firms' expected sales growth, x^e_{iy}, can-

not ex-ante forecast their forecast error, otherwise firms would have incorporated this information in their expectation to reduce the forecast error. The same must hold then also for any estimates, $\tilde{x}^e_{iy}$, of firms' sales growth expectations. Having proven that $\mathbb{E}\big[x^{fe}_{iy}\big|\tilde{x}^e_{iy}\big] = 0$ allows me to ignore the unobserved forecast error that is an omitted variable.

Unobserved Firm Heterogeneity

Having established the forecast error's mean independence of $\tilde{x}^e_{iy}$, and in order to obtain consistent estimates of the parameters in equation (3.6), it remains to be shown that $\mathbb{E}\big[\xi_{iy}\big|\tilde{x}^e_{iy}\big] = 0$. A sufficient condition for mean independence of the error term, $\mathbb{E}\big[\xi_{iy}\big|\tilde{x}^e_{iy}\big] = 0$, to hold is that $\mathbb{E}[\xi_{iy}|\{XS^e_{im}\}_{m\in y}] = 0$. In Appendix D.2 I provide a formal proof that this is indeed a sufficient condition. This leaves me with the task to control for the unobserved firm heterogeneity that is likely to make ξ_{iy} correlated with $\{XS^e_{im}\}_{m\in y}$. I do so here.

Equation (3.4) shows that the numerator of ξ_{iy} depends on the error terms ν^+_{im}, ν^-_{im}, and on XS^e_{im}. In this section I show how to control for the effect of the unobserved heterogeneity hidden in this numerator. For that, my second identifying assumption (ID2) is to assume that the error term in equation (3.4) can be decomposed as

$$\xi_{iy} = \frac{\psi_i + \vartheta_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} \qquad \text{with} \qquad \mathbb{E}[\vartheta_{iy}|\{XS^e_{im}\}_{m=1,\ldots,T_i}] = 0, \qquad \text{[ID2]}$$

where ψ_i captures the effect of unobserved firm heterogeneity on sales growth. ϑ_{iy} is an idiosyncratic error which is mean-independent of XS^e_{im} for all m. Note that the notation $\{XS^e_{im}\}_{m=1,\ldots,T_i}$ in ID2 denotes the entire history of months m for variable XS^e_{im}, where T_i is firm i's total number of monthly observations.[5] This stems from the fact that the

[5]This notation is distinct from $\{XS^e_{im}\}_{m\in y}$, used above, which refers to all months m in year y.

unobserved firm heterogeneity is firm-specific and as such affects the entire history of the firm forecasts. Therefore, what is left in the error after the firm-specific effect is necessarily mean independent from the firm specific heterogeneity.

ID2 is a standard assumption to deal with unobserved heterogeneity. Essentially, ID2 says that the composite error term in the numerator of equation (3.4) can be broken down into two components. The first one is the firm-specific unobserved heterogeneity that is likely to cause endogeneity. The second one is an idiosyncratic error which is exogenous to the firm behavior. The unobserved firm heterogeneity, ψ_i, is in fact an omitted variable and is endogenous. The reason is that firm heterogeneity is related to the entire history of XS^e_{im}, so that

$$\mathbb{E}[\psi_i + \vartheta_{iy}|\{XS^e_{im}\}_{m=1,\ldots,T_i}] = \mathbb{E}[\psi_i|\{XS^e_{im}\}_{m=1,\ldots,T_i}] \neq 0,$$

from assumption ID2. To control for unobserved heterogeneity, I need to approximate $\mathbb{E}[\psi_i|\{XS^e_{im}\}_{m=1,\ldots,T_i}]$.

The structure of the non-linear equation (3.4) that I want to estimate does not allow me to derive an estimator for ψ_i analytically, and I cannot use dummy variables either, because the cross-sectional dimension is very large. Another possibility would be to linearize (3.4) with Taylor series expansion. However, Taylor expansion around a specific point holds locally, only in a small area around this point, otherwise the higher order terms that will appear into the linearized regression error will be endogenous to the lower order terms included in the estimation. To avoid this endogeneity problem, I would need to use local polynomial fitting methods which are too complex, both algebraically and computationally (see Fan and Gijbels (1996)).

A widely used approximation for this purpose is the one suggested in Mundlak (1978) (see for example, Bartelsman et al. (1994), Semykina and Wooldridge (2010), Kosova

(2010) and Triguero and Córcoles (2013)). This approximation is the standard tool used in non-linear models in panel data. In linear models, it is equivalent to the least squares dummy variable and the standard within estimator (when the time periods are 'infinite'). The original Mundlak (1978) specification is linear, but in the following I additionally include a second-order term due to the non-linearity of equation (3.4). Therefore, my third identifying assumption is that the conditional expectation of the unobserved firm heterogeneity in the error term ξ_{iy} is

$$\mathbb{E}[\psi_i|\{XS^e_{im}\}_{m=1,\ldots,T_i}] \;=\; \delta_1\overline{XS^e_i} + \delta_2(\overline{XS^e_i})^2, \qquad \text{[ID3]}$$

where δ_1 and δ_2 are coefficients. This results in the following auxiliary regression for ψ_i

$$\psi_i \;=\; \delta_1\overline{XS^e_i} + \delta_2(\overline{XS^e_i})^2 + \omega_i, \tag{3.8}$$

where ω_i is the the part of the firm-specific heterogeneity that is mean independent from the survey expectations, that is $\mathbb{E}[\omega_i|\{XS^e_{im}\}_{m=1,\ldots,T_i}] = 0$; and $\overline{XS^e_i} = \frac{1}{T_i}\sum_{m=1}^{T_i} XS^e_{im}$ is the simple arithmetic mean of the survey variable XS^e_{im} across time for each firm i. I now substitute equation (3.8) for ψ_i in the numerator of ξ_{iy}, obtaining

$$\xi_{iy} = \frac{\delta_1\overline{XS^e_i} + \delta_2(\overline{XS^e_i})^2 + \omega_i + \vartheta_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{3.9}$$

The Final Equation to be Estimated

In the preceding steps I showed that the observed realized growth rate can substitute for the unobserved annual forecast on the left hand side to fit the non-linear regression. Additionally, I provided a way to approximate the unobserved firm heterogeneity. With these I can now derive the final equation that I will estimate. I substitute equation (3.9)

into equation (3.6) and obtain

$$x_{iy} = \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1 \overline{XS_i^e} + \delta_2 (\overline{XS_i^e})^2}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \tilde{\xi}_{iy}, \tag{3.10}$$

where

$$\tilde{\xi}_{iy} = x_{iy}^{fe} + \frac{\omega_i + \vartheta_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{3.11}$$

Overall equation (3.10) is estimable because the error term $\tilde{\xi}_{iy}$ is mean-independent of the explanatory variables. I provide a formal proof of this statement in Appendix D.2. The variables $\overline{XS_i^e}$ and $(\overline{XS_i^e})^2$ proxy and control for the unobserved heterogeneity in equation (3.6). With equation (3.10), I can obtain consistent estimates of the coefficients of interest, α, β, γ_1 and γ_2 which I can then use in equation (3.4) to obtain the quantified forecasts. Additionally, the error term, $\tilde{\xi}_{iy}$, in equation (3.10) is likely to be heteroscedastic and autocorrelated within each firm. Therefore, I will use the heteroscedasticity robust estimator for the standard errors which addresses both problems — this robust estimator treats errors as clustered within cross-sectional units. Finally, I estimate equation (3.10) separately for the boom period (1998-2008) and for the bust period (2009-2015), to allow for the parameters to be state dependent.

3.2.1 The Quantification Method in a Nutshell

I have derived two nonlinear equations. First, equation (3.4) shows how the unobserved quantitative expected sales growth can be a non-linear function of the observable survey variables. This equation shows that once one has estimates of the parameters, they can compute the quantitative forecasts. Second, equation (3.10) shows that the observed quantitative annual sales growth can be used as a dependent variable to fit the non-linear equation. Moreover, the identifying assumptions ID2 and ID3 ensure that the coefficient

estimates obtained from (3.10) are consistent. Essentially, both of these equations have the same parameters. I showed that, I can estimate the parameters from equation (3.10) using Nonlinear Least Squares (NLS), and use these estimated parameters in equation (3.4) to derive imputed values for quantitative expectations on sales growth. To reiterate, this methodology to quantify forecasts is generally applicable to variables other than sales growth. It is applicable to any qualitative (survey based) variable on future developments, as long as a quantitative corresponding variable on realization is available.

The practical implementation of the estimation methodology to derive quantitative forecasts on sales growth can be summarized in the following steps:

1. Compute the weighted shares of months per year that record a rise (decline) in expected sales P_{iy} (N_{iy}) from survey data, using equation (3.5).

2. Compute the firm heterogeneity proxies, $\overline{XS_i^e}$ and $(\overline{XS_i^e})^2$, based on the arithmetic mean (across time for each firm i) of the qualitative survey variable XS_{im}^e .

3. Estimate equation (3.10) using NLS. Run the estimation separately for the boom ($y \leq 2008$) and bust period ($y > 2008$).

4. Use the NLS estimated coefficients α, β, γ_1 and γ_2 and the observed values of P_{iy} and N_{iy} to compute the quantified forecasts, $\hat{x}_{iy}^e$, from equation (3.4). Note that these estimates do not include the firm heterogeneity proxies.

The difference between the sales growth rate available from the financial statements, x_{iy}, and the quantified forecast on sales growth for the corresponding year, $\hat{x}_{iy}^e$, then gives the quantified forecast error on sales growth, $\hat{x}_{iy}^{fe}$. In what follows, I will drop the hat from the expression for the forecasts and the forecast errors to ease notation.

Table 3.1: NLS Estimation of Equation (3.10).

	(1)	(2)
Coefficients	Dependent Variable: x_{iy}	
α	0.190**	0.104**
β	0.151*	0.238***
γ_1	-0.366	-0.446
γ_2	-0.179	0.0712
Firm-Year Observations	2,471	1,397
R^2	0.043	0.057
Period	$y \leq 2008$	$y > 2008$

Fixed effects proxies of equation (3.10) are omitted – but are included in the estimation – to maintain a simple representation. I use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance. Column (1) shows results for the boom period up to 2008 and column (2) for the following recession.

3.3 Estimates of the Quantification Model

I report the parameter estimates of the NLS estimation of equation (3.10) in Table 3.1. Column (1) shows estimation results for the boom period up to 2008 and column (2) for the following recession. As a reminder, α and $-\beta$ are respectively the constant terms in the positive and negative continuous monthly forecasts of ID1 (equation (3.1)). The table shows that the constant of the positive monthly forecasts is larger during the boom than in the bust which is consistent with the economic intuition. Moreover, the constant of the negative monthly forecast is lower during the bust than in the boom, which is also consistent with our economic intuition.

In Appendix E and in Table 1.E, I show that these estimates are robust to using alternative weighting for equation (3.3). There, I use the firm-year simple arithmetic mean of the firm-moth survey responses. I also show that the obtained quantified values are remarkably close to the ones I obtain in my baseline estimates.

Descriptive Statistics on the Quantified Forecast Errors

After having obtained the estimates of the parameters, I use them in equation (3.4) and compute the quantified forecasts on sales growth. I also compute the quantified forecast errors on sales growth using the realized sales growth in the financial statements. I will now provide an overview of the characteristics of the estimated forecast errors that I will study in Chapter 4. Subsequently, in section 3.4, I will examine the accuracy and the external validity of my estimates.

Figure 3.4 shows the distribution of forecast errors. I report the moments of this distribution in Table 3.2. Unless I explicitly describe it differently, in what follows I report the values of forecast errors on sales growth rates with base 1 — a rate of 1 is 100%. The average forecast error in my sample is 0 zero and is slightly larger than the median (-0.03). This implies that for half of the forecast errors, the sales growth forecast is at least three percentage points more optimistic than the subsequent realization. Overall, some of forecast errors made by firms are small (in absolute value), as these are centred close to zero, but still a the majority of the forecast errors made are quite substantial. Table 3.2 also shows that the distribution of the forecast errors is very stable across the boom and the bust period.

Since for my analysis I will be concerned with major forecast errors, I also provide some statistics about these. For this purpose, I classify the top and bottom 26% of forecast errors to be major, which is in line with the estimates for this threshold obtained in Chapter 4. With this classification, at the lower 26 percentiles, firms expected sales growth to be 14.3 percentage points higher than subsequently realized. Similarly, at the upper 26 percentiles, firms expected sales growth to be 8.6 percentage points lower than subsequently realized. Hence, the remaining 48% of forecast errors in the center of the distribution, which I call minor, are still economically significant. Interestingly, Table

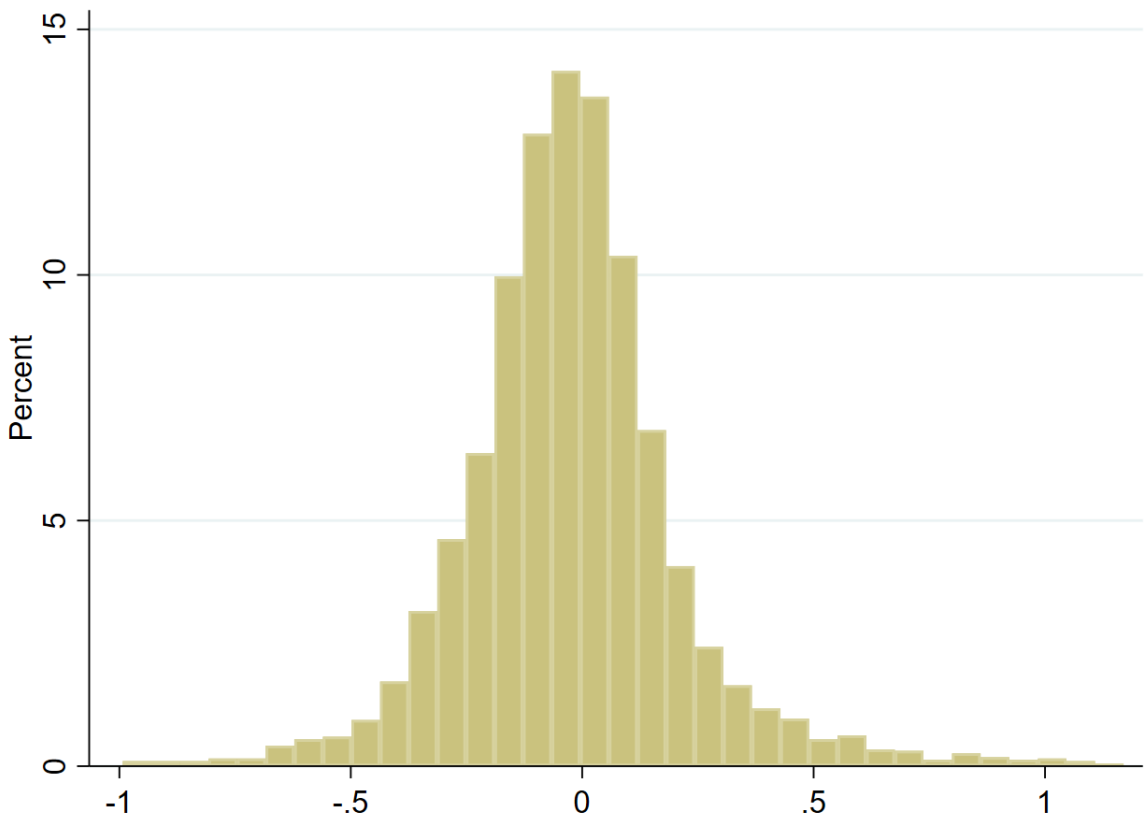

Figure 3.4: **Distribution of Annual Quantified Sales Growth Forecast Errors.** The 1% of forecast errors at the top of the distribution are omitted to ease visibility. Forecast error values are reported with base 1.

3.2 shows that the major and minor forecast errors are evenly distributed across the boom and the severe depression periods in the sample. During both periods their shares are close to the 26% of the full sample which is imposed as a result of the threshold estimates of Chapter 4.

How are these major and minor forecast errors distributed over different dimensions the sample? Panel A of Table 3.3 sorts the sample according to the share of major forecast errors in a firm's total number of observations. For a large number of firms (460 of the total of 785 firms) this share is positive and up to 80%, so that they make major as well as minor forecast errors. These firms are present in the sample for a relatively long time as they also account for the vast majority of the firm year observations (3,119 out

Table 3.2: Descriptive Statistics for Quantified Sales Growth Forecast Errors.

	Mean	Median	Stand. Dev.	Share of Forecast Errors (in %) Major Negative	Minor	Major Positive
Full Sample	0.00	-0.03	0.34	26	48	26
Boom	0.01	-0.02	0.34	24	49	27
Bust	-0.02	-0.05	0.35	30	46	24

Forecast error values are reported with base 1. Major forecast errors are defined for the purpose of this table as the 26% of forecast errors at the top and bottom of the distribution. The boom (bust) period spans the years 1998-2008 (2009-2015).

Table 3.3: Major Forecast Errors (MaFE) and Different Cuts of the Sample.

Panel A: Sorting: Share of MaFE in Firm's Observations				**Panel B:** Sorting: Total Net Assets	
Share of MaFE	# of Firms	# firm-year obs. with MaFE	Total firm-year observations	Percentile of Total Net Assets	Share of MaFE in firm-year obs.
0%	111	0	194		
(0%,40%]	129	258	963	(0%,40%]	54.74%
(40%,80%]	331	1238	2156	(40%80%]	51.13%
(80%,99%]	35	195	235	(80%,100%]	48.09%
100%	179	320	320		

Major forecast errors are defined for the purpose of this table as the upper or lower 26% of forecast error distribution. The percentile of total net assets has been determined using firm's average percentile in the pool distribution.

of total 3,868). 179 (111) firms make exclusively major (minor) forecast errors, however these account only for 320 (194) firm-year observations in the sample and hence are quite short lived. Panel B of Table 3.3 shows that the share of major forecast errors in the total observations of a firm is relatively constant across different firm sizes as well. It varies between 48% and 55% across the firm size distribution where the larger firms make slightly fewer major forecast errors.

The evidence in Table 3.3 shows that, independent of their size, most firms make major as well as minor forecast errors. Table 3.4 provides the average year-on-year transition matrix among minor, positive and negative major forecast errors for the pooled data. It suggests that firms do not tend to make many consecutive major positive forecast errors, but that major and minor forecast errors are likely to alternate. Following a negative (positive) major forecast error in year $y-1$, the probability of making another

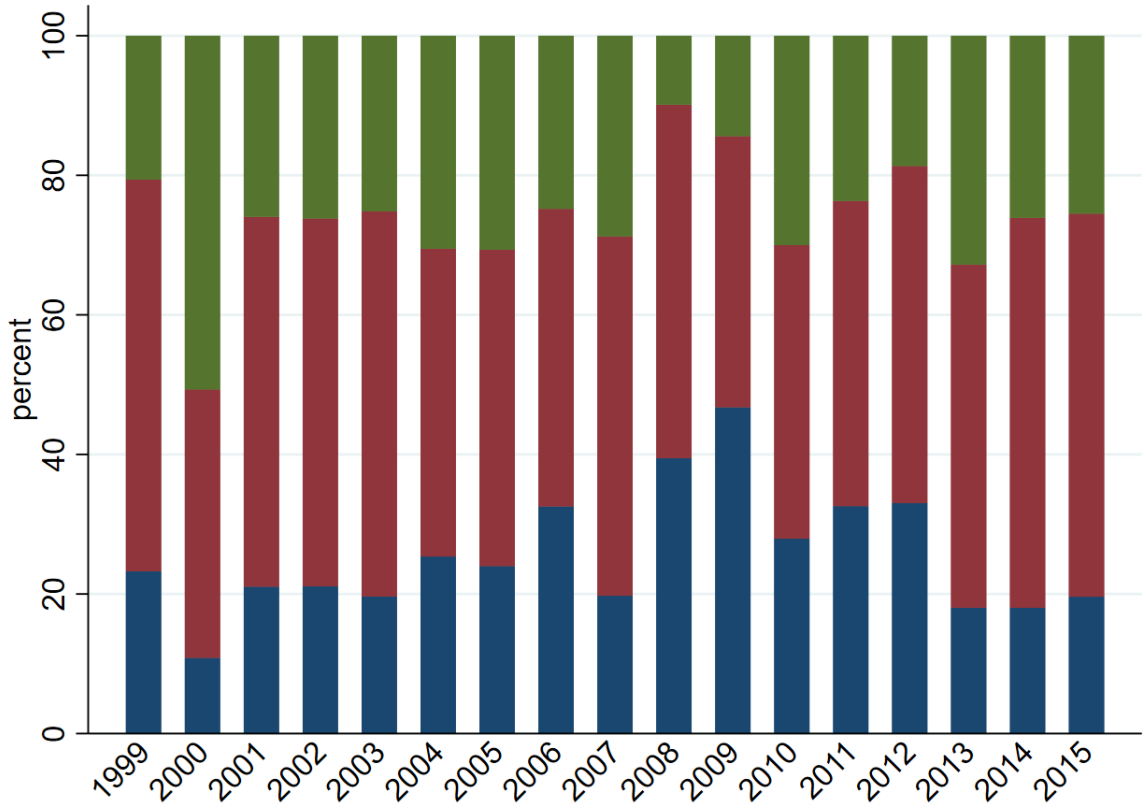

Figure 3.5: **Distribution of Quantified Sales Growth Forecast Errors across years.** Blue (green) indicates the share of major negative (major positive) forecast errors and red stands for the share of minor forecast errors. Major forecast errors are defined for the purpose of this figure as the 26% of forecast errors at the top and bottom of the distribution.

major negative (positive) forecast error in year y is always lower than the probability of making a minor forecast error. Furthermore, the likelihood of being in the left or right tail is approximately equal. Finally, in Figure 3.5, I show the share of observations classified as major positive/negative or minor forecast errors per year. It is evident that the share across these classifications can vary somewhat across years. Notably, in 2009, the first year of the Greek depression, the share of negative major forecast errors increased.

Overall the above evidence suggests that major forecast errors are distributed relatively evenly across all firms (when sorted by size) and across the within-firm obser-

Table 3.4: Transition matrix of Major Forecast Errors (MaFE) and Minor Forecast Errors over Time.

	Negative MaFE in y	Minor FE in y	Positive MaFE in y	Total
Negative MaFE in year $y-1$	30.41%	40.40%	29.19%	100.00%
Minor FE in year $y-1$	22.26%	53.79%	23.96%	100.00%
Positive MaFE in year $y-1$	27.79%	44.99%	27.22%	100.00%

Major positive (negative) forecast errors are defined for the purpose of this table as the upper (lower) 26% of forecast error distribution.

vations. Forecast errors are not highly persistent and both major and minor forecast errors tend to alternate. Additionally, in this section I documented that the share of major positive and negative forecast errors is stable across the boom and bust periods in the sample.

Now that the estimates of quantitative forecasts are available, how can I assess the performance of the quantification model? To do so, I would need to observe quantitative forecasts. However, the lack of this data is why I need a quantification method. To overcome this limitation and assess the performance of the model, I run two exercises in the following section. First, I show that the quantitative forecasts of my quantification massively outperform the estimates from the ordered response models. Second, I run a Monte Carlo simulation on artificial data mimicking the characteristics of my data. With the simulation results I show that the quantification model I propose is very accurate in imputing the quantitative forecasts. These exercises add validity to the quantification model.

3.4 External Validity and Accuracy of the Methodology

In this section, I conduct two exercises to demonstrate the external validity of the quantification methodology. In the first exercise, I use the qualitative firm forecast data from

the survey as a benchmark and test whether my quantified estimates are accurate in terms of the sign of expected sales growth. I also perform a horse race with ordered response models that are an alternative quantification approach. In the second exercise, I test the accuracy of my quantification methodology in terms of the magnitude of firm growth forecasts by conducting a Monte Carlo experiment using artificial datasets. In the following, I discuss these two exercises in turn.

3.4.1 Directional Consistency of Estimated Forecasts with the Survey Data

For the first exercise, I use the observed survey data on the direction of expected sales growth to benchmark how well the quantified forecasts match the direction of expected sales growth. To facilitate the comparison of the monthly survey data with the annual forecast estimates, I annualize the survey responses by computing a weighted yearly average $\sum_{m \in y} W_{im}[XS^e_{im}]$, where the weights are based on equation (3.3). While the annualized survey forecasts cannot provide a detailed indication about the size of the forecasts, as they are based on trinomial and purely qualitative monthly data, they can still be informative about the direction of the observed forecasts.

To benchmark the quantified forecasts against the annualized survey-based qualitative forecasts, I split responses in each of these two variables into three categories — positive, zero or negative — and cross-tabulate them. In Panel A.1 in Table 3.5 I report how well the quantified forecasts match the direction of the annualized observable ones. The main diagonal shows the share of observations that are directionally consistent across the two variables when classified as either positive, zero or negative. Overall, the direction of my quantified forecasts are highly consistent with the one of the annualized survey responses — their direction coincides for 93.98% of all observations (the sum of

Table 3.5: Directional consistency between survey-based sales forecasts and forecasts based on different quantification methodologies (share in total observations)

	Entire Sample			**Restricted Sample**		
	Panel A.1: NLS			**Panel A.2: NLS**		
	Negative	Zero	Positive	Negative	Zero	Positive
Negative Forecasts	23.94%	0.00%	1.45%	11.21%	0.00%	0.00%
Zero Forecasts	0.26%	14.71%	0.34%	0.00%	56.96%	0.00%
Positive Forecasts	3.98%	0.00%	55.33%	0.00%	0.00%	31.83%
	Directional Consistency: 93.98%			Directional Consistency: 100.00%		
	Panel B.1: Ordered Logit			**Panel B.2: Ordered Logit**		
	Negative	Zero	Positive	Negative	Zero	Positive
Negative Forecasts	4.89%	17.67%	3.54%	5.73%	5.62%	0.54%
Zero Forecasts	0.42%	8.85%	6.18%	1.51%	32.86%	22.92%
Positive Forecasts	1.71%	25.96%	30.77%	0.54%	9.62%	20.65%
	Directional Consistency: 44.51%			Directional Consistency: 59.24%		
	Panel C.1: Ordered Probit			**Panel C.2: Ordered Probit**		
	Negative	Zero	Positive	Negative	Zero	Positive
Negative Forecasts	4.41%	18.18%	3.51%	5.19%	6.16%	0.54%
Zero Forecasts	0.37%	8.94%	6.15%	1.30%	33.19%	22.81%
Positive Forecasts	1.46%	26.41%	30.57%	0.54%	9.73%	20.54%
	Directional Consistency: 43.92% %			Directional Consistency: 58.92%		

Rows refer to forecasts on sales growth based on annualized weighted average of the firm-month survey responses. Variables in columns refer to estimates for quantified sales growth forecasts using Non-Linear Least Squares (Panel A). For the ordered choice models (Panel B and C) I used the direction of the sales growth predicted by the model at the firm-month level instead of the latent variable and then I took their annualized weighted average. The restricted sample only considers annualized survey observations for which, in a given year, all underlying monthly observations report forecasts in the same direction.

the main diagonal).

The small share of observations for which the directions do not coincide can be explained by the absence of information on scale in the qualitative survey data. In practice, even if the majority of all monthly forecasts in one year point in the same direction, a single large monthly forecast in the opposite direction could dominate the annual response. This however cannot be captured by annualizing purely qualitative monthly forecasts. For this reason, I also report in Table 3.5 results based on a restricted sample that only includes annualized observations for years in which all underlying monthly survey responses indicated sales forecasts in the same direction. This ensures

that the direction implied by the annualized survey data is accurate for all considered observations. Panel A.2 shows results for this restricted sample which comprises 26% of the observations of the full sample used in Panel A.1. It is evident that now the direction of all quantified forecasts is consistent with the ones of the annualized survey responses. What is more, results are fully directionally consistent even if I consider annualized observations for which at least 67% of all underlying monthly survey responses of a particular year indicated sales forecasts in the same direction. This comprises 39% of the observations of the full sample used in Panel A.1.

These results provide a first indication of the quality of my quantified estimates. Next, I run a horse race with alternative ways to quantify sales growth forecasts — namely, ordered response models such as logit and probit. To economize on space and maintain the reading flow, I only outline the details of these alternatives in Appendix F. Panels B.1 and C.1 in Table 3.5 show the fit of forecasts based on ordered response models with the annualized survey data. Again, the observations in each variable have been split into three categories — positive, zero or negative — before I cross-tabulate the three directions. The overall share of observations that exhibit directional consistency between the annualized survey data and the forecast estimates is only about 45% for both ordered logit and probit. For the restricted sample shown in Panels B.2 and C.2, these shares only rise to 59%, pointing to substantial directional differences between forecasts based on logit or probit and the observable survey responses. As I have already pointed out, an important drawback of relying on estimates based on ordered response models is that these are conditional on the information contained in the right hand side variables. It is very likely however that, due to data limitations, the econometrician's information set is much smaller than the information set actually available to firms when they make forecasts.

Overall, this first exercise shows that quantified forecasts derived from my model are fully consistent with the direction of sales growth implied by the qualitative survey responses. Furthermore, my estimates massively outperform the alternative quantifications based on ordered response models. This is strong evidence for the accuracy of my quantification methodology. I next turn to a Monte Carlo exercise that uses simulated data to infer how precisely my estimates in the artificial data match the magnitude of the artificial forecast errors.

3.4.2 Matching the Magnitude of Forecast Errors — Monte Carlo Simulation

It is important to understand how well forecast errors based on my methodology match, in terms of magnitude, the true quantitative forecast errors. In practice, this is challenging due to the unavailability of data on quantitative firm-level expectations. The vast majority of surveys contain only qualitative questions about firms' future developments. If quantitative survey-based expectations are available at all, then they either focus on aggregate rather than firm-specific variables or have an extremely limited sample size. This is the key motivation for developing the quantification methodology I propose in this chapter, in the first place. To overcome this obstacle, I perform a Monte Carlo exercise that provides a benchmark based on simulated data. In particular, I simulate data on firm (continuous) annual sales growth realizations, as well as corresponding qualitative and quantitative expectations. I then use the data on realized sales growth and qualitative expectations as inputs to the quantification methodology of Section 3.2 and generate estimates for quantified sales growth expectations. Subsequently, I evaluate the accuracy of the estimated forecast errors in the artificial data when compared to those that are now observable in the artificial data.

I generate 1,000 sets of random artificial data, each one of which mimics the structure of the true dataset in terms of number of firms and its unbalanced nature of firm-year-month observations. I provide details about the data generation in Appendix G. This Appendix documents that the underlying processes and their calibration to generate the artificial data are carefully guided by the characteristics and statistics of the observable financial statements and the survey data. I further highlight in this appendix that the simulated datasets match closely moments and statistics in the empirical data that have not been targeted during the calibration.

Table 3.6 shows the distribution of the difference between the true forecast error and the estimated one, both based on the artificial datasets. The mean and median of this distribution are very close to zero — for both moments the difference is roughly only 0.01. This is very small compared to what we observed in Figure 3.4 and Table 3.2. Figure 3.4 shows that the empirical distribution has non-negligible mass at forecast error values as large as ± 0.5. Table 3.2 shows that the median forecast error in my data is as large as -0.03. In general, the distribution of the difference between the artificial true and the artificial estimated forecast error is rather tight. For the 10th (90th) percentile the difference is -0.052 (0.069); and even for the 5th (95th) percentile, it is still reasonably small at -0.071 (0.085). The close correspondence between the estimated and the true forecast error can also be illustrated in a scatter plot. Figure 3.6 contains the scatter plot for one artificial dataset (randomly chosen among the 1,000 draws). The forecast error pairs conform to the 45 degree line (red) quite closely.

The empirical results I shall obtain in Chapter 4 endogenously establish three segments of the distribution of forecast errors that display different statistical properties in terms of their autocorrelation and predictability. These three segments are delimited by the lower and upper 26% of observations in the forecast error distribution. It is

Table 3.6: Distribution of the difference between the estimated quantitative forecast error and the true quantitative forecast error (both based on artificial data)

5%	10%	25%	Median	Mean	75%	90%	95%
-0.071	-0.052	-0.019	0.013	0.011	0.042	0.069	0.085
(0.015)	(0.013)	(0.011)	(0.010)	(0.009)	(0.009)	(0.009)	(0.011)

I report the average across 1,000 random samples of artificial data of the descriptive statistics. Standard deviations across the 1,000 sets for these statistics are reported in parentheses.

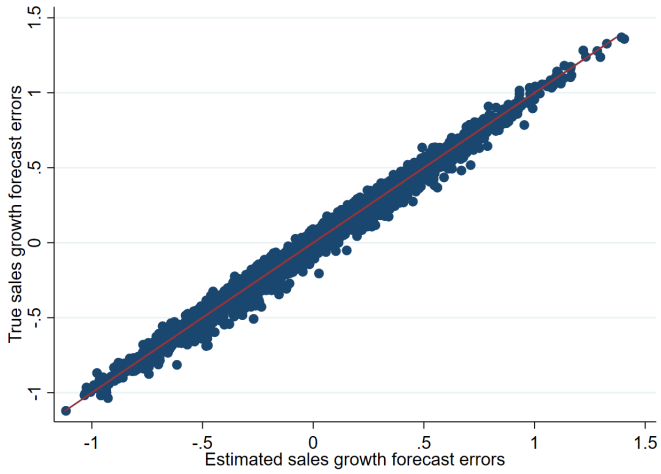

Figure 3.6: **Pairs of true and estimated sales growth forecast errors based on artificial data.** The figure shows all points in the dataset (I randomly selected one of the 1,000 draws for the datasets). The 45° line is shown in red.

therefore important that I also check whether a forecast error based on my methodology is mapped into the same segment of the corresponding 'true' forecast error (both based on artificial data). In the Monte Carlo simulation, I find this is the case for the vast majority; 94% (average across all across all 1,000 draws with 0.007% standard deviation across them) of such pairs.

Overall, I have shown in the first exercise in this section that my estimates for quantified sales growth forecasts are fully consistent with the qualitative information contained in the underlying survey data. In the second exercise, the Monte Carlo simulation has further demonstrated that my quantitative forecast error estimates are highly accurate.

3.5 Concluding Remarks

In this Chapter I provide a quantification model that computes firm-level quantitative forecasts on sales growth from their qualitative counterparts observed in the survey. The availability of quantitative forecasts is key in distinguishing the major forecast errors from the minor ones. A distinction that in my next chapter proves to be important. For the quantification model I develop, I draw upon the 'regression approach' developed by Pesaran (1987), as extended by Smith and McAleer (1995). My contribution is that I redesign the regression method to maintain the panel data structure. This quantification methodology is broadly applicable as long as there is quantitative data on realizations and higher frequency qualitative expectations.

Overall, the quantification comprises of two key equations. The first equation shows how the observed qualitative forecasts map into the overall annual forecast. The second equation in necessary to obtain estimates for the parameters of the first one. Naturally, this is possible because, as I show, the two equations have the same parameters. By using the estimated parameters and the survey-based forecasts, I compute the quantitative firm-year forecasts on sales growth. With the quantitative forecasts, I can now use the realized sales growth rate and compute the ex-post forecast error that I will use in Chapter 4.

I provide evidence of external validity and accuracy for my methodology in a number of ways. First, I show that my quantified forecast estimates are fully consistent in terms of sign with the corresponding qualitative survey-based expectations. Additionally, my methodology also substantially outperforms ordered response models which are potential alternatives for obtaining quantified predictions. Second, I provide evidence for the accuracy of my quantification methodology in terms of the magnitude of firm growth forecasts. I perform a Monte Carlo exercise that provides a benchmark based on simu-

lated data. I find forecast errors based on my methodology using the artificial data are highly accurate when compared with the observable in the artificial data forecast errors.

Having quantitative forecast errors, I now proceed in the next chapter and I show how the major forecast errors behave differently from the minor ones. More specifically, I find robust evidence that only the major ones are predictable and autocorrelated, that is the least accurate ones. I endogenously estimate that the major forecast errors are those that lie at the lower or upper 26% of their empirical distribution. To estimate the point in the distribution where this shift takes place I adapt and implement the Seo and Shin (2016) Dynamic Panel Threshold estimator. To rationalize these empirical findings, I also develop a simple theoretical model of rational inattention. In particular, I show that firms optimally chose to pay less than full attention because attention is costly. Large spikes in the cost of attention push firms to make major forecast errors and to underweight (mistrust) the information they observe. These major forecast errors show predictability and autocorrelation. Low information cost leads to minor forecast errors that are not predictable nor autocorrelated.

Chapter 4

Predictability and Autocorrelation of Forecast Errors

Expectations and assumptions on their formation play a central role to key results in economic models and can be of primary importance for policy makers to understand the implications of their decisions. For example, in Mankiw and Reis (2002), the information stickiness of firms is what enables monetary policy to affect output in the short-run – also known as, non-neutrality of money. Therefore, it is crucial to understand how firms form their expectations. While the literature has provided first guidance, important questions remain open. Are forecast errors always serially autocorrelated? Past realizations can always predict forecast errors? In this chapter I make a first step to answering these questions.

To answer these questions, I will use the quantitative expectations I imputed from the qualitative survey expectations. As a reminder, my data is a combination of firm level survey responses from Greece's Manufacturing Sector matched with the official financial statements of the surveyed firms from 1998-2015. The survey data is monthly and records expectations and realizations that are qualitative; responses indicate positive,

negative or no change. The details of the data are in Chapter 2. The availability of matched financial statement data enabled me to quantify the otherwise qualitative survey responses and use them for my analysis. The details of the quantification are in Chapter 3. As we will see, this quantification is what allows me to study the forecast errors in depth. I examine firm's forecasts and forecast errors of its own future sales growth, because the realized sales are directly observable in the financial statements.

Based on the quantified forecasts, I find robust evidence that only the major forecast errors are predictable and autocorrelated, these are the least accurate ones. I endogenously estimate that the major forecast errors are those that lie at the lower or upper 26% of their empirical distribution. In particular, agreeing with existing empirical findings concerning both firm-level and aggregate variables (see for example see Gennaioli et al. (2016), Massenot and Pettinicchi (2018), Bordalo et al. (2018), Bordalo et al. (2020), Tanaka et al. (2020)), I find that forecast errors of sales growth are serially correlated and that past realizations can predict forecast errors. However, I contribute to the literature by showing that these empirical results are statistically significant only following major forecast errors. In other words, major forecast errors are responsible for the existing empirical findings that lead to the rejection of the rational expectations hypothesis. Finally, I use the qualitative monthly survey responses to show that my empirical results on the autocorrelation and persistence properties of large forecast errors are also a feature in the directly observable survey data.

To identify the threshold above which forecast errors become predictable and autocorrelated I adapt and implement the Seo and Shin (2016) Dynamic Panel Threshold estimator. I adapt this Seo and Shin (2016) estimator by using the Arellano and Bover (1995) Forward Orthogonal Transformations to eliminate the unobserved firm heterogeneity (panel fixed effects) instead of the original Arellano and Bond (1991) First Dif-

ferences. This adaptation is necessary to accommodate the unbalanced nature of my panel data; first differences result in many dropped observations. I endogenously estimate the shift in the behavior of the forecast errors to happen when they take values at the tail of their distribution, at the lower or upper 26%. Clearly, the availability of quantified forecast errors is what enabled me to identify their largest values and their importance in driving the empirical findings. Qualitative survey responses are inherently bounded within a specific interval and cannot identify large positive or negative values. Therefore, they cannot distinguish major values from minor ones.

Following my empirical work, in the spirit of Gabaix (2014), I develop a simple theoretical framework of rational inattention that models the expectation formation process and rationalizes my findings. In particular, I show how optimal forecasts are obtained when the forecasted variable follows an auto-regressive process with costly attention. In this model environment, noisy signals limit the attention of the firms to lagged realizations and make it costly. I show that firms optimally chose to pay less than full attention because being more careful bears higher costs. An increase in the cost of attention pushes firms to make major forecast errors and to underweight (mistrust) the information they observe. Without a positive attention cost, firms accurately observe the lagged values, do not make major forecast errors and form rational and fully informed forecasts. Therefore, only in the presence of major forecast errors that originate from rational inattention, agents violate the full information rational expectations hypothesis. Before unravelling my rational inattention model, I build its micro-foundations using a simple signal-extraction problem similar to that of Lucas (1973).

My work contributes to two strands of the literature. First, my findings contribute to the empirical literature working with surveys to understand the expectation formation. More specifically, I relate to the work of Gennaioli et al. (2016) who study CFO's

forecasts of their firm's future earnings and find that the forecast errors are predictable by the past realizations. However, they do not provide a model explaining this finding nor identify the importance of major forecast errors in driving these results. Similarly, Bordalo et al. (2018) study forecasts of the aggregate credit spread and find that past realizations predict future forecast errors, a behavior resulting from the agents' 'diagnostic expectations'. Bordalo et al. (2020) also find evidence that agents' forecast errors are predictable by past forecast revisions, which also constitutes a violation of the rational expectations hypothesis. I deviate from Bordalo et al. (2018) and Bordalo et al. (2020) by focusing on firm's forecasts errors about its own sales growth and showing that only the major forecast errors result in predictability. Finally, Massenot and Pettinicchi (2018) also study expectations of firm's own 'business conditions' from the IFO survey and find that forecast errors are predictable by past realizations, but provide no model to track the origins of this behavior. Overall, I contribute to these findings by showing that past realizations can predict the forecast errors only following major forecast errors, while I also provide the theoretical background that explains these differences between the major and the minor forecast errors.

In addition to the predictability of forecast errors, there is also abundant evidence that forecast errors are serially autocorrelated. Mankiw et al. (2003) study firms' forecasts of aggregate inflation and Tanaka et al. (2020) study firm's forecasts of aggregate GDP growth. Although, these empirical findings document that forecast errors are serially correlated, they report this finding without attempting to explain its origins. I contribute by showing that this autocorrelation is statistically significant only following major forecast errors, and I provide a model that allows for this behavior.

Second, I contribute to the theoretical work that aims at modelling the expectation formation process and how informed and rational expectations are. My approach shows

that positive attention costs can generate major forecast errors and disorientate agents leading them to inattention. First, I build my framework in the spirit of Gabaix (2014) to study how optimal forecasts are obtained in an environment of costly attention and information. Second, I use it to model and explain the behavior of forecast errors.

The rational inattention model that I propose offers both analytical tractability and comprehensive simplicity that make it preferable to more sophisticated rational inattention models following Sims (2003)'s model (see also Dhami (2016)). A 'sticky information' environment of Mankiw and Reis (2002) would also explain the autocorrelation and the predictability of the forecast errors in my data, but cannot explain why they are the result of major forecast errors only. Another alternative explanation, would be the 'diagnostic expectations' of Bordalo et al. (2018) (a result of the stereotypes of Bordalo et al. (2016)) which allow for autocorrelation and predictability of forecast errors. However, models with 'diagnostic expectations' cannot explain why only the major forecast errors display that behavior. Capistrán and Timmermann (2009) also provide an asymmetrical loss function of forecasters in which the cost of positive forecast errors differs from that of negative ones. This behavior, however, results in positively autocorrelated forecast errors, which is not supported by my data. I argue that the rational inattention is the most suitable explanation for two reasons. First and foremost, it is supported by my empirical findings. Second, the aforementioned alternative models cannot explain why the autocorrelation and the predictability of the forecast errors occur only in the face of major forecast errors.

Overall, we observe that there is ample empirical evidence that forecast errors show autocorrelation and are predictable by past realizations. These two stylized facts are a direct violation of the full information rational expectations hypothesis, which can have immense impact on identifying the effects of adopted policies. I contribute to these

findings in two ways. First, I provide solid evidence that the autocorrelation and the predictability of the forecast errors are driven by major forecast errors only. Without the major forecast errors, the existing empirical findings no longer hold. Second, I contribute by showing how rational inattention can be used in the study of the forecasting behavior of the firms in explaining my findings.

In the following section I provide evidence that forecast errors are predictable by past realizations and that this predictability holds only following major forecast errors. Next, in Section 4.2, I study the negative autocorrelation of the forecast errors which is also entirely attributed to the largest forecast errors. In both of these sections, I endogenously estimate the threshold above which the behavior of the forecast errors changes. Finally, in Section 4.4, I provide my intuitive model of rational inattention that rationalizes these findings.

4.1 Predictability of forecast Errors

In this section, I study the predictability of the forecast errors. I provide evidence that only following major forecast errors, firms make predictable forecast errors. To provide some context on the predictability of forecast errors, assume that a firm-level variable evolves as a first order auto-regressive process, $z_t = \rho z_{t-1}$ (without loss of generality, I omit the error term for simplicity). The firm uses the lagged value to form a forecast on its future evolution, z_t^e. Circumstances such as behavioral biases or noisy signals may affect firms' weight attached to the lagged value, i.e. $z_t^e = \rho\Lambda z_{t-1}$. Here, $\Lambda > 1$ could capture behavioral biases as for example in Bordalo et al. (2018), while $\Lambda < 1$ would be in line with noisy signals as in Gabaix (2014). The forecast error is $z_t - z_t^e = \rho(1-\Lambda)z_{t-1}$. If $1 - \Lambda = 0$, then firms' forecasts correctly extrapolate without any bias, and forecast errors are not predicable from past realizations and are purely random. If $1-\Lambda \neq 0$, this

is a violation of the *efficiency property* (Pesaran (1987)) of the full information rational expectations (FIRE) hypothesis. For $1-\Lambda > 0$, firms' forecasts under-weight the lagged values of the predicted variable, while for $1-\Lambda < 0$ they over-weight them.

4.1.1 Methodology

Given the intuitive examination above, I can define $\varphi \triangleq \rho(1-\Lambda)$, and I can estimate the extrapolation bias in the forecasts of the firms using the following equation

$$x_{iy}^{fe} = \varphi x_{i,y-1} + \Psi_y + \Psi_i + \eta_{iy}, \tag{4.1}$$

where Ψ_i and Ψ_y control for unobserved firm heterogeneity and aggregate annual effects, respectively, and η_{iy} is an idiosyncratic error. If φ is statistically significant, firms' forecasts extrapolate incorrectly. To evaluate the effects of major forecast errors on sales growth, I further estimate the following threshold regression which allows for shifts in the extrapolation bias

$$x_{iy}^{fe} = \varphi_1 x_{i,y-1} * (1 - MAJ_{i,y-1}^q) + \varphi_2 x_{i,y-1} * MAJ_{i,y-1}^q + \varphi_3 MAJ_{i,y-1}^q + \Psi_y + \Psi_i + \eta_{iy}, \tag{4.2}$$

where MAJ_{iy}^q takes the value 1 when there is a major forecast error. A major forecast error occurs when a forecast error lies at either the lower or upper $q\%$ of the distribution. Accordingly, I call all forecast errors in the center of the distribution minor forecast errors. The extrapolation bias following minor forecast errors is φ_1, whereas following a major forecast error, that bias is φ_2. Given the estimated cut-off $q\%$, if $\varphi_1 = 0$ and $\varphi_2 \neq 0$, then forecast errors are predictable only following major forecast errors. φ_3 indicates whether the occurrence of a major forecast error has any effect on the forecast error in the following period.

As far as the estimation is concerned, due to the dynamic nature of equation (4.2), the standard fixed effects estimator (also known as the 'least squares dummy variables' estimator, LSDV) is severely negatively biased when the time dimension of the panel is finite (Pesaran (2015)) as with my data. That is the the Nickell (1981) bias. To see why this is the case for equation (4.2), observe that the major forecast error indicator on the right hand side, $MAJ^{q}_{i,y-1}$, is a function of the lagged dependent variable, $x^{fe}_{i,y-1}$. Moreover, the lagged realization, $x_{i,y-1}$, is also correlated with the lagged forecast error by construction, $x^{fe}_{i,y-1} \triangleq x_{i,y-1} - x^{e}_{i,y-1}$. As a result, to estimate equation (4.2) I need to apply the Dynamic Panel Threshold estimator of Seo and Shin (2016). As Asimakopoulos and Karavias (2016) succinctly note, the original Seo and Shin (2016) estimator consists of the following two steps:

1. For all the values of $q\%$ in a pre-determined interval, I estimate equation (4.2) using the Arellano and Bond (1991) First-Difference GMM (FD) and obtain the value of the objective function of the GMM.

2. The estimated value of $q\%$ is the one that minimizes the objective function of the Two Step Optimal GMM process. Naturally, that is from where I obtain the estimates of φ_1, φ_2 and φ_3 as well.

Although the estimator of Seo and Shin (2016) is straightforward, my application poses some challenges that makes it different from previous applications like those of Asimakopoulos and Karavias (2016) and Polemis and Stengos (2019). The main challenge of the first difference GMM (FD) estimator for dynamic panels is that it drops observations from taking first differences and this problem is more intense in severely unbalanced panels like mine (Roodman (2009), Gorbachev (2011)). For example, if one

observation is preceded and followed by a missing value, then the FD estimator cannot take first differences and completely ignores this observation. To address this problem, Arellano and Bover (1995) develop the forward orthogonal transformation (FOT) which eliminates the unobserved firm-heterogeneity by subtracting from each observation the firm-specific arithmetic mean of its future values, instead of taking first differences. This significantly limits data loss in dynamic panels and allows me to use all available data. This is the only adaptation I make to the original Seo and Shin (2016) estimator. For completeness, I implement the original Seo and Shin (2016) with FD in my robustness checks to re-estimate the threshold and the coefficients of equation (4.2) and show that it yields biased coefficient estimates.

The second challenge I face and is inherent to all the GMM estimators is that the system can be 'excessively' over-identified — moment conditions proliferate. Excessive over-identification leads to biased estimates and to a biased Hansen statistic (for over-identifying restrictions) — see Roodman (2009). Therefore, to limit the over-identification, I collapse the instruments as Roodman (2009) suggests (see also Gorbachev (2011) and Caselli and Tesei (2016)). Third, standard errors might also be downwards biased because I have a large number of moment conditions compared to the number of firms, so I use the Windmeijer (2005) corrected standard errors (see Roodman (2009), Gorbachev (2011) and Caselli and Tesei (2016)), as well as the two-step GMM to improve efficiency. In order to further limit the risks of over-identification I do not use all available lags as instruments. For my baseline results I restrict the lag length of instruments to 5. Estimation results are robust to using fewer lags as instruments, and I show this in the robustness checks.

Overall, my application of the Dynamic Panel Threshold estimator of Seo and Shin (2016) is as follows:

1. For my baseline results I use the Arellano and Bover (1995) FOT Two-Step GMM estimator with collapsed instruments and Windmeijer (2005) corrected standard errors. As instruments, I use the lagged values of the right hand side variables with lag length of 5 lags (lags dated from $y-2$ to $y-6$). With my FOT estimator, my instruments and the chosen number of lags, I estimate the threshold regression (equation (4.2)) for all the values of $q\%$ in the pre-determined interval $q\% = 6\%, 7\%, 8\%...45\%$ to obtain the value of the objective function of the estimator.[1]

2. I obtain the estimate of $q\%$ as well as of the rest of the coefficients when the GMM objective function is minimized.

Finally, to maintain comparability with the existing literature findings, I also estimate the dynamic equation (4.1). I do so by using the Arellano and Bover (1995) FOT Two-Step GMM estimator with collapsed instruments, Windmeijer (2005) corrected standard errors and instruments lagged from $y-2$ to $y-6$.

4.1.2 Baseline Empirical Findings

In Table 4.1, I present the results from the estimation of equations (4.1) and (4.2) using the quantified forecast errors. In column (1), I report the estimation results from the simple linear equation (4.1). In line with the literature (Gennaioli et al. (2016), Massenot and Pettinicchi (2018) and Bordalo et al. (2018)) I also find that the forecast errors of the firms are predictable by past realizations.[2] In other words, firms' forecasts violate the FIRE hypothesis.

[1] There is no specific guidance in the literature on the choice of interval, but it will become apparent below that my estimates turn out to be well in the middle of the interval. I remain agnostic and specify a fairly wide interval that covers up to 80% of all observations.

[2] To be precise, the existing estimates in the literature use the biased LSDV, which I also show in the Appendix Section H, in Table 1.H. These estimates also indicate that forecast errors are predictable by past realizations.

In column (2) of Table 4.1, I report the estimates of the non-linear threshold equation (4.2). We observe that the coefficient of the lagged realization is statistically significant only following a major forecast error. That is, when it interacts with the indicator $MAJ^q_{i,y-1}$. In other words, firms form biased predictions of the sales growth only following major forecast errors. When firms make more accurate forecasts, they remain more aligned with the FIRE hypothesis. Finally, the estimated cut-off value is 26%, which means that major forecast errors are estimated to occur when their values lie at the upper or at the lowest 26% of their distribution. In both columns (1) and (2) of Table Table 4.1 the Hansen p-values are remarkably high (0.782 and 0.982, respectively), which can be an alarming indication of 'excessive over-identification' as Roodman (2009) notes. However, I have taken all necessary steps to avoid this problem as I explained earlier. In addition, I show in Appendix H that my estimates are fully robust to using a GMM system without any over-identifying restrictions. These demonstrate that my baseline estimates do not suffer from 'excessive over-identification'.

Most importantly, the estimated threshold of 26% suggests that the forecast errors resulting to predictability are sizeable. Forecast errors at the lower 26% of the distribution take values lower or equal to -0.143 which means that sales growth was expected to be 14.3 percentage points higher than the subsequent realization. Similarly, those at the upper 26% take values larger that 0.086, which means that growth was expected to be 8.6 percentage points lower than the subsequent realization.[3]

Finally, for the non-linear model (threshold) both Hansen p value and Arellano-Bond test of autocorrelation of order 2 (m2 test) strongly reject the null that my specification

[3] I attempted to estimate the threshold equation with asymmetric thresholds for upper and lower cut-off % (output omitted). The resulting cut-off values and coefficient estimates were not robust to using different lag lengths. This is to be expected as by introducing a further non-linearity in the only one variable that I have on the right-hand side, I sacrifice efficiency and accuracy. To estimate this model one needs more observations and particularly a much larger cross-sectional dimension. Moreover, with the added non-linearity the instruments can be very weakly correlated with the right-hand side variables.

Table 4.1: Predictability of firms' forecast errors of sales growth.

	(1)	(2)
Estimation	FOT	FOT
Stand. Errors	2-step, Windmeijer corrected	2-step, Windmeijer corrected
Lags as Instruments	2-6	2-6
Estimated Threshold q	N.A.	26%
	Dependent Variable: Sales Growth Forecast Error, x_{iy}^{fe}	
$x_{i,y-1}$	-0.161***	–
$x_{i,y-1} * (1 - MAJ_{i,y-1}^{q})$	–	-0.0583
$x_{i,y-1} * MAJ_{i,y-1}^{q}$	–	-0.146**
$MAJ_{i,y-1}^{q}$	–	0.0164
$\overline{x}_{IND,y}$	0.837***	0.794***
Observations	2,805	2,069
# of Firms	590	432
Over-identified	Yes	Yes
Hansen p-value	0.782	0.982
m2 test p-value	0.857	0.936

Column (1) shows estimates of equation (4.1) without the threshold. Column (2) is the Dynamic Panel Threshold estimator of Seo and Shin (2016) using the Arellano and Bover (1995) FOT Two-Step GMM for equation (4.2). Instruments are collapsed in both specifications. Instruments are with lags dated from $y-2$ to $y-6$. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate annual effects with the NACE two-digit industry, IND, year average of sales growth from the entire sample of the financial statements, $\overline{x}_{IND,y}$. x_{iy}^{fe} is the forecast error of sales growth for year y; $x_{i,y-1}$ is lagged realized sales growth. MAJ_{iy} takes value one when the forecast error lies at the lower or upper $q = 26\%$ of its empirical pool distribution. *** and ** indicate statistical significance at the 1% and 5% level, respectively.

is weak. This indicates that my model specification is valid so the non-linearity indeed exists.[4]

4.1.3 Robustness Checks

In addition to my baseline results, I run some robustness checks to solidify my findings in Appendix H that I discuss here. I confirm that my baseline estimates are robust to using shorter lag length as instruments. I also justify my choice of using the FOT GMM by showing that the original estimator remains biased in my data.

[4]Seo and Shin (2016) propose a bootstrap procedure as a test for the non-linearity suitable for balanced panels. However, it is not clear how this testing algorithm can be extended to unbalanced panels.

First, in Table 1.H, I show alternative estimations of the predictability equation without the thresholds; i.e. linear equation (4.1). In Column (1), the GMM system is just identified without over-identifying restrictions and the predictability coefficient is -0.158 which is remarkably close to the baseline estimate of -0.161 in Table 4.1. In column (2) I estimated equation (4.1) using the biased LSDV, following the standard practice of the literature. We observe that forecast errors are predictable, but the LSDV estimates suffer from the Dynamic Panel bias. The LSDV estimate of the predictability coefficient is -0.221 which is much lower than the GMM estimates. Evidently, using the GMM FOT corrects this bias.

Second, in Table 2.H, I show alternative estimations of the nonlinear equation (4.2) with the threshold. In column (1), I estimate the threshold regression with the biased LSDV classifying the major forecast errors based on the baseline threshold estimate of 26%. The LSDV coefficients are much lower than the baseline estimates; -0.236 in Table 2.H versus my baseline estimate of -0.146 in Table 4.1. This demonstrates that my choice of estimator corrected the bias. In column (2) of Table 2.H I use my baseline estimation strategy but with fewer lags, that is without any over-identifying restrictions. There, the estimated coefficients are close to my baseline estimates; -0.157 versus the baseline estimate of -0.146. The estimated threshold cut-off is not much larger either; 28% versus 26% in my baseline results. Column (3) is with the original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$. The estimated threshold is 27% which is close to the baseline estimate of 26%. However, the coefficient is biased and is close to the LSDV biased one of column (4): -0.198 and -0.235, respectively.

Overall, I find that the extrapolation bias becomes non-zero following only a major forecast errors. I endogenously estimate that major forecast errors are those that lie at

the upper or lower 26% of the distribution. These results are robust to using shorter lag length as instruments. Moreover, my threshold estimates are also robust to using the original Seo and Shin (2016) threshold estimator with FD GMM. Finally, I provide evidence that using the original estimator with my data does not sufficiently correct the Nickell (1981) bias in the estimated coefficients.

4.2 Autocorrelation of forecast Errors

In addition to the predictability of forecast errors, there is empirical evidence in the literature that forecast errors are autocorrelated. Under the full information rational expectations, forecast errors should be neither predictable by past realizations nor autocorrelated. In this section I study their autocorrelation and I show that these findings hold only following major forecast errors.

4.2.1 Methodology

To examine the autocorrelation of the forecast errors, I need to estimate the following equation

$$x_{iy}^{fe} = \varrho x_{i,y-1}^{fe} + \Psi_y + \Psi_i + \eta_{iy}, \tag{4.3}$$

where ϱ is the autocorrelation coefficient, Ψ_i and Ψ_y control for unobserved firm heterogeneity and year fixed effects, and η_{iy} is an idiosyncratic error.

As with the predictability of the forecast errors, I want to evaluate whether the size of forecast errors matters for their autocorrelation. To allow for asymmetries in the autocorrelation coefficient I additionally estimate the following threshold regression

$$x_{iy}^{fe} = \varrho_1 x_{i,y-1}^{fe} * (1 - MAJ_{i,y-1}^{q}) + \varrho_2 x_{i,y-1}^{fe} * MAJ_{i,y-1}^{q} + \varrho_3 MAJ_{i,y-1}^{q} + \Psi_y + \Psi_i + \eta_{iy}, \tag{4.4}$$

where MAJ_{iy}^{q} again is a dummy variable that takes the value one when there is a major forecast error. A major forecast error is defined as a forecast error in the top and bottom $q\%$ of the distribution. The persistence following minor forecast errors is given by ϱ_1, while following a major forecast error, forecast errors are autocorrelated with coefficient ϱ_2. If only ϱ_2 is statistically significant for the estimated threshold $q\%$, then forecast errors show persistence only following a major forecast error.

To estimate equation (4.4) I adapt and apply the Dynamic Panel Threshold estimator of Seo and Shin (2016). In Section 4.1, I describe in detail the challenges that I face with the application of the Seo and Shin (2016) threshold estimator for the predictability of the forecast errors. It is evident that for the autocorrelation I face the same challenges. Therefore, based on the analysis in Section 4.1, I apply the Seo and Shin (2016) threshold estimator for the autocorrelation as follows:

1. For my baseline results I use the Arellano and Bover (1995) FOT Two-Step GMM estimator with collapsed instruments and Windmeijer (2005) corrected standard errors. As instruments, I use the lagged values of the independent variables, and I use lags dated $y-2$ to $y-6$. With my FOT estimator, my instruments and the chosen number of lags, I estimate the threshold regression (equation (4.4)) for all the values of $q\%$ in the pre-determined interval $q\% = 5\%, 6\%, 7\%...45\%$.

2. I obtain the estimate of $q\%$ as well as of the rest of the coefficients when the GMM objective function is minimized.

Additionally, I estimate the dynamic equation (4.3). I do so by using the Arellano and Bover (1995) FOT Two-Step GMM estimator with collapsed instruments, Windmeijer (2005) corrected standard errors and instruments lagged from $y-2$ to $y-6$.

4.2.2 Baseline Empirical Findings

In Table 4.2, I show the results of the estimation of equations (4.3) and (4.4). First, we observe that sales growth forecast errors are serially negatively autocorrelated (column (1)). Clearly, this violates the FIRE hypothesis. Second, from column (2) in Table 4.2, is that the negative autocorrelation of the forecast errors is only the result of major forecast errors. The autocorrelation coefficient is statistically significant only when lagged values of forecast errors are interacted with the major forecast error indicator. In addition, the estimated threshold cut-off is 26% which is also the estimated value of the cut-off in the predictability of forecast errors (Table 4.1).[5] Note that I independently estimate the two equations, and the two estimates coinciding adds to the robustness of my results. Finally, for the non-linear model (threshold) both Hansen p value and Arellano-Bond test of autocorrelation of order 2 (m2 test) strongly reject the null that the specification is weak. Similarly to the predictability equation, this indicates that my model specification is valid; i.e. the non-linearity indeed exists.

As with the estimates for the predictability equations, the estimates in Table 4.2 also have remarkably high Hansen p-values. Roodman (2009) suggests that this can be an indication of excessive over-identification. However, I have taken all necessary steps to avoid this problem as I explained earlier. Most importantly, I show in Appendix I that my estimates are fully robust to using the GMM FOT without any over-identifying restrictions. These demonstrate that my baseline estimates do not suffer from 'excessive over-identification'.

[5]For reference, forecast errors lower than -0.143 are at the lower 26%, while those that are larger than 0.086 are at the upper 26%.

Table 4.2: Autocorrelation of firms' forecast errors on sales growth.

	(1)	(2)
Estimation	FOT	FOT
Stand. Errors	2-step, Windmeijer corrected	2-step, Windmeijer corrected
Lags as Instruments	2-6	2-6
Estimated Threshold q	N.A.	26%
	Dependent Variable: Sales Growth Forecast Error, x_{iy}^{fe}	
$x_{i,y-1}^{fe}$	-0.164***	–
$x_{i,y-1}^{fe} * (1 - MAJ_{i,y-1}^{q})$	–	0.213
$x_{i,y-1}^{fe} * MAJ_{i,y-1}^{q}$	–	-0.167**
$MAJ_{i,y-1}^{q}$	–	0.0305
$\overline{x}_{IND,y}$	0.811***	0.797***
Observations	2,069	2,069
# of Firms	432	432
Over-identified	Yes	Yes
Hansen p-value	0.935	0.99
m2 test p-value	0.892	0.936

Column (1) shows estimates from equation (4.3) without the threshold. Column (2) is the Dynamic Panel Threshold estimator of Seo and Shin (2016) using the Arellano and Bover (1995) FOT Two-Step GMM for equation (4.4). In both specifications instruments are collapsed. In columns (1) and (2) lags dated from $y-2$ to $y-6$. The Arellano-Bond p-value (m2 test) shows no autocorrelation of order two in the errors. I proxy the aggregate annual effects with the NACE two-digit industry, IND, year average of sales growth from the entire sample of the financial statements, $\overline{x}_{IND,y}$. x_{iy}^{fe} is the forecast error of sales growth for year y. MAJ_{iy} takes value one when the forecast error lies at the lower or upper $q = 26\%$ of its empirical pool distribution. *** and ** indicate statistical significance at the 1% and 5% level, respectively.

4.2.3 Robustness Checks

As with the predictability equation, I run some robustness checks to solidify my baseline findings in Appendix I that I discuss here. I confirm that my baseline estimates are robust to using shorter lag length as instruments. I also show that the original Seo and Shin (2016) estimator delivers biased coefficient estimates in my data, which justifies my choice of using the FOT GMM.

First, in Table 1.I, I show alternative estimations of the predictability equation without the thresholds; i.e. linear equation (4.3). In Column (1), the GMM system is just identified without over-identifying restrictions and the autocorrelation coefficient is

-0.163 which is remarkably close to the baseline estimate of -0.164 in Table 4.2. In column (2) I estimated equation (4.3) using the biased LSDV, following the standard practice of the literature. We observe that forecast errors are predictable, but the LSDV estimates suffer from the Dynamic Panel bias. The LSDV estimate of the autocorrelation coefficient is -0.238 which is much lower than the GMM estimates. Evidently, using the GMM FOT corrects this bias.

Second, in Table 2.I, I show alternative estimations of the nonlinear equation (4.4) with the threshold. In column (1), I estimate the threshold regression with the biased LSDV classifying the major forecast errors based on the baseline threshold estimate of 26%. The LSDV coefficients are much lower than the baseline estimates; -0.24 in Table 2.I versus my baseline estimate of -0.167 in Table 4.2. This demonstrates that my choice of estimator corrected the bias. In column (2) of Table 2.I I use my baseline estimation strategy but with fewer lags, that is without any over-identifying restrictions. There, the estimated coefficients are close to my baseline estimates; -0.168 versus the baseline estimate of -0.167. The estimated threshold cut-off is not much different either; 20% versus 26% in my baseline results. Column (3) is with the original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$. The estimated threshold is 26% which is the same as the baseline estimate. However, the coefficient is biased and is close to the LSDV biased one of column (4): -0.217 and -0.240, respectively.

Overall, I find that the autocorrelation coefficient becomes non-zero following only major forecast errors. I endogenously estimate that major forecast errors are those that lie at the upper or lower 26% of the distribution. Most importantly, the threshold estimate is independently estimated and is the same between the two equations, the predictability and the autocorrelation. This adds to the robustness. My empirical results

are also robust to using shorter lag length as instruments. Moreover, my threshold estimates are robust to using the original Seo and Shin (2016) threshold estimator with FD GMM. Finally, I provide evidence that using the original estimator with my data does not sufficiently correct the Nickell (1981) bias in the coefficients.

Given that my results have so far been based on the computed forecast errors, In the following section I give evidence that this behavior is also present in the directly observable survey data as well. Subsequently, I proceed with a simple model of rational inattention that rationalizes my empirical findings.

4.3 Predictability and Persistence in Survey-Based Forecast Errors

One may wonder whether the empirical findings I obtained using my quantification model are also a feature of the qualitative survey data. Below, I provide some evidence that this is the case. I construct monthly forecast errors from qualitative expectations data in a manner employed by Bachmann et al. (2013) and Massenot and Pettinicchi (2018). I then document that the probability of large survey-based forecast errors is affected by past realizations and past large forecast errors only in years that my quantification methodology flags as involving a major forecast error. This means that in the survey data too the FIRE hypothesis is violated only when my findings suggest there is a major forecast error. Since the survey responses are qualitative, binary choice models permit me to identify patterns of predictability and autocorrelation similar to those of their annual quantified counterparts.

In the monthly surveys, by subtracting the expectational responses from the corresponding realizations, I construct monthly forecast errors with values $XS_{im}^{fe} = \{-2, -1,$

$0, +1, +2\}$. Compare the survey questions A.2 and D.2 in Section 3.1 where *increased/ unchanged/ decreased* responses are labelled as -1/0/+1. Essentially, these responses indicate the direction of change of sales, expected and realized. For the outcomes labelled -2 and $+2$, the firm's forecast completely mis-predicted the direction of change of its sales and those errors are likely large. I label them XS_{im}^{Lfe} and analyze below whether their occurrence is predictable or shows persistence.

I identify years in which firms made a major forecast error from the annual quantified values based on the threshold regressions contained in Tables 4.1 and 4.2. The months that belong to these years are assigned $MAJ_{iy} = 1$. I test for the predictability using the following probit model which is analogous to the continuous regression (4.1)

$$\begin{aligned} \mathbb{P}\Big\{XS_{im}^{Lfe} = 1 \big| XS_{im}, MAJ_{iy}\Big\} &= G\Big(\varphi_1' XS_{im} * (1 - MAJ_{iy}) + \varphi_2' XS_{im} * MAJ_{iy} \\ &\quad + \varphi_3' MAJ_{iy} + \Psi_y' + \Psi_i'\Big) \end{aligned} \tag{4.5}$$

where Ψ_i' and Ψ_y' control for firm and year fixed effects, MAJ_{iy} takes the value 1 when there is a major forecast error (lower or upper 26% of the annual quantified forecast error distribution). $\mathbb{P}\{\cdot|\cdot\}$ is the conditional probability and $G()$ is the standard normal distribution which gives the probit model.

Even though equations (4.1) and (4.5) are analogous, the interpretation of the coefficients is vastly different. In (4.5), φ_1' measures how the probability of a large survey-based forecast error is affected by past survey-based realizations in a year when a major forecast error took place. A non zero value of φ_1' suggest predictability whether this results in a decrease or an increase of the probability of the forecast error. As a result, the sign of φ_1' is not comparable to the sign of the linear coefficient φ_1 in the threshold equation (4.1). Similarly, a non-zero φ_2' suggests predictability only during a year with

a minor forecast error.

Analogous to the predictability test I have the following dynamic probit for the persistence of the survey forecast errors, analogous to the continuous version, equation (4.4),

$$\mathbb{P}\left\{XS_{im}^{Lfe} = \pm 2 \middle| XS_{i,m-3}^{Lfe}, MAJ_{iy}\right\} = G\Big(\rho_1' XS_{i,m-3}^{Lfe} * (1 - MAJ_{iy}) + \rho_2' XS_{i,m-3}^{Lfe} * MAJ_{iy} + \rho_3' MAJ_{iy} + \Psi_y' + \Psi_i'\Big). \tag{4.6}$$

As before, a non-zero ρ_1' suggests persistence of large survey-based forecast errors only during a major forecast error. While, a non-zero ρ_2' suggests persistence only during minor forecast errors. To reiterate the sign of the coefficients is not comparable to the linear coefficients in equation (4.4).

To proxy for the firm fixed effects I follow Wooldridge (2010). In equation (4.5), I use the firm specific cross-time average of the right hand side variables XS_{im} and MAJ_{iy}, and in (4.6), the firm-specific average of MAJ_{iy}. To address the initial conditions problem I also include the observation of the dependent variable for each firm, $XS_{i,0}^{Lfe}$, on the right hand side (see Wooldridge (2010)).

Table 4.3 shows that only during years when annual quantified forecast errors are classified as major (i.e. $MAJ_{iy} = 1$) are survey-based forecast errors predictable and show persistence. This means that they affect the probability of future large survey-based forecast errors, which is a violation of the FIRE hypothesis. The negative sign and the statistical significance of the coefficient of $XS_{im} * MAJ_{iy}$ in Panel A of Table 4.3 suggests that a positive growth in the past-three months sales ($XS_{im} = 1$) will reduce the probability of a large survey-based forecast error concerning the following three month period. Accordingly, a negative growth rate in the past-three months ($XS_{im} = -1$)

Table 4.3: Predictability and Persistence of firms' forecast errors of sales growth in the qualitative survey data. Probit Estimates.

Panel A: Predictability		Panel B: Persistence	
$XS_{im} * (1 - MAJ_{iy})$	-0.0398	$XS^{Lfe}_{i,m-3} * (1 - MAJ_{iy})$	-0.0554
$XS_{im} * MAJ_{iy}$	-0.133***	$XS^{Lfe}_{i,m-3} * MAJ_{iy}$	0.326***
MAJ_{iy}	0.407*	MAJ_{iy}	0.0543
$XS^{Lfe}_{i,0}$	0.918***	$XS^{Lfe}_{i,0}$	0.422*
$\overline{XS}_i$	0.0177		–
$\overline{FEL}_i$	-0.207	$\overline{FEL}_i$	-0.231
Constant	-1.527***	Constant	-1.548***
Observations	8,659	Observations	5,592
Number of firms	411	Number of firms	328

Probit estimation of the conditional probability of large survey-based forecast error of sales growth, $\mathbb{P}\{XS^{Lfe}_{im} = 1\}$. The definition of XS^{Lfe}_{im} is in the main text. Panel A shows estimates of predictability and Panel B of persistence. XS_{im} is the survey-based realization. MAJ_{iy} takes value one when the annual quantified forecast error lies in the lower or upper 26% of its empirical pool distribution. $XS^{Lfe}_{i,0}$ is the first observed survey forecast error of firm i and addresses the initial conditions problem (see Wooldridge (2010)). $\overline{XS}_i$ and $\overline{FEL}_i$ are firm-specific cross-time averages and control for the firm fixed effects. Fixed year effects are also included. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level, respectively.

will increase the probability. In both cases, this effect violates the FIRE hypothesis. In Panel B, the coefficient of $XS^{Lfe}_{i,m-3} * MAJ_{iy}$ suggests that in years with major forecast errors, past large survey-based forecast errors increase the probability of future large survey-based forecast errors. Again, this violates the FIRE hypothesis.

The coefficients of past realizations (in Panel A) and past forecast errors (in Panel B) in years with minor forecast errors (i.e. $(1 - MAJ_{iy}) = 1$) are insignificant. This means that they do not affect the probability of future large survey-based forecast errors, which is not a violation of the FIRE hypothesis. These results based on the monthly qualitative survey forecasts are consistent with the ones I obtained from the annual quantified forecasts.

4.4 Model of a Firm with Rational Inattention

In this section, I outline a simple framework in which forecast errors result from the fact that a firm cannot perfectly observe its current sales growth, but has to solve a signal-extraction problem. This framework is subsequently extended, in the spirit of Gabaix (2014), to endogenize the firm's choice on signal precision. The firm can choose its degree of attention to information which potentially comes at a cost. I subsequently show that a simple model with limited attention to information and variations in the cost for attentiveness can rationalize the empirical findings of Sections 4.1 and 4.2.

4.4.1 Forecasts in a Simple Signal-Extraction Framework

I first use a simple signal-extraction model of a firm to provide micro-foundations behind a framework with behavioral inattention. A firm i cannot observe its current sales growth x_y, but only a noisy signal $s_y = x_y + \epsilon_y$, where the noise term is i.i.d. with $\mathbb{E}\epsilon_y = 0$, $\mathbb{E}\epsilon_y^2 = \sigma_\epsilon^2$ and $\mathbb{E}x_y\epsilon_y = 0$, for all years y. I abstain from a subscript i for the remainder of this section to ease notation. I show in Appendix Section J.1 that realized sales growth follows an AR(1) process. That is,

$$x_y = \rho_0 + \rho x_{y-1} + u_y, \tag{4.7}$$

with i.i.d. shocks $u_y \sim N\big(0, \sigma_u^2\big)$. It follows that the mean of x_y is $\mu \triangleq \mathbb{E}[x_y] = \rho_0/(1-\rho)$, and that its variance is $\sigma_x^2 \triangleq V[x_y] = \sigma_u^2/(1-\rho)$. Without loss of generality, I assume henceforth that $\mu = 0$. Finally, I assume that the shock, u_y, and the noise term, ϵ_y, are independent.

At time y the firm wants to obtain a one period ahead forecast, $\tilde{x}_{y+1}$, that minimizes the expected squared forecast error, but its information set only includes the most recent

noisy signal s_y and not the true value x_y. This assumption about the information set is consistent with managerial practice. When, towards the end of the (financial) year, forecasts are made about next year's sales, the financial statements are not yet finalized so that managers have to rely on intermediate reports or not yet fully compiled information which only provide an imperfect signal. It is important to understand that firms can observe the true value of sales growth but with delay, so this is not a Kalman filter problem. Kalman filter is applicable when the true value is never observable and only the signal can be observed.

Based on these assumptions, the optimal forecast, x^e_{y+1}, is[6]

$$x^e_{y+1} = \arg\min_{\tilde{x}_{y+1}} \mathbb{E}\left[\frac{1}{2}\big(\tilde{x}_{y+1} - x_{y+1}\big)^2 | s_y\right].$$

The first order condition yields $x^e_{y+1} = \mathbb{E}\big[x_{y+1}|s_y\big]$ and using the fact that x_{y+1} follows the AR(1) process (4.7), I obtain

$$x^e_{y+1} = \rho\mathbb{E}\big[x_y|s_y\big] + \mathbb{E}\big[u_{y+1}|s_y\big],$$

where $\mathbb{E}\big[u_{y+1}|s_y\big] = \mathbb{E}\big[u_{y+1}|x_y + \epsilon_y\big] = 0$, because u_y and ϵ_y are independent and $\mathbb{E}\big[u_{y+1}|x_y\big] = 0$. In line with Gabaix (2014), and given the linear process for the signal and normally distributed errors, Bayesian updating implies the following linear decomposition of the conditional expectation $\mathbb{E}\big[x_y|s_y\big]$,

$$x^e_{y+1} = \rho\mathbb{E}\big[x_y|s_y\big] = \rho\lambda_0 + \lambda\rho s_y, \qquad \text{where} \qquad \lambda = \frac{Cov(x_y, s_y)}{V(s_y)}, \quad \text{and} \quad \lambda_0 = (1-\lambda)\mu = 0. \tag{4.8}$$

Since I further know that $Cov(x_y, s_y) = \mathbb{E}\big[x_y s_y\big] = \mathbb{E}\big[x_y(x_y + \epsilon_y)\big] = \mathbb{E}\big[x_y^2\big] = \sigma_x^2$, and

[6]Minimizing the quadratic forecast error implies that on average firm's predictions will be correct, i.e. the mean forecast error will be zero.

that $V(s_y) = \mathbb{E}\left[s_y^2\right] = \sigma_x^2 + \sigma_\epsilon^2$ due to the independence of x_y and ϵ_y, it follows that

$$\lambda = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\epsilon^2}. \tag{4.9}$$

Equation (4.9) shows that in the presence of noise, $\sigma_\epsilon^2 > 0$, λ is strictly between 0 and 1. This has implications on the optimal forecast (4.8), which is, when applying the definition for the signal,

$$x_{y+1}^e = \lambda \rho s_y = \lambda \rho x_y + \lambda \rho \epsilon_y. \tag{4.10}$$

Equation (4.10) links the firm's optimal sales growth forecast with the current value of sales growth. It shows that if the signal is contaminated by noise the optimal forecast understates the persistence of sales growth, since $0 < \lambda < 1$. Under perfect information (in the absence of noise $\sigma_\epsilon^2 = 0$), $\lambda = 1$ and equation (4.10) becomes the full information optimal forecast.

Another interpretation of the discussed simple setup with a noisy signal is provided by the literature on rational inattention: the firm can potentially perfectly observe all information on current sales growth, but it would choose not to pay attention to all information when making a forecast, e.g. because information processing is costly. The degree of limited attention to information is captured in an abstract way by the noise. In this section the noise variance, and hence the degree of attention, was given exogenously. In the next section, I will endogenize this choice. Then the firm can choose its level of attention to information by determining the parameter λ; which is equivalent to choosing the information content in the signal by varying the noise variance σ_ϵ^2. If the firm pays attention to all information the noise variance equals zero and $\lambda = 1$. For a positive noise variance, attention to information is limited and $0 < \lambda < 1$. I will develop in the next subsection the simple signal-extraction framework into a model with rational

inattention in which the firm can endogenously determine the level of attention, λ.

4.4.2 Introducing Rational Inattention

While the firm's level of attention to information was determined exogenously in the above signal extraction framework, it will now be endogenized. Based on the discussion in the previous section the firm applies the following utility metric to make an optimal forecast

$$W\left(\tilde{x}_{y+1}, \lambda s_y\right) \triangleq -\frac{1}{2}\left(\tilde{x}_{y+1} - \rho\lambda s_y\right)^2,$$

where the parameter λ determines the firm's degree of attention to the observed signal about current sales growth. The forecast is based on the information set at time y and, as in the model in the previous section, the firm does not observe current sales growth but makes decisions based on the noisy signal — as before this signal is the sum of actual sales growth and the noise. The full information optimal forecast ($\lambda = 1$) would be $x^e_{y+1} = \rho x_y$ which is in line with the underlying AR(1) process (4.7) for sales growth. For $0 < \lambda < 1$ the firm pays limited attention to the signal and for $\lambda = 1$ the firm pays full attention to all information. I define the value of $\tilde{x}_{y+1}$ that maximizes firm's utility as

$$x^e_{y+1}(\lambda) \triangleq \arg\max_{\tilde{x}_{y+1}} W(\tilde{x}_{y+1}, \lambda s_y),$$

which is now a function of the attention parameter λ. If I substitute the optimal forecast, $x^e_{y+1}(\lambda)$, into the utility function, I obtain the indirect utility function

$$U(\lambda) = W\left(x^e_{y+1}(\lambda), \lambda s_y\right), \tag{4.11}$$

which transforms the firm's problem to one that requires the choice of the attention parameter λ. Increasing the precision of the signal through information accumulation, reflected in the choice of λ, potentially comes at a cost. I assume the cost function

$$C(\lambda, c_y) = c_y K(\lambda), \tag{4.12}$$

where $K(\lambda)$ is a continuous increasing and convex function in λ. Note that this function depends on the cost shock c_y, which is assumed to be independently and identically distributed across time and is bounded between zero and a positive upper bound.[7] The firm observes the cost shock at the beginning of the period prior to any choice on the level of attention.

Given the above assumptions, the firm first chooses an optimal level of attention λ^*, and conditional on this choice, it obtains in a second step the optimal forecast for sales growth.[8] I will look at these two steps in turn. First, the firm's objective is to choose the attention parameter so that it maximizes the difference between the expected indirect utility (4.11) and the cost function (4.12). This can be formalized as

$$\max_{\lambda} \Big[\mathbb{E}U(\lambda) - C(\lambda, c_y) \Big].$$

One can show (detailed steps are provided in Appendix K.1) that the firm obtains the optimal level of attention, λ^* by solving the following intratemporal problem

$$\lambda_y^* \triangleq \arg\max_{\lambda} \Big[-\frac{1}{2}\sigma_s^2(1-\lambda)^2 - c_y K(\lambda) \Big], \tag{4.13}$$

[7] I make minimal assumptions about the stochastic process for c_y. The actual choice of the upper bound may depend on the functional form of $K(\lambda)$ as can be seen from equations (4.14) below. The only requirement on the positive upper bound on c_y is that it is specified to ensure that $\lambda > 0$.

[8] The reason why I can write this as a two-step approach is that in the first step the decision is independent of sales growth, x_y.

where σ_s^2 denotes the variance of the signal. It becomes apparent now that, given the time varying cost c_y, also the optimal level of attention fluctuates over time. The first order condition is then

$$\sigma_s^2(1-\lambda_y^*) - c_y K'(\lambda_y^*) = 0,$$

where $K'(\cdot)$ denotes the first derivative of $K(\cdot)$.

The results that follow in this section below do not require a particular functional form for $K(\cdot)$. In fact, my assumptions on $K(\lambda)$ are consistent with several specific functional forms used in the literature. For example, $K(\lambda) = \frac{1}{2}log_2((1-\lambda)^{-1})$ that was introduced in the seminal work of Sims (2003) and is based on the Shannon entropy. I refrain from making any more assumptions about the cost function in order to maintain its simplicity. However, to briefly provide intuition about how the optimal level of inattention depends on the information cost and the variance of the signal, I specify $K(\lambda) = \lambda^a$ where $a \geq 1$ as in Gabaix (2014). Then the first order condition has, for the cases $a = 1$ and $a = 2$, the following simple analytical solutions

$$\begin{aligned} \lambda_y^* &= \frac{\sigma_s^2 - c_y}{\sigma_s^2}, \qquad \text{for} \quad a = 1. \\ \lambda_y^* &= \frac{\sigma_s^2}{\sigma_s^2 + 2c_y}, \qquad \text{for} \quad a = 2. \end{aligned} \tag{4.14}$$

These two parameterizations exemplify that the optimal level of attention is negatively related to the cost shock, c_y. In other words, everything else equal, the firm reduces the level of attention in light of an increase in the cost of information acquisition. In general, for $c_y = 0$ there is no cost for information and $\lambda_y^* = 1$. Given the parameterization for a, I assumed an upper bound for c_y that guarantees $0 < \lambda_y^* < 1$.

Having chosen the optimal level of attention via (4.13), the firm's optimal forecast

is given by

$$x^e_{y+1} \triangleq \arg\max_{\tilde{x}_{y+1}} \left[-\frac{1}{2}(\tilde{x}_{y+1} - \lambda^*_y \rho s_y)^2 \right],$$

so that the optimal forecast is

$$x^e_{y+1} = \lambda^*_y \rho s_y. \tag{4.15}$$

As in the simple signal extraction problem above, the forecast understates the persistence of the sales growth in the case of imperfect information ($0 < \lambda^*_y < 1$). If $\lambda^*_y = 1$ the firm makes the full information rational forecast. In the above, I extended the simple framework of Section 4.4.1 so that the firm may pay more attention to information or may more accurately observe sales, but at a cost. This will be key to explaining my empirical facts on predictability and autocorrelation of forecast errors, which I will show next.

4.4.3 The Size of Forecast Errors, their Predictability and Autocorrelation

Next, I show, based on the above framework, how rational inattention leads to large (absolute) forecast errors and that these are serially correlated and predictable by past sales growth.

Using the process for sales growth (4.7) and the optimal forecast (4.15), the ex-post forecast error in the framework with rational inattention is given by

$$x^{fe}_{y+1} \triangleq x_{y+1} - x^e_{y+1} = (1 - \lambda^*_y)\rho x_y - \lambda^*_y \rho \epsilon_y + u_{y+1}, \tag{4.16}$$

where I used that $s_y = x_y + \epsilon_y$. I will use this equation to derive three results from my model.

Result 1. An increase of the cost c_y from zero to a positive value results in larger absolute forecast errors and a violation of the full information rational expectations hypothesis.

Without costs for attention, $\lambda_y^* = 1$ and firms make rational forecasts since the absolute forecast error is given by

$$\left|x_{y+1}^{fe}\right| = \left|x_{y+1} - x_{y+1}^{e}\right| = \left|\rho x_y + u_{y+1} - \lambda_y^* \rho(x_y + \epsilon_y)\right| = \left|u_{y+1}\right|,$$

which is purely random. Note that the noise, ϵ_y, is zero for $\lambda^* = 1$ as implied by equation (4.9). A positive cost, $c_y > 0$, reduces λ_y^* to positive values strictly lower than unity. In this case the absolute forecast error is

$$\left|x_{y+1}^{fe}\right| = \left|x_{y+1} - x_{y+1}^{e}\right| = \left|(1-\lambda^*)\rho x_y - \lambda^* \epsilon_y + u_{y+1}\right| \leq \left|(1-\lambda^*)\rho x_y - \lambda^* \epsilon_y\right| + \left|u_{y+1}\right|.$$

Since $\left|(1-\lambda^*)\rho x_y - \lambda^* \epsilon_y\right|$ is typically larger than zero, this absolute forecast error for $0 < \lambda^* < 1$ is larger than the one for the case $\lambda^* = 1$.[9] In presence of positive cost, $0 < \lambda^* < 1$ subsequently understates the persistence of sales growth. The forecast error's dependence on understated persistence of sales growth, rather than solely on the random variables, implies firms violate the FIRE hypothesis. This finding is consistent with my empirical results in Section 4.1 on differences between major and minor forecast errors.

[9] Only in the exceptional case of zero sales growth and at the same time a zero realization for the noise shock, this term would be exactly zero and the forecast error would be not strictly larger, but of the same size as the one without information costs.

Result 2. For a strictly positive cost c_y, forecast errors are predictable by past realizations, and the forecast error is negatively correlated with lagged sales growth if (and only if) $\rho < 0$. Forecast errors are not predictable if $c_y = 0$.

This result follows from equation (4.16) and the discussion of Result 1. As explained above, following an increase in the cost c_y from zero to a positive value, the attention parameter λ_y^* reduces from unity to positive values strictly lower than one. For $\lambda_y^* = 1$, the forecast error as given in equation (4.16) is not predictable as it only depends on the i.i.d. shock u_{y+1}. For $0 < \lambda_y^* < 1$, the forecast error is predictable as it additionally depends on sales via the term $(1-\lambda_y^*)\rho x_y$ and is also larger (in absolute value) as shown in Result 1. The forecast error can only be negatively correlated with lagged sales growth if the coefficient $(1-\lambda_y^*)\rho$ in equation (4.16) is negative, which only is the case if $\rho < 0$. In Appendix J.1 I provide empirical evidence from my dataset that the autocorrelation of sales growth is indeed negative — this is also consistent with evidence in the literature, see e.g. Barrero (2019). Taken together, Results 1 and 2 are also consistent with the empirical findings based on equation (4.2) in Section 4.1. In this section, I document the predictability of major forecast errors as well as a negative relation between these major forecast errors and lagged sales growth. I further find that minor forecast errors are not predictable.

Result 3. For a strictly positive cost c_y the autocorrelation of forecast errors is negative if (and only if) $\rho < 0$. Zero cost for attention, $c_y = 0$, implies the autocorrelation of forecast errors is zero.

Substituting $x_y^e + x_y^{fe}$ for x_y in equation (4.16) and using that $x_y^e = \lambda_{y-1}^* \rho s_{y-1}$ as well

as the definition of the signal, I obtain

$$x_{y+1}^{fe} = (1 - \lambda_y^*)\rho x_y^{fe} + (1 - \lambda_y^*)\lambda_{y-1}^*\rho^2(x_{y-1} + \epsilon_{y-1}) - \lambda_y^*\rho\epsilon_y + u_{y+1}.$$

The coefficient on the forecast error, x_y^{fe}, in the equation above is negative for $0 < \lambda_y^* < 1$ only if $\rho < 0$, and it is zero for $\lambda_y^* = 1$. I showed with Results 1 and 2 that a positive value of the cost c_y implies $0 < \lambda_y^* < 1$, and that $c_y = 0$ implies $\lambda_y^* = 1$. Also Result 3 is consistent with my empirical findings. The estimation results of equation (4.4) in Section 4.2 show that major forecast errors are negatively autocorrelated. I further document in that section that the autocorrelation of minor forecast errors is not significantly different from zero.

Overall, my model shows that at times without the attention cost, the firm is fully informed and makes decisions in line with the FIRE hypothesis. In this case, forecast errors on sales growth are neither predictable nor autocorrelated. As soon as the cost for information occurs in the market environment in which the firm operates the FIRE hypothesis will be violated, absolute forecast errors will increase, forecast errors are predictable (negative correlation) by past sales growth and they exhibit negative autocorrelation. All these implications of my theoretical model are consistent with my empirical results documented in Section 4.1. The model can also rationalize my negative estimates of the coefficients on persistence and autocorrelation for major forecast errors. I show that the negative sign of these estimates is the result of the negative autocorrelation of sales growth in my data.

The above has shown that, despite its simplicity, my model is able to rationalize my main empirical findings. Key for the model results to hold are variations in firm's optimal level of rational inattention, λ_y^*, that depends on the cost for information governed by c_y.

The literature on rational inattention often remains agnostic about the specific drivers of the cost for information in such models. The empirical evidence in Chapter 2 documents that major forecast errors are not specific to a particular time or selected firms, but occur relatively evenly throughout the panel. They further do not have a very high persistence and hence have a tendency to alternate with minor forecast errors. This suggests that changes in rational inattention through variations in c_y would, in my case, be less likely to capture macroeconomic or low-frequency shocks, but may be linked to high-frequency effects.[10] A variety of reasons, and a combination of these, can be behind the latter, for example changes in the specific market environment, regulatory changes, uncertainty about sales markets or supply chains, and adaptations to firm internal processes that temporarily limit the firm's attention to information. The information cost is — as typically used in the literature on rational inattention — an abstract way of capturing changes in firms' behavior over time. Given that my aim was to develop a parsimonious model to rationalize my empirical results, a model with a fully endogenous cost function goes beyond the scope of this thesis and I leave it for future research.

4.5 Concluding Remarks

In this chapter I document that only major errors in firms' sales forecasts are predictable and autocorrelated. In contrast, minor forecast errors are neither predictable nor autocorrelated. To arrive at this result, I rely on the quantified forecast errors I obtained from my quantification model in Chapter 3. In order to interpret my empirical results that show that the Full Information Rational Expectations hypothesis is violated only following major forecast errors, I also provide a simple model of rational inattention.

[10] Time variation in the level of attention increases the complexity of solving the information problem substantially if one goes beyond a simple setup such as mine. Other papers in the literature develop theoretical frameworks where attention varies at business cycle frequencies and use simplifying assumptions to keep the problem tractable. See for example Macaulay (2019).

Firms optimally limit their degree of attention to information when operating in market environments where information processing is more costly. This limited attention leads to larger forecast errors that are predictable and autocorrelated.

Chapter 5

Decomposing the Survey-based Forecast Error Variance into a Forecastable and a Non-Forecastable Component

Firms make forecasts to plan their activities ahead. Information inefficiencies and uncertainty make it harder for firms to accurately forecast their future and plan accordingly. I construct a simple and intuitive measure of uncertainty that builds on the fact that uncertainty makes forecasts more difficult. Additionally, I construct a test of information inefficiencies in firms' forecasts that also measures the magnitude of these inefficiencies. Using firm-level survey-based forecasts, I decompose the widely used forecast error variance into two components; a forecastable and an unforecastable one. To do so I apply the Auto-regressive Conditional Heteroscedasticity (ARCH) methodology on firm-level survey forecast errors after adapting its estimation to accommodate my panel data structure. The unforecastable component of the variance is my proposed measure

of micro-uncertainty and has the advantage of having firm-month frequency. With the unforecastable component I test whether firms efficiently use the available information when they form their expectations.

In particular, using the two components of the survey-based forecast error variance, the forecastable and the non-forecastable, I engage in three exercises. First, using the unforecastable component as a measure of micro-uncertainty, I reproduce stylized facts of the literature, namely the negative correlation of uncertainty with the aggregate manufacturing production growth (Bachmann et al. (2013)). Second, I show that larger micro-uncertainty increase the probability of making forecast errors as well as that of large forecast errors in the survey data. Altig et al. (2019) also find evidence that micro-level uncertainty has predictive power over the magnitude of future forecast errors. Third, using the forecastable component of the forecast error variance and the law of total variance, I construct a test of the Full Information Rational Expectations Hypothesis. The law of total variance states how the conditional expectation and the conditional variance can be summed up to make the unconditional variance.

Uncertainty has gained a lot of attention over the last few years, since the seminal works of Bloom (2009) and Bachmann et al. (2013). The innovative uncertainty measure that I propose contributes in the literature in a number of ways ways. First and foremost, it is of firm-month frequency while most existing measures and proxies are time-series. Altig et al. (2019) developed a novel measure of micro-uncertainty of firm-month frequency that requires data on the firm-month subjective conditional distribution. However, subjective conditional distributions are not recorded by most firm-level surveys. Second, the measure of uncertainty I propose can be widely applied to all surveys that collect forecast-errors (Business Climate Surveys, Surveys of Professional Forecasters etc.) without restricting assumptions and without much computational nor

analytical effort. I call my proposed uncertainty measure "micro-uncertainty", because it is related to firm's own production, and not to aggregate macro-economic variables (see for example Bloom (2014)). I should note, however, that my methodology is still applicable to forecasts concerning aggregate variables as well. My uncertainty measure is also a 'subjective' one in the sense of Bloom et al. (2017) and Altig et al. (2019). This means that it captures the subjective uncertainty the firm faces in its own forecasting.

With my data I have the advantage that I observe and use the firm' forecasts that are conditional on their private information set. Hence, I do not attempt to forecast the firm's performance myself as in Gilchrist et al. (2014). Uncertainty in Gilchrist et al. (2014) is based on the authors' forecasts of firm's performance. Gilchrist et al. (2014)'s forecasts are conditional on the publicly available information and rely on the assumption that the firm's performance cannot be forecasted by other variables or by any other specification. As I directly observe the forecasts that the firm has made, I do not impose any such assumption. Additionally, my micro-uncertainty measure stems from the core concept of uncertainty that the future becomes less predictable, while it does not rely entirely on the empirical variance as in Bachmann et al. (2013), Bloom (2014), Bloom et al. (2018) and Bachmann et al. (2019). A large (if not exhaustive) variety of existing uncertainty measures can be found in Kozeniauskas et al. (2018). All the measures in their study however are aggregate time-series relying on the cross-sectional variance. By decomposing the widely used empirical variance, I show that it has two components and only the unforecastable one is associated with uncertainty. This substantiate my proposed uncertainty measure even more.

In addition to identifying uncertainty, survey data has also been the focus of attention as a mean of testing the full information rational expectations (FIRE) hypothesis. Pesaran (1987) and Pesaran and Weale (2006) examine methodologies to test for the

FIRE hypothesis using survey data, but they do not distinguish whether the rejection of the FIRE hypothesis comes from information inefficiencies. Pesaran (1987) defines 'information inefficiency' as the predictability of forecast errors by information that was observable by the forecaster at the time of the forecast. Intuitively, the forecaster should be able to foresee their forecast error and adjust their forecast accordingly. Coibion and Gorodnichenko (2015) take the seminal step and give us a way to test whether the rejection of the FIRE hypothesis arises from information inefficiencies. They show that that either a signal extraction mechanism or information stickiness can result in forecast errors predictable by past forecast updates, and they also derive a measure of the magnitude of the information inefficiencies.

What hinders the wide applicability of the Coibion and Gorodnichenko (2015) methodology is that it requires observations of forecast updates which are not available in most surveys. An advantage of the FIRE hypothesis test that I develop is that it does not rely on any assumptions about how the predictability of the forecast errors arise, which makes it broadly applicable. However, I show that the predictability from past forecast updates or past realizations can be special cases of my more general test. My proposed test can be carried away without forecast updates and can be widely applicable to any survey-based forecast errors, either qualitative or quantitative. Most importantly, I show that my test rejects the FIRE hypothesis only when there are information inefficiencies, while it measures the magnitude of these inefficiencies.

In the next section I discuss the sample and its quality. Then, I present the ARCH model that will decompose the forecast error variance to a forecastable and an unforecastable component, I discuss the estimation methodology and present the estimation results. Next, I introduce my measure of uncertainty being the unforecastable component, I reproduce the stylized facts, and I validate it as an uncertainty proxy. Finally,

I use the law of total variance and the forecastable component to test the rational expectations hypothesis and measure the magnitude of information inefficiencies in firms' expectations.

5.1 Data

Before getting into the details of the decomposition of the survey based forecast errors, I discuss my data in this section. For this chapter I only use the survey data without matching them with the financial surveys and I also extend the time horizon: I will study the survey data from 1998 until 2019. In my earlier applications, the sample horizon was limited by the Financial Statements that are only available from 1998-2015. In particular, I use the categorical survey responses on production from the Manufacturing sector of Greece for a period from January 1998 to December 2019 and of firm-month frequency. The (translated) questions from the survey that I will use are

> Question A.1: *During the previous 3 months, your total production, has increased/ remained unchanged/ decreased.*
>
> Question D.1: *During the next 3 months, you expect your total production to increase/ remain unchanged/ decreased.*

These qualitative survey responses are coded in the data as +1/0/-1 indicating an increase/remain unchanged/decrease, respectively. In the following, I label the variables that include the responses of firm i in month m to questions A.1 as QS_{im}, and to question D.1 as Q^e_{im}. In Chapter 2, I have the details of the survey data collection. The only cleaning step I take for this chapter is to eliminate the imputed survey responses. This leaves me with 1,737 firms and 44,767 firm-month observations.

In Figure 5.1 I show the survey responses to question D.1, indicating an expected

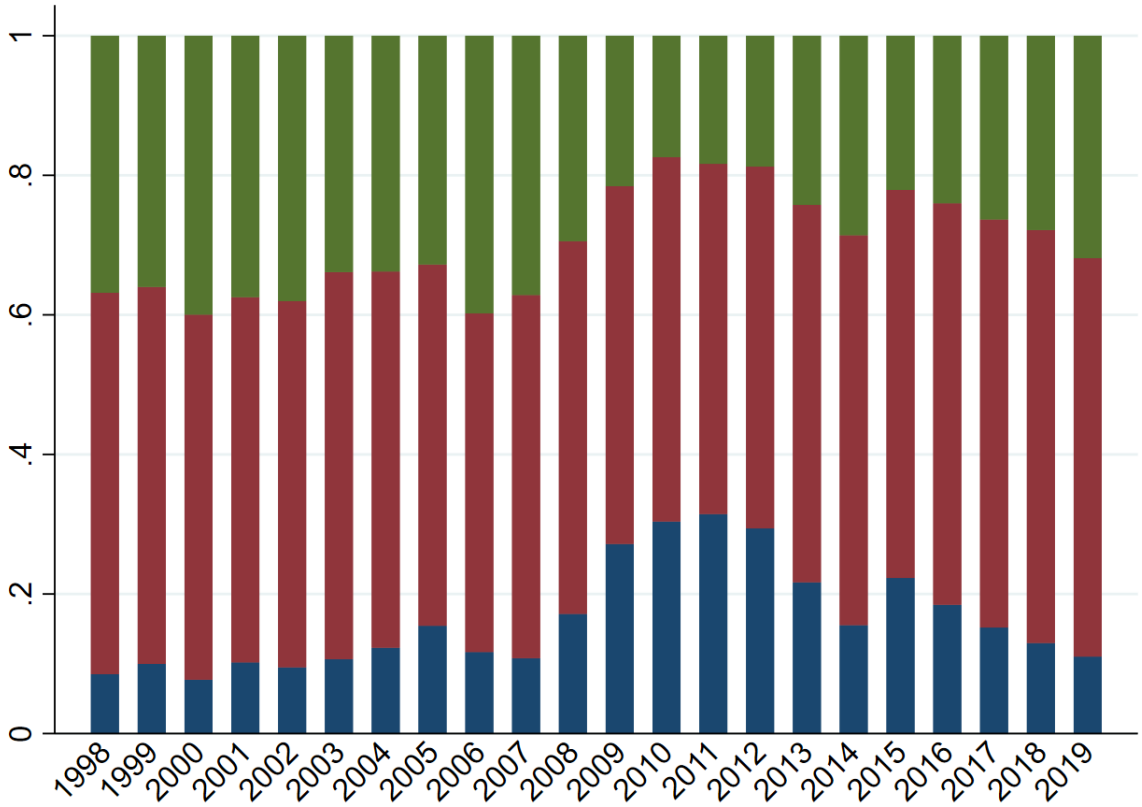

Figure 5.1: **Distribution of the Survey Forecasts on Production across Years.** The figure shows the survey responses indicating an increase/unchanged/decrease in green/red/blue as share of total monthly observations per year.

increase/ unchanged/ decrease in green/red/blue as share of total monthly observations per year. Readers may observe that the responses expecting a decrease (blue) are much higher during the years of prolonged contraction, namely 2009-2015, than during the years of economic expansion. Responses expecting an increase (green) are also more frequent during the economic expansion ended in 2008. These are in accordance with our economic intuition.

As the survey questions indicate, each firm i reports each month m the performance of production during the preceding three months as a whole. This is survey variable QS_{im} (Question A.1). This firm also posts its forecast for $QS_{i,m+3}$ in month m. This is survey variable QS^e_{im} (Question D.1), where superscript ^{e} denotes the expectation.

Table 5.1: Possible Values of Forecast Error.

		$QS_{i,m+3} =$		
		-1	**0**	**1**
	-1	0	+1	+2
$QS^e_{im} =$	**0**	-1	0	+1
	1	-2	- 1	0

Rows refer to expected direction of change in production for the following three-month period. Columns refer to the realized direction of change in production of the preceding three-month period. -1 indicates a decrease, 0 no change, $+1$ an increase. I compute the survey forecast error as $QS^{fe}_{im} \triangleq QS_{i,m+3} - QS^e_{im}$.

Therefore, the forecast error made in month m for the following three-month period is defined as $QS^{fe}_{im} \triangleq QS_{i,m+3} - QS^e_{im}$. The survey-based forecast error takes only values $\{-2, -1, 0, +1, +2\}$ and I show their computation in Table 5.1. Intuitively, -2 ($+2$) forecast error values indicate very optimistic (pessimistic) forecasts that completely missed the direction of change. In Figure 5.2, I show the distribution of the survey-based forecast errors. We observe that the majority of the qualitative forecasts are accurate with more than 60% of them being 0. Also, forecast errors with values ± 2 are particularly infrequent.

Information Quality of the Survey Data

As I did in Chapter 2, I proceed here with two exercises and show that the extended survey data on output growth shows both internal and external validity.

First, I examine the consistency of the survey-based response on output with the real production growth computed from the financial statements. For this exercise, I have matched my survey data on production with the production growth from the financial

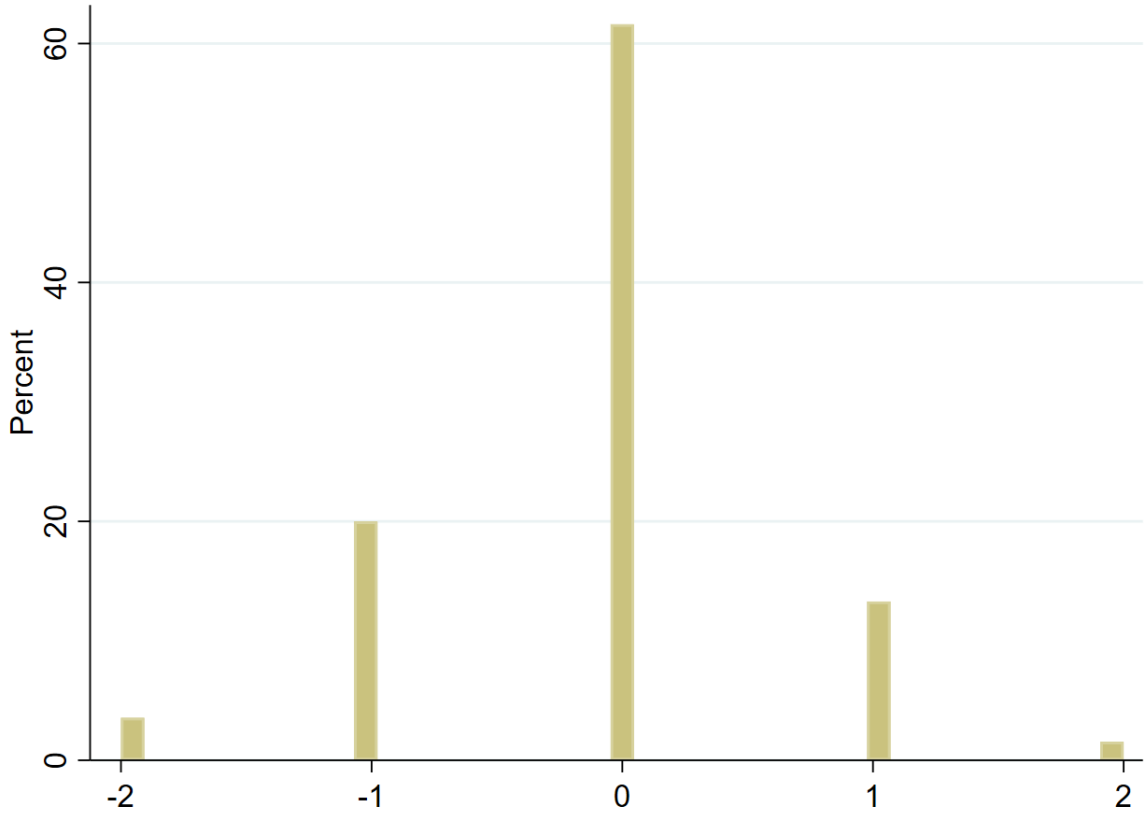

Figure 5.2: **Distribution of the Survey Forecast Errors on Production.**

statements for the years 2000-2015 when this data is available.[1] I expect that the annual growth of real production, q_{iy}, in the financial statements, must be positively correlated with the survey question concerning the evolution of production, QS_{im}. This is equivalent to the linear equation

$$QS_{im} = \beta_0 + \beta_1 q_{iy} + \psi_i + \psi_y + \eta_{im}, \tag{5.1}$$

where QS_{im} indicates the direction of change in past production and q_{iy} the real growth rate of sales; ψ_i and ψ_y control for firm and year fixed effects respectively; η_{im} is the idiosyncratic error.

I estimate equation (5.1) in two ways: (i) with standard fixed effects tools; (ii) with NACE sector dummies to substitute for the panel fixed effect. I report the estimation results in Panel A of Table 5.2. There, we observe that the monthly survey responses about the change in production are on average positively and significantly (sig. 1%) correlated with the real production growth rate from the financial statements. In other words, survey based responses for production are also consistent with the financial statements.

Second, I re-estimate equation (2.2) for the new sample period (1998-2019). This equation examines the consistency between on production and capacity utilization. I check that an increase in the production reported in the surveys is positively associated with an increase in capacity utilization. As a reminder to check this relationship, I estimate the following equation which is exactly the same as (2.2)

$$QS_{im} = \beta_0 + \beta_1[U_{i,m-1} - U_{i,m-3}] + \psi_i + \psi_y + \eta_{im},$$

[1] I compute production from the financial statements as the sum of sales plus the contemporaneous first difference of final goods inventories. I deflated the firm-year output of the financial statements using the ratio of the nominal over real (chain linked volumes) gross value added at the NACE 2-digit level. The first two years of the financial statements (1998 and 1999) are dropped because I take first differences of final goods inventories to compute production level and then I compute their growth rates.

Table 5.2: Consistency of survey responses

PANEL A: Dependent Var. QS_{im}			**PANEL B:** Dependent Var. QS_{im}		
q_{iy}	0.146***	0.151***	$U_{i,m-1} - U_{i,m-3}$	0.00511***	0.00505***
Constant	0.220***	0.302***	Constant	0.264***	0.329***
RE/FE	FE	RE	RE/FE	FE	RE
NACE FE	NO	YES	NACE FE	NO	YES
Observations	24,353	24,353	Observations	15,663	15,663
Overall R^2	0.06	0.072	Overall R^2	0.052	0.07
Number of firms	1,101	1,101	Number of firms	890	890

Estimations with NACE FE were made with Random Effects GLS in pooled data (RE). All variables (apart from NACE 2-digit code) are survey questions. NACE fixed effects are taken at the 2-digit level. Fixed year effects are omitted to simplify representation but are included in the estimation. U_{im} is the percentage capacity utilization; QS_{im} is the recent change of production. q_{iy} is real production growth from financial statements. Complete details about the exact wording of the questions are in the text of this section. *** denotes significance at the 1% level.

where β_0 and β_1 are parameters to be estimated, QS_{im} indicates the change in past production for firm i in month m, ψ_i and ψ_y control for firm and year fixed effects respectively, and η_{im} is the idiosyncratic error. U_{im} corresponds to survey question E.3: *'During the ongoing period, what is your percentage (%) utilization of your production capacity?*, and firms respond to this question with a quantitative answer.

As previously, I estimate equation (5.1) in two ways: (i) with standard fixed effects tools; (ii) with NACE sector dummies. I report the estimation results in Panel B of Table 5.2. These agree with my economic intuition: a reported increase in production is positively and significantly correlated with a three-month increase in % capacity utilization (from $m-3$ to $m-1$).

Overall, based on the results in Tables 5.2, I can conclude that survey responses are both consistent with each other within the questionnaire, but also with the financial statements. These results add validity to the the data that I use for this chapter. Having documented the quality of my dataset, I now proceed with showing how the variance of the survey-based forecast error can be decomposed in a forecastable and a

non-forecastable component.

5.2 Decomposing the Forecast Error Variance and the ARCH Model

To distinguish the forecastable from the unforecastable part of the variance of the survey-based forecast errors I apply the ARCH methodology. The ARCH models can predict a part of the forecast error variance. This is the forecastable component of the variance of the forecast errors that also measures the persistence of the variance. Using the forecastable component, I will test the FIRE hypothesis. The remaining part of the variance is the unforecastable component. I will measure uncertainty with the unforecastable component (error) of the auto-regressive variance.

To implement Engle (1982)'s ARCH analysis I need to make two assumptions. First, the firm posts its monthly survey forecasts using an information set $\mathscr{F}_{i,m-1}$. Therefore, the conditional mean is the forecast (for the following three-month period) made by the firm i in month m: $QS^e_{im} \triangleq \mathbb{E}\big[QS_{i,m+3}|\mathscr{F}_{i,m-1}\big]$ which is directly observable in the survey data. Accordingly, I identified the forecast error as $QS_{i,m+3} - QS^e_{im} = QS^{fe}_{im}$. Second, I assume that the forecast error variance follows an ARCH process or order 1.[2] The auto-regressive conditional heteroscedasticity is the 'forecastable variance' component, and I denote it with FV_{im}.

The survey data gives me an important advantage. I can directly observe the conditional expectations in the surveys, so I do not need to estimate any first stage equation to obtain the forecast errors. In other words, the forecast error is entirely attributed to

[2] I test the assumption that the forecast error variance is an ARCH process of order 1, by including an additional term on the right hand side of (5.2),namely $\big(QS^{fe}_{i,m-6} - \overline{QS^{fe}_{m-6}}\big)^2$. The 2nd order autocorrelation coefficient is not significant at 10%. In Appendix L and in Table 1.L I show the estimation results.

the firm and it is not the result of the information or the specification chosen by the econometrician.

Having assumed that the variance is autoregressive of order 1, I obtain the following equation

$$\left(QS_{i,m}^{fe} - \overline{QS_m^{fe}}\right)^2 = FV_{im} + \psi_m + \psi_i + \eta_{im} = \phi\left(QS_{i,m-3}^{fe} - \overline{QS_{m-3}^{fe}}\right)^2 + \psi_m + \psi_i + \eta_{im}, \quad (5.2)$$

where ϕ is the auto-correlation coefficient, ψ_i is the unobserved firm heterogeneity, ψ_m month-specific aggregate effects and η_{im} the idiosyncratic error. $\overline{QS_m^{fe}}$ is the cross-sectional arithmetic mean of the forecast errors. To compute the variance, one needs the squared deviations from the mean, so I deduct the monthly average of the forecast errors to allow for the mean to be time-dependent.[3]

On the right hand side, the squared forecast error $\left(QS_{i,m-3}^{fe} - \overline{QS_{m-3}^{fe}}\right)^2$ is lagged by three months so that the three-month horizon that it covers does not overlap with the horizon of the left hand side. Month-specific effects are essential to account for aggregate uncertainty that would lead to correlated η_{im} across firms. Essentially, equation (5.2) breaks down the forecast error variance in two components. First, I obtain the forecastable variance

$$FV_{im} \triangleq \phi\left(QS_{i,m-3}^{fe} - \overline{QS_{m-3}^{fe}}\right)^2. \quad (5.3)$$

Second, after identifying the forecastable variance, I can obtain the unforecastable component of the forecast error variance

$$UNCERT_{im} \triangleq \left(QS_{i,m}^{fe} - \overline{QS_m^{fe}}\right)^2 - \phi\left(QS_{i,m-3}^{fe} - \overline{QS_{m-3}^{fe}}\right)^2, \quad (5.4)$$

[3] I show later that my estimates are robust to using the squared forecast error without demeaning it. This is to be expected as demeaning (deducting a first moment) does not affect the variability of any variable (second moment).

which will be my measure of micro-uncertainty.

The uncertainty the firm faces in forecasting its own production is naturally a mixture of time-specific (aggregate) factors, firm-specific factors and idiosyncratic factors as well.[4] These terms collectively affect the firm's uncertainty in forecasting its own production, and the uncertainty measure I propose successfully incorporates all of them. As equation (5.2) demonstrates, the micro-uncertainty measure of equation (5.4) also includes the aggregate monthly error ψ_m, the firm-specific error ψ_i, and the idiosyncratic error η_{im}. Finally, in equations (5.3) and (5.4), both variables of the forecastable and unforecastable variance are of firm-month frequency. That is, I provide a measure of uncertainty that has firm-month frequency as I mentioned earlier.

5.2.1 ARCH Estimation

To estimate the dynamic equation (5.2), the standard fixed effects estimator (also known as the 'least squares dummy variables' estimator, LSDV) suffers from a severe negative bias, the 'Nickell (1981) bias', when the time dimension of the panel is finite as with my data (see also Pesaran (2015)).[5] The next unbiased alternatives are the Anderson and Hsiao (1981) and the Arellano and Bond (1991) First-Differences estimators. As I discussed more thoroughly in Chapter 4, the first difference estimators, although they do not suffer for the Nickell bias, drop a lot of observations by taking first differences to eliminate the fixed effects. This problem becomes particularly severe in unbalanced panels with frequent time gaps like mine (see Roodman (2009) and Gorbachev (2011)). To

[4]As a reminder, the term 'micro' in micro-uncertainty is used to link the uncertainty of the firm's forecasts on its own variable and not on a macroeconomic variable. That is, 'micro' does not refer to the source of information in firm's forecasts, rather to the nature of the variable at hand. Kozeniauskas et al. (2018) explore and model different measures of uncertainty and show that they all increase as a result of macro-economic volatility.

[5]Although the survey-based forecast error $QS^{fe}_{i,m}$ takes discrete values, the dependent variable of equation (5.2), $\left(QS^{fe}_{i,m} - \overline{QS^{fe}_m}\right)^2$, takes continuous values. This allows me to treat equation (5.2) as linear.

address this problem, Arellano and Bover (1995) suggest the forward orthogonal transformations (FOT) to eliminate the unobserved firm heterogeneity. The FOT subtracts from each observation the firm-specific arithmetic mean of its future values, instead of taking first differences. This significantly limits data loss, allows me to use all available data points, and it takes advantage of the efficiency of the two-step GMM estimator.

Moreover, to avoid the proliferation of moment conditions in the GMM estimator I take some additional steps following the suggestions of Roodman (2009) and the standard practice (Gorbachev (2011) and Caselli and Tesei (2016)). I collapse the instruments, and I use only one lag length as instrument. Additionally, as standard errors can also be downwards biased when the number of firms is not large enough compared to the number of moment conditions, I also use the Windmeijer (2005) corrected standard errors.

In Table 5.3, I present the estimation results of equation (5.2) with the Arellano and Bover (1995) Two-Step GMM estimator, for the entire period Jan. 1998 - Dec. 2019. I observe that the persistence of the forecast error variance is statistically significant at 1%. This validates my intuition that the forecast error variance has a forecastable component, its persistence, and therefore it does not entirely capture uncertainty. Additionally, the coefficient of the proxy for the month-specific aggregate uncertainty is significant at 1% indicating that there is a time-specific aggregate component in uncertainty as Kozeniauskas et al. (2018) suggest.

The sample size used in the GMM estimation in Table 5.3 is smaller than the one I described in the previous section for a number of reasons. First, to compute the survey-based forecast errors I need both expectations and realizations with specifically three months-difference. If a firm for example has posted a forecast in January and a realization in June of the same year, I cannot use these observations. What is more, if these are the only observations of that firm, that panel will also disappear from the

estimation. Second, to estimate the equation, I need the squared demeaned forecast errors with also three period lags. Given the unbalanced nature of the survey sample, it is reasonable for the sample to be used in the estimation to be smaller. I have confirmed that that the reduced sample used in the estimations of Table 5.3 also shows internal and external validity in the same exercises I did in the previous section (output is omitted but can be made available upon request).

Finally, in Figure 5.3 I show the histograms of the forecastable variance and of the uncertainty. Given that the persistence coefficient ϕ has a small value, 0.0695, I expect the unforecastable component to take larger values as shown on the right of Figure 5.3. Some values are larger than 4, because during these months the month-specific average forecast error was negative, so the difference $QS^{fe}_{i,m} - \overline{QS^{fe}_{m}}$ was larger than 2, and when squared larger than 4.

In Appendix Section M, I run some robustness checks on the GMM estimates of the Dynamic Panel equation (5.2). I show that using more lags as instruments, as it is standard in the literature, leads to weak identification, which fully justifies my choice of instruments. I also show that my baseline estimate of the persistence coefficient ϕ (0.0695) is much higher than the LSDV biased estimate (0.0281 in Table 1.M). This demonstrates that I have addressed the Nickell (1981) bias. Finally, I show that my estimates are robust to using the directly observed squared forecast error without demeaning it.

Table 5.3: Estimated ARCH Model for the Variance of the Survey-based Forecast Errors on Production.

Estimation	FOT
Stand. Errors	2-step, Windmeijer corrected
Lags as Instruments	$m-4$
Dependent Variable: $\left(QS^{fe}_{i,m} - \overline{QS^{fe}_{m}}\right)^2$	
$\left(QS^{fe}_{i,m-3} - \overline{QS^{fe}_{m-3}}\right)^2$	0.0695***
$\hat{\psi}_m$	1.004***
Constant	N.A.
Observations	14,019
# of Firms	678
Over-identified	No
Hansen p-value	N.A.
m2 test p-value	0.811

Table shows estimates of equation (5.2). Instruments are collapsed and with lags dated at $m-4$. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. $\hat{\psi}_m$ is the proxy of the aggregate monthly effects and is equal to the monthly cross-sectional arithmetic mean of the dependent variable. The survey-based forecast error on production is defined as $QS_{i,m+3} - QS^{e}_{im} = QS^{fe}_{im}$. *** indicates statistical significance at the 1%.

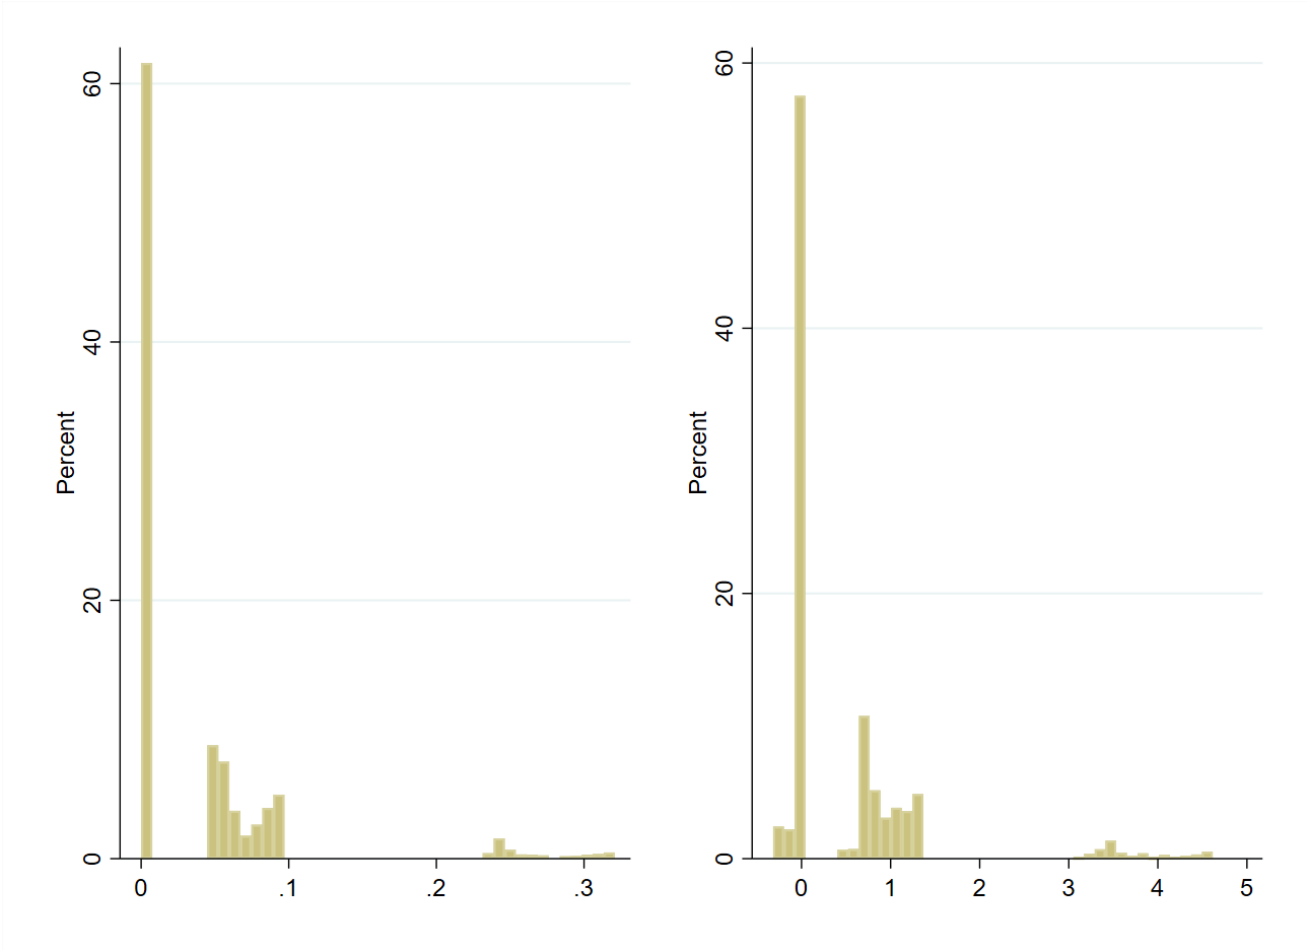

Figure 5.3: **Distribution of the Forecastable Forecast Error Variance and of the uncertainty for Production Growth.** The figure on the left shows the distribution of firm-month observations of the forecastable component of the survey-based forecast error variance. The figure on the right shows the distribution of the firm-month unforecastable component.

5.3 A Measure of Micro-uncertainty

In order to properly model micro-uncertainty I firstly need to carefully define its concept. As Jurado et al. (2015) correctly note, rather than relying on the increased volatility of a variable, higher uncertainty coincides with the fact that the future is becoming less predictable, that it is more uncertain. I built my uncertainty measure inspired by the definition of uncertainty offered by Jurado et al. (2015).

We know that as future becomes more uncertain, forecast error variance increases. However, with the ARCH model I showed that a part of this variance is forecastable and is the result of persistence. The distinction between forecastable and unforecastable variance is of vital importance. The most widely used uncertainty measures are based on the cross-sectional dispersion of the forecast errors, and purely statistical dispersions measures are not suitable without controlling for their forecastable part. By definition, uncertainty is necessarily hidden in the unforecastable component of the forecast error variance. This is the variable $UNCERT_{im}$ that I introduced with equation (5.4).

In the following subsection I cross-examine the time-series properties of my proposed measure of micro-uncertainty with respect to key existing uncertainty proxies. By reproducing stylized facts of the literature, I add validity to my proposed measure. Finally, I further show that my micro-uncertainty measure results in increased probability of forecast errors.

5.3.1 Time-series Properties of Uncertainty and Forecastable Variance

In this subsection I am interested in documenting how closely my proposed measure of micro-uncertainty follows the stylized facts appearing in the literature. For this analysis I will aggregate the cross-section and focus on the time-series properties. That is, I

create time-series using the survey-based forecasts and forecast errors and the micro-uncertainty. I also include the forecastable variance to show that it materially differs from the unforecastable component, because it fails to reproduce the stylized facts.

First, for each month, I compute the cross-sectional weighted means of the micro-uncertainty ($\overline{UNCERT}_m$) and of the forecastable variance ($\overline{FV}_m$). For any variable X_{im} with firm-month frequency, the general formula for the weighted mean is

$$\overline{X_m} = \sum_i \frac{w_{im}}{\sum_i w_{im}} X_{im}. \tag{5.5}$$

The weights w_{im} are the ratio of the Real Gross Value Added (GVA) of the 2-digit industry over the total GAV of the whole Manufacturing sector (see Bachmann et al. (2013)).[6] This weighting gives sectors that are more important in manufacturing higher weights than the less important ones and controls for over- or under- representation in the sample.

Second, I create another set of time series variables from the survey expectations and the survey forecast errors. For each month, I create the cross-sectional weighted dispersions of the forecasts ($FDISP_m$) and the forecast errors ($FEDISP_m$). The weights are the same as in the weighted means. Overall, for any variable X_{im} with firm-month frequency, the general formula for the weighted dispersion is

$$\sqrt{V[X_m]} = \sqrt{\overline{X_m^2} - \left(\overline{X_m}\right)^2}, \tag{5.6}$$

where $\overline{X_m}$ and $\overline{X_m^2}$ are weighted means computed using the formula in equation (5.5). The forecast (forecast error) dispersion is the dispersion of the survey-based forecast QS_{im}^e (forecast error QS_{im}^{fe}). The two dispersion proxies of uncertainty were introduced

[6]Source: Eurostat Table nama_10_a64 for Greece, annual frequency, in 2005 Chain Linked Volumes.

by Bachmann et al. (2013).

In Table 5.4 I report the pairwise correlation among the various uncertainty proxies. The table shows a very strong and positive correlation (0.81) between the two forecast-error based uncertainty proxies. However, their correlation with the forecast dispersion is much lower: 0.28 between forecast dispersion and forecast error dispersion; 0.21 between forecast dispersion and my proposed uncertainty proxy.

Table 5.4: Time Series Pairwise Correlations of the Uncertainty Proxies

	$FEDISP_m$	$FDISP_m$	$\overline{UNCERT}_m$	$\overline{FV}_m$
$FEDISP_m$	1			
$FDISP_m$	0.28	1		
$\overline{UNCERT}_m$	0.81	0.21	1	
$\overline{FV}_m$	0.10	0.10	-0.03	1

The table shows the time series pairwise correlations among the various uncertainty proxies. $FEDISP_m$ is the cross-sectional weighted dispersion of the forecast errors on production; $FDISP_m$ is the cross-sectional weighted dispersion of the forecasts on production; $\overline{UNCERT}_m$ is the cross-sectional weighted average of the unforecastable variance of the forecast errors on production; $\overline{FV}_m$ is the cross-sectional weighted average of the forecastable variance of the forecast errors on production. All the monthly frequency variables are linearly de-trended and cleared from month-specific effects.

Next, In Table 5.5 I show the correlations between each one of the aforementioned variables with the growth rate of the Industrial Production Index (IPI) from the Hellenic Statistical Authority (ELSTAT). The sampling dates are from January 2000 to December 2019 in monthly frequency.[7] All time series have been linearly detrended as it is standard in the literature (see Bachmann et al. (2013)), and I removed the month-specific effects (seasonality) using month indicator variables. I use the rolling three-month growth rates of the IPI in each month ($\Delta\%IPI_m$), because this is a smoother series than the month-over-month growth rates.[8] I also examine their correlation in a quarterly frequency with

[7] IPI is only available as of January 2000.

[8] $\Delta\%IPI_m \triangleq \frac{IPI_m - IPI_{m-2}}{IPI_{m-2}}$

the quarter-over-quarter growth rates of the IPI ($\Delta\%IPI_q$). Again, I linearly detrend the time-series and eliminate the quarter-specific effects. Additionally, I report the time series descriptive statistics of the various uncertainty proxies.

First, my micro-uncertainty measure is negatively correlated with the economic activity by -14% in the monthly data. This implies that as manufacturing production grows, uncertainty decreases. Bachmann et al. (2013) proxied uncertainty with two survey variables, the cross-sectional dispersion of the forecasts and that of the forecast errors and also documented a negative but weaker correlation with the industrial production activity of Germany. These results indicate that my measure of micro-uncertainty reproduces the stylized facts. Readers can also observe that the forecastable component of the forecast error variance is very weakly correlated compared to the unforecastable one which captures the micro-uncertainty. This validates my point that forecast error dispersion as a whole is not uncertainty entirely. In Table 5.4, I report the pairwise correlations among the various measures of uncertainty.

Second, in a quarterly frequency the negative correlation becomes weaker for all proxies with the exception of the dispersion of the forecasts. This indicates that the uncertainty might be short-lived. Despite these weak correlations, in Figure 5.4 I show that the error-based uncertainty if more strongly related to the business cycle than the forecast-based one. The upper panel of Figure 5.4 shows a strong positive effect of the business cycle on the forecast error based uncertainty proxies. My uncertainty proxy and the forecast error dispersion, are much higher in the shaded periods indicating the recessions than in the rest of the periods. Most importantly, they both increase as soon as the cycle turns to recession in all occasions. However, the effect of the business cycle on the forecast dispersion is not as strong, as shown on the lower panel of Figure 5.4. The forecast dispersion is not persistently increased during recessions and starts rising

Table 5.5: Time series properties of the Uncertainty Proxies

Panel A: Monthly Frequency				
	$FEDISP_m$	$FDISP_m$	$\overline{UNCERT}_m$	$\overline{FV}_m$
Min	-0.19	-0.11	-0.34	-0.02
Mean	0.00	0.00	0.00	0.00
Max	0.41	0.21	0.84	0.05
St. Deviation	0.08	0.04	0.14	0.01
Skewness	0.71	0.82	1.35	1.65
Corr w/ $\Delta\%IPI_m$	-0.15	-0.09	-0.14	-0.04
Panel B: Quarterly Frequency				
	$FEDISP_q$	$FDISP_q$	$\overline{UNCERT}_q$	$\overline{FV}_q$
Corr w/ $\Delta\%IPI_q$	-0.06	-0.18	-0.07	0.01

The table shows the time series descriptive statistics for the various uncertainty proxies. $FEDISP_m$ is the cross-sectional weighted dispersion of the forecast errors on production; $FDISP_m$ is the cross-sectional weighted dispersion of the forecasts on production; $\overline{UNCERT}_m$ is the cross-sectional weighted average of the un-forecastable variance of the forecast errors on production; $\overline{FV}_m$ is the cross-sectional weighted average of the forecastable variance of the forecast errors on production. $\Delta\%IPI_m$ is the rolling three-month growth rate of the IPI of the Industrial Production Index of ELSTAT. Data on IPI_m is available from 2000 and on. All the monthly frequency variables are linearly de-trended and cleared from month-specific effects. $FEDISP_q$, $FDISP_q$, $\overline{UNCERT}_q$, $\overline{FV}_q$ are the arithmetic means of their monthly counter-parts within each quarter. For all the quarterly arithmetic means I used the monthly time-series linearly de-trended and with excluded month fixed effects. $\Delta\%IPI_q$ is the quarter-over-quarter growth rates of the Industrial Production Index, linearly detrended and cleared from quarter-specific effects.

well after the start of the recession.

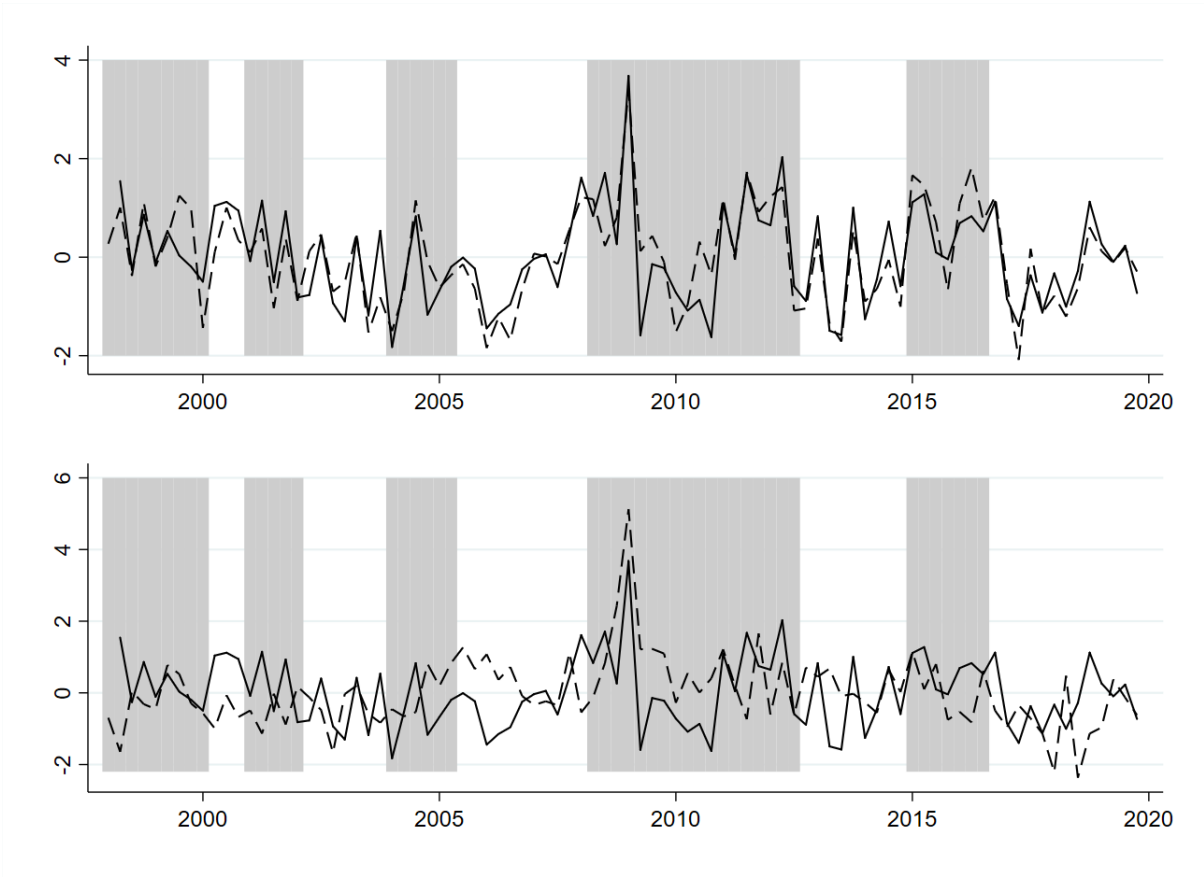

Figure 5.4: **Uncertainty Proxies across the Business Cycles.** In the upper panel, the dash line shows the quarterly average of the monthly time series $FEDISP_m$. In the lower panel, the line shows the quarterly average of the monthly time series $FDISP_m$. The solid line shows the quarterly average of the monthly time series $\overline{UNCERT}_m$ in both panels. The sample period is $1998Q1 - 2019Q4$. Before taking the quarterly averages shown, all the monthly time-series were linearly detrended, I eliminated the monthly fixed effects and normalized them by their standard deviation, as this is standard practice. The shaded periods indicate recessions and they are the quarters between a peak of the cycle and a trough as dated by OECD (2020).

5.3.2 Uncertainty and Forecast Errors

In this subsection, I examine the effect of micro-uncertainty on the probability that a firm posts a non-zero forecast error and on the probability that it posts a large forecast error. To do so, I create two dummy variables: $\mathbb{1}\left[QS_{im}^{fe} \neq 0\right]$ takes value 1 if firm i in month m made a forecast error; $\mathbb{1}\left[QS_{im}^{fe} = \pm 2\right]$ takes the value 1 if firm i in month m made a large forecast error, that is $QS_{im}^{fe} = -2$ or $+2$. In the survey data, the only possible forecast error outcomes are $0, \pm 1, pm2$ — see Table 5.1. I also examine the effects of the forecastable component and show that it does not result in larger forecast errors. Both $UNCERT_{i,m-1}$ and $FV_{i,m-1}$ are with one period lag to exclude the effects of simultaneity; their sum is the square of the contemporaneous forecast error that I use on the dependent variable. I control for the firm-specific unobserved heterogeneity using the Mundlak (1978) approximation.[9]

Table 5.6 reports the estimation results for the forecast errors on production. It shows that micro-uncertainty significantly increases the probability that a firm makes a forecast error in production. In fact, it increases the probability of any forecast error and that of a large one. This result agrees with Altig et al. (2019) who find that micro-uncertainty can predict the magnitude of future forecast errors. The forecastable component has only a positive and less significant effect on the probability that the firm posts a forecast error. These probit results validate the distinction between the forecastable and the unforecastable components of the forecast error variables. Further, they demonstrate that the unforecastable one is the strongest negative effect on the predictability of production, which is how Jurado et al. (2015) define uncertainty.

[9] As a reminder, the Mundlak (1978) approximation is the standard tool used in non-linear models in panel data. See e.g. Bartelsman et al. (1994), Kosova (2010), Semykina and Wooldridge (2010) and Triguero and Córcoles (2013)).

Table 5.6: Probit estimations of the Survey-Based Forecast Errors.

	(1)	(2)
Dep. Variable	$\mathbb{1}\left[QS_{im}^{fe} \neq 0\right]$	$\mathbb{1}\left[QS_{im}^{fe} = \pm 2\right]$
$UNCERT_{i,m-1}$	0.151***	0.154***
$FV_{i,m-1}$	0.618**	0.406
Constant	-0.0820	-2.835***
Observations	10,394 (694 firms)	10,394 (694 firms)

Year fixed effects and Mundlak (1978) proxies for the unobserved heterogeneity are omitted to maintain a simple representation, but they are included in the estimation. In column (1) I estimate a probit with the dependent variable taking value 1 when firm i in month m makes a non-zero forecast error, $\mathbb{1}\left[QS_{im}^{fe} \neq 0\right]$. In column (2) I estimate a probit with the dependent variable taking value 1 when firm i in month m makes a large survey-based forecast error, $\mathbb{1}\left[QS_{im}^{fe} = -2, +2\right]$.
*** and ** indicate statistical significance at the 1% and 5% level, respectively.

5.4 Testing the Information Efficiency of the Rational Expectations Hypothesis

In this section, using the forecastable component of the forecast error variance, I develop a test of the full information rational expectations (FIRE) hypothesis. My proposed test has two key advantages. First, it only rejects the FIRE hypothesis when firms violate the information efficiency, that is when they make predictable forecast errors, and measures the extent of the deviation from the FIRE hypothesis. More importantly, I do not need to make any assumptions about the predictability of the forecast errors; for instance, in Chapter 4 I modelled predictability using past realizations. Also, Coibion and Gorodnichenko (2015) model predictability stemming from preceding forecast updates. This results in the second key advantage: my test can be straightforwardly applied to any survey data, with either categorical or continuous forecast errors, without the need of forecast updates as in Coibion and Gorodnichenko (2015).

I test the FIRE hypothesis hypothesis by using the law of total variance (LTV). In its most general form without time dimension, LTV states that the variance of any random variable X can be decomposed as follows:

$$\mathbb{V}\big[X\big] = \mathbb{E}\big[\mathbb{V}\big[X|\mathscr{F}\big]\big] + \mathbb{V}\big[\mathbb{E}\big[X|\mathscr{F}\big]\big],$$

where $\mathscr{F}$ is the information set.

In my context, I apply the LTV on the forecast error made in month m by firm i, QS_{im}^{fe}, using $\mathscr{F}_{i,m-3}$ as the information set available to the firm in $m-3$. Decomposing the variance of the forecast error using the LTV, I obtain

$$\mathbb{V}\Big[QS_{im}^{fe}\Big] = \mathbb{E}\Big[\mathbb{V}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\Big] + \mathbb{V}\Big[\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\Big]. \tag{5.7}$$

I use the three-month lag in the information set, because later I relate equation (5.7) to the ARCH estimates where I use three-month lag on the right had side. Before showing how I estimate the terms of equation (5.7) from my data and from the forecastable component of the variance, I will show first how equation (5.7) relates to the FIRE hypothesis.

Equation (5.7) and the FIRE hypothesis

If the FIRE hypothesis holds, and firms efficiently use the available information, then the forecast error should not be predictable, i.e. $\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big] = \mathbb{E}\big[QS_{im}^{fe}\big] = 0$. If the expected forecast error is zero, its variance will also be zero. Although not sufficient, this is a necessary condition for the FIRE hypothesis to hold. If firms use information efficiently and forecast errors are not predictable but show some bias then the forecast error has a constant non-zero mean, i.e. $\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big] = \mathbb{E}\big[QS_{im}^{fe}\big] = c$. In that case

my test does not reject the FIRE hypothesis, because the variance of the constant bias is still zero. That is, my test only rejects the FIRE hypothesis when my forecasters are not efficient.[10] Formally, if the the FIRE hypothesis holds, then $\mathbb{V}\Big[\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\Big] = \mathbb{V}[0] = 0$ which implies in equation (5.7) that

$$\mathbb{V}\Big[QS_{im}^{fe}\Big] - \mathbb{E}\Big[\mathbb{V}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\Big] = \mathbb{V}\Big[\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\Big] = 0. \tag{5.8}$$

As I will document below, $\mathbb{V}\big[\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\big] \neq 0$, so the FIRE hypothesis does not hold.

What is more, I argue that the quantity $\mathbb{V}\big[\mathbb{E}\big[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\big]\big]$ is a measure of the magnitude of the inefficiency in firms' forecasts. I showed that under the FIRE hypothesis hypothesis, the expected forecast error is zero, so its variance must also be zero. If the firm deviates from efficiency, the expected forecast error becomes stochastic and its variance increases. The information that the firm ignores (uses inefficiently) at the time of the forecast stacks on the forecast error and is what makes that error predictable. The higher the variance of the expected forecast error the more information the firm has ignored, i.e. the more it has deviated from efficiency. As a result, the variance of the expected forecast error is a measure of the magnitude of the inefficiency in expectations.

Observe here that this measure of the magnitude of the inefficiencies has a very general form. I have not made any assumptions about the origin of the predictability of the forecast errors. My proposed generalized measure of inefficiency can be easily specialized to the inefficiency measure introduced by Coibion and Gorodnichenko (2015). Assuming either a sticky information environment or a signal extraction mechanism, Coibion and Gorodnichenko (2015) derive a reduced form for the expected forecast error

[10] Researchers can always independently test whether $\mathbb{E}\big[QS_{im}^{fe}\big] = 0$ to make sure there is no bias.

that is a linear function of the preceding forecast update

$$\mathbb{E}\big[QS_{im}^{fe}|UPDATE_{i,m-3}\big] = \alpha + \beta\big[UPDATE_{i,m-3}\big], \tag{5.9}$$

where $UPDATE_{i,m-3}$ is the forecast update from 3 months before.[11] α and β are parameters; β is the Coibion and Gorodnichenko (2015) measure of inefficiency capturing the extent of information rigidities that their model assumes. If I plug equation (5.9) into the variance of the expected forecast error I obtain

$$\mathbb{V}\Big[\mathbb{E}\big[QS_{im}^{fe}|UPDATE_{i,m-3}\big]\Big] = \beta^2\mathbb{V}\big[UPDATE_{i,m-3}\big]. \tag{5.10}$$

In equation (5.8), if I use the conditioning on the update and the result of equation (5.10), I will obtain

$$\mathbb{V}\Big[QS_{im}^{fe}\Big] - \mathbb{E}\Big[\mathbb{V}\big[QS_{im}^{fe}|UPDATE_{i,m-3}\big]\Big] = \beta^2\mathbb{V}\big[UPDATE_{i,m-3}\big],$$

which can give an estimate of the parameter β, the Coibion and Gorodnichenko (2015) measure of inefficiency. Clearly, this is a special case of equation (5.8).

Alternatively, I can also model the forecast errors using the simple Rational Inattention model that I developed in Chapter 4. Following the same steps as above, one can see that this will be another special case of equation (5.8). Evidently, by refraining from assumptions about the form of the expected forecast error, I obtain the advantage of the general applicability. My methodology does not require any additional data, such as forecast updates, other than the ex-post forecast errors. The forecast errors can almost always be observed!

[11] My survey data does not record forecast updates in any way. I coin the variable here for expositional purposes.

Estimating Equation (5.7) from my Data

In equation (5.8), the first term on the left hand side is the unconditional variance of the forecast error and is directly observable from the the survey data. It is the empirical cross-sectional weighted variance of the survey-based forecast errors in each month m, $\left(FEDISP_m\right)^2$. The second term on the left hand side of equation (5.8) is the mean forecastable variance and I can obtain it from the ARCH model and equation (5.3). It is the the cross-sectional weighted mean of the forecastable component in each month, $\overline{FV}_m$. An implicit assumption here is that only lagged values of the forecast error variance can predict the forecast error variance and not other variables from the firm's information set — I test this assumption at a later stage. This allows me to equate $\mathbb{V}\left[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\right]$ with the forecastable component from the ARCH model, FV_{im}.

If I substitute the aforementioned variables from the data for their theoretical counterparts in equation (5.8) I can test the FIRE hypothesis in my data. That is, I need to test in the data that the following equation holds

$$\left(FEDISP_m\right)^2 - \overline{FV}_m = 0. \tag{5.11}$$

From equations (5.8) and (5.11) I know that the difference $\left(FEDISP_m\right)^2 - \overline{FV}_m$ is an estimate of the variance of the expected forecast error $\mathbb{E}\left[QS_{im}^{fe}|\mathscr{F}_{i,m-3}\right]$. The lower its value, the lower the variance of the conditional expectation of the forecast error and the closer I get to the FIRE hypothesis. As I showed, the quantity $\left(FEDISP_m\right)^2 - \overline{FV}_m$ measures the extent of the deviation from the FIRE hypothesis.

A simple one-sample t-test would be enough to test (5.11), but the time-series $\left(FEDISP_m\right)^2 - \overline{FV}_m$ is most likely heteroscedastic and serially correlated. As a result

I estimate the following linear equation

$$\left(FEDISP_m\right)^2 - \overline{FV}_m = \beta_0 + u_m, \tag{5.12}$$

and I test the hypothesis that $\beta_0 = 0$ using the Newey-West (NW) standard errors that are robust to heteroscedasticity and to serial autocorrelation. If I find that $\beta_0 = 0$, then the monthly values of $\left(FEDISP_m\right)^2 - \overline{FV}_m$ are purely random (u_m) and do not constitute information inefficiencies.

For the NW standard errors I choose 12 lags.[12] Also, because of some gaps in the time-series (missing data) of $\left(FEDISP_m\right)^2 - \overline{FV}_m$, I follow Datta and Du (2012) and Greenwood and Shleifer (2014), and I treat all existing monthly observations as 'equally spaced' for the Newey-West estimator. That is I ignore the gaps completely.

Table 5.7 shows the estimation results for equation (5.12). I report the Newey-West standard errors and t-statistics calculated using them. I reject the null that $\beta_0 = 0$ at significance level of 1%, which results in rejecting the FIRE hypothesis. These results are robust to: (i) using fewer lags (6 lags); (ii) to including month-specific effects to control for seasonalities (results are not shown but can be made available upon request).

Next, in Table 5.8 I show some additional evidence that corroborate that the rejection of the FIRE hypothesis is a result of information inefficiencies. To begin with, in Panel A, I show that the constant term is non-zero even after controlling for uncertainty. This implies that the variance of the conditional expectation of the forecast errors remains non-zero even after controlling for the average value of the unforecastable component.

In order to apply the LTV I made the implicit assumption above that only the lagged squared forecast errors in the firm's information set can predict the forecast error variance. Now I provide evidence that this assumption is legitimate. If the firm's

[12] Using fewer lags did not materially change the results.

Table 5.7: Rational Expectations Hypothesis Test: Newey-West Estimates.

	Output
Dep. Var:	$\left(FEDISP_m\right)^2 - \overline{FV}_m$
β_0	0.455
NW SE	0.0135
NW t-stat	33.62
Obs.	195 Months

The table shows the estimation results for equation (5.12). I report the Newey-West standard errors with 12 lags and t-statistics estimated using them.
$FEDISP_m$ is the cross-sectional weighted dispersion of the forecast errors on production; $\overline{FV}_m$ is the weighted cross-sectional average of the forecastable variance of the forecast errors on production.
Results are robust to using either 6 or 2 lags for the Newey-West standard errors (output omitted). Including month fixed effects does not affect the results either (output omitted).

information — which I cannot observe — can predict the forecast error variance and I ignored it from the ARCH estimation, the unforecastable component will have this information. In that case, what my test would capture as information inefficiency would basically be the information that I ignored when I forecasted the variance. The estimates in Panel A of Table 5.8 show that there is still inefficiency even after eliminating the information hidden in the unforecastable component — the constant term is statistically significant at 1%.

Next, in Panel B, I show that the variance of the conditional expectation is not resulting from the variance of the firm specific biases in the forecast errors. The firm-specific bias is the optimism and pessimism that is inherent to each firm's forecasts, that is $\mathbb{E}_i\big[QS_{im}^{fe}\big] = FEBIAS_i \neq 0$, where $\mathbb{E}_i$ is the firm specific mean. Because my panel data is unbalanced with firms (re)entering and (re)exiting the sample quite frequently, it might be the case that the non-zero variance of the conditional expectation that I documented is the result of the appearance and disappearance of the firm-specific forecast biases. To test that, I compute for each firm the average survey-based forecast error that it makes across all months. Then, for each month I calculate he variance of the firm-specific means, $FEBIAS_i$, of the firms that participated in that month. This is the independent variable in Panel B in Table 5.8: $\mathbb{V}\big[FEBIAS_i\big]$. Its coefficient is not statistically significant, which implies that the non-zero variance of the conditional forecast error is not at all the result of the firm-specific optimism or pessimism.

Finally, in Table 5.9 I show that my test results are also robust to using time-series with non-overlapping horizons. As I described, the survey based-forecast errors concern the three-month period ahead. As such, the forecastable variance and the cross-sectional forecast error variance also have a horizon of three months ahead. The rejection of the FIRE hypothesis using equation (5.12) might result from the fact that the monthly time-

Table 5.8: Violation of Information Efficiency: Further time-series results.

Dep. Var: $FEDISP_m - \overline{FV}_m$			
Panel A		Panel B	
$\overline{UNCERT}_m$	0.759***	$\mathbb{V}\left[FEBIAS_i\right]$	0.740
Constant	0.111***	Constant	0.422***
Observations	195	Observations	195
R-squared	0.733	R-squared	0.012

The table shows the estimation results for equation (5.12) with one additional explanatory variable in each column. I report the statistical significance using Newey-West standard errors with 12 lags.
$FEDISP_m$ is the cross-sectional weighted dispersion of the forecast errors on production; $\overline{FV}_m$ is the weighted cross-sectional average of the forecastable variance of the forecast errors on production.
In column (1) the additional independent variable $\overline{UNCERT}_m$ is the monthly cross-sectional weighted average of the unforecastable component of the forecast error variance. In column (2) the independent variable $\mathbb{V}\left[FEBIAS_i\right]$ is the cross-sectional weighted variance of the firm-specific forecast biases for the firms that participated in that month.
*** indicates statistical significance at the 1%, from Newey-West standard errors.
Using 6 lags for the Newey-West standard errors: the significance of the coefficient of $\overline{UNCERT}_m$ does not change; coefficient of $\mathbb{V}\left[FEBIAS_i\right]$ becomes significant at 10% (output omitted).

series $\left(FEDISP_m\right)^2 - \overline{FV}_m$ has overlapping horizons. Intuitively, the inefficient use of information in one month might show in the following two months' forecasts that overlap and this would inflate the estimated inefficiency. To control for this, I split the time-series sample of $\left(FEDISP_m\right)^2 - \overline{FV}_m$ in three sub-samples that do not have overlapping horizons. These sub-samples are: (I) January, April, July, October; (II) February and November (we cannot observe forecasts in August since no surveys are collected and as a result the forecast error made in May cannot be observed either because the realization would be recorded in August); (III) March, June, September, December.

I estimate equation (5.12) distinctly and independently for each of these three sub-samples and I report the results in Table 5.9. I use only 2 lags for the Newey-West standard error because more would imply that the correlation in the error terms of (5.12) extends well in the past. However, results are robust ot using 6 or 12 lags. Evidently, the rejection of the FIRE hypothesis from my test is not the result of the overlapping horizon of the monthly forecast errors. In addition, the estimated magnitude of the inefficiencies across the three sub-samples is close to the one of the whole sample. For the whole sample the estimated magnitude is 0.455 (Table 5.7), while for the three sub-samples it is 0.452, 0.414 and 0.467.

Overall, I construct a test of the FIRE hypothesis that identifies whether the rejection is the result of inefficient use of information and measures the magnitude of the inefficiencies. This test is generally applicable to any survey-based forecast errors and does not rely on assumptions about the origin of the inefficiencies. By testing the firm expectations in my data, I find that they suffer from information inefficiencies. Additionally, I provide further evidence that the information inefficiencies that my test captures is not the result of missing information in forecasting the variance of the forecast errors. I also provide evidence that firm specific optimism and pessimism does not drive

Table 5.9: Rational Expectations Hypothesis Test: Newey-West Estimates for sub-samples with non-overlapping horizons.

Sub-Samples	I	II	III
Dep. Var: $\left(FEDISP_m\right)^2 - \overline{FV}_m$			
β_0	0.452	0.414	0.467
NW SE	0.0154	0.0258	0.0143
NW t-stat	29.37	16.03	32.71
Obs.	87 Months	21 Months	87 Months

The table shows the estimation results for equation (5.12) for three distinct sub-samples: (I) January, April, July, October; (II) February and November (we cannot observe forecasts in August since no surveys are collected and as a result the forecast error made in May cannot be observed either because the realization would be recorded in August); (III) March, June, September, December. I report the Newey-West standard errors with 2 lags and t-statistics estimated using them.
$FEDISP_m$ is the cross-sectional weighted dispersion of the forecast errors on production; $\overline{FV}_m$ is the weighted cross-sectional average of the forecastable variance of the forecast errors on production.
Results are robust to using either 6 or 12 lags for the Newey-West standard errors (output omitted).

the results of my test. Finally, I show that the overlapping three-month horizon of the forecast errors recorded each month does not drive my results either.

5.5 Concluding Remarks

In this section I use firm-level survey-based forecast errors on production from Greece. With this data I show how the forecast error variance can be decomposed into a forecastable and an unforecastable component. My decomposition is an application of the Auto-Regressive Conditional Heteroscedasticity model on the survey-based forecast errors. This decomposition allows me two applications: a novel test of information efficiency in expectations and a measure of micro-uncertainty.

Using the forecastable component I devise a novel test for the full information rational expectations hypothesis. Importantly, the test I propose has four key advantages. First, it rejects the the rational expectations indicating that forecasts suffer from information inefficiencies. Second, it measures the magnitude of these inefficiencies. Third, it does not rely on any assumptions about how the information inefficiencies can be modelled. Fourth, my test is generally applicable to any survey-based forecast errors, qualitative and quantitative, without the need of forecast updates or other additional data. This is particularly important as most surveys do not have forecast updates.

Additionally, I argue that the unforecastable component of the forecast error variance is a new measure of micro-uncertainty. This measure is of firm-month frequency and can be easily obtained from any survey-based forecast errors. To validate my proposed measure I engage in two exercises. First, I show that the micro-uncertainty is negatively correlated to the growth rate of aggregate manufacturing production. Second, I find that larger micro-uncertainty increases the probability that the forecaster makes not only a forecast error, but also a large forecast error. These two findings are consistent with

existing literature findings concerning uncertainty.

Chapter 6

Conclusion

As economists, we know that expectations are a key component in understanding how economies function. Therefore it is important to understand the properties of the expectations of the agents. To improve our understanding of the agents' expectations, in this thesis, I focus on their forecast errors. I study an innovative combination of data comprising of two datasets. First, the survey data that directly records the expectations of the firms. Second, I match the survey responses with the respondents' financial statements. In Chapter 2, I provide a detailed description of the dataset and I give evidence about the quality of the information it provides. I demonstrate that the survey responses show consistency in two ways. First, the survey responses of a firm in a month are on average consistent with each other for the same firm in the same month. Second, responses concerning sales growth in the survey are positively correlated with the sales growth observed in the financial statements. The former also demonstrates that the respondents are aware of the state of affairs in their firm; financial statements for a given year are published the year after the firms post their responses.

In order to thoroughly study the forecast errors of the firms, I need quantitative forecast errors. As I show, qualitative ones are not suitable for this analysis. To obtain

quantified forecast errors, I firstly develop in Chapter 3 a quantification model that allows me to impute quantified forecast errors on sales growth from the survey-based expectations and the realized sales growth from the financial statements. Without this quantification, the distinction between major and minor forecast errors is impossible, and as I demonstrate later this distinction is pivotal. The contribution of my quantification model is that it delivers quantitative annual forecasts at the firm-level.

Studying the firm-level quantified forecast errors on sales growth, I offer some novel insights into the full information rational expectations hypothesis. Namely, I show in Chapter 4 that only major forecast errors show systematic patterns that lead to the rejection of the rational expectations. Minor forecast errors show no systematic patterns. To arrive at this finding, I adapt and implement a dynamic panel threshold estimator that endogenously estimates where the behavior of the forecast errors changes. In fact, I obtain robust estimates that only the forecast errors that are at the upper or lower 26% of their empirical distribution are predictable and autocorrelated.

To rationalize these empirical findings, in the second part of Chapter 4, I provide a simple intuitive model of rational inattention that explains this change in the behavior of the forecast errors. In the presence of information costs, agents optimally choose lower quality signals to use for their forecasts. By paying a higher cost firms can acquire more accurate signals, so there is an optimal point where firm balances the accuracy with the cost. When firms face these information costs, they acquire less accurate signals, which results in larger forecast errors that show systematic behavior. Without these costs, firms can observe the variable of interest and make forecasts more accurate and aligned with the full information rational expectations. These agree with my findings that only major forecast errors show systematic behavior.

In the final chapter, Chapter 5, I focus on the survey-based forecast errors on pro-

duction and show how their variance can be decomposed into a forecastable and an unforecastable component. Using the forecastable component I devise a novel test for the full information rational expectations hypothesis. In fact, this test constitutes a measure of the information inefficiencies in the forecasts of the firms. Importantly, this test is widely applicable to all survey data with minimal assumptions, as long as the forecast errors are observable in the data. Next, I argue that the unforecastable component of the forecast error variance is a new measure of micro-uncertainty. This measure is of firm-month frequency and can be easily obtained from any survey-based forecast errors.

My findings imply that systematic forecast errors can inherently be avoided by improving the quality of the information that the firms use. This can help firms avoid sub-optimal decisions. The question that naturally arises at this point is whether policy makers can take steps to improve the forecasting behavior of the firms. Evidently, this is a question open to future investigation as it requires appropriate models. In addition, future research can focus on examining which firm-level or aggregate variables drive the accuracy of the forecasts of the firms.

Appendix A

Cleaning the Survey Data

IOBE uses imputation techniques for missing monthly responses and for August, a month for which they do not send out surveys. I set to missing all the survey variables of the firm-month observations that were imputed.[1]

Finally, I have set to missing all firm-month observations in one particular year if I have less than three monthly survey responses of this firm within the year. This was necessary because my quantification in Chapter 3 aggregates (and quantifies) the firm-month observation to the firm-year frequency. The informativeness of this aggregation is rather limited when during the year, a firm has responded only once or twice. These cleaning steps leave me with 1,093 firms

[1]Lui et al. (2011) report that for the UK business climate survey, the CBI (who administer the survey) also implements imputation techniques for missing data; they document that, on average, this does not jeopardize the credibility of the survey data.

Appendix B

Cleaning the Financial Statements Data

I have financial statements data available from ICAP. In the following I outline the consecutive steps undertaken to prepare and clean the financial statements database. Prior to these steps this data comprised 1,219 firms with 18,786 firm-year observations in the manufacturing sector. After the cleaning I retained all 1,219 firms and 18,213 firm-year observations.

1. The way the data is recorded, Net Worth is included in Total Liabilities. Therefore, Total Net Assets should equal Total Liabilities, i.e. $TotalNetAssets_{i,y} = TotalLiabilitiess_{i,y}$, for every the firm i, year y. For the firm i-year y observations for which $TotalNetAssets_{i,y} \neq TotalLiabilitiess_{i,y}$, I replaced their values with those from an alternative Balance Sheet data-base of Hellastat S.A.[1,2] I confirmed that for the replaced values of $TotalNetAssets_{i,y}$ and $TotalLiabilitiess_{i,y}$

[1]Non-satisfaction of the accounting identity is entirely due to human error, and since the data providers are different, the person making the error is also different, so I can assume that the two data-bases do not include the same errors.

[2]Hellastat S.A. is a private consultancy firm collecting and digitalizing the financial statements from official and publicly available sources. This database is very similar to my ICAP data, but includes a less detailed break-down of financial statement variables.

the equality holds, and that the net value of the sub-categories included in the Assets sum up to the Total Net Assets. If these variables did not add up, I set to missing all the financial statement variables of these firm-year observations.

2. The following equality should hold:

   $TotalGrossSales_{i,y} = GrossOpertingProfit_{i,y} + CostOfSoldGoods_{i,y}$, for every the firm i, year y. For the observations for which the above equality does not hold, I replaced their values with those from Hellastat. Then I confirmed that for the replaced values of $TotalGrossSales_{i,y}$, $GrossOpertingProfit_{i,y}$ and $CostOfSoldGoods_{i,y}$ the equality holds. If these variables did not add up, I set to missing all the financial statement variables of these firm-year observations.

3. I set to missing all the financial statement variables for the firm-year observations for which the following equality does not hold.

$$
\begin{aligned}
&TotalNetValueOfFixedAssets_{i,y} + TotalAccumulatedDepreciation_{i,y} \\
=&GrossValueOfMachinery\&Equipment_{i,y} + GrossValueOfBuilding\&Facilities_{i,y} \\
&+ GrossValueOfIntangibleAssets_{i,y} + ValueOfLand_{i,y} + ValueOfHoldings_{i,y} \\
&+ ValueOfLongTermReceivables_{i,y}
\end{aligned}
$$

4. For some firm-year observations the NACE classification was the version 1 or its Greek analogue, STAKOD 2003. I used ELSTAT (2002), EUROSTAT (2008a) and EUROSTAT (2008b) to translate all NACE classifications to NACE v. 2.

5. $GrossDepreciablePropertyValue_{i,y}$ is defined as the sum of the Gross Values of Building & Facilities, Machinery & Equipment and Intangible Assets, for every firm i, year y. I set to missing all the financial statement variables for the firm i-year y observations for which at least one of the Gross Depreciable Property,

the Gross Sales, the Total Net Fixed Assets, the Total Net Assets or the Owner's Equity is lower or equal to 0, as this would indicate that the firm was under dissolution in that year.

6. To derive values of Real Total Net Assets, Real Owner's Equity, Real Total Sales I used the annual implicit gross added value deflator (ratio of nominal over real value) from Eurostat Table nama_10_a64 for Greece. To derive Real Total Net Fixed Assets and Real Gross Depreciable Property I used the implicit deflator of capital stocks from Eurostat Table nama_10_nfa_st.

7. In the final cleaning steps, I deal with extreme observations that likely result from miscoding. When the growth rate of any the following variables was at the lower 0.5% of its empirical distribution I set to missing all the financial statement variables: Real Total Net Assets, Real Total Net Fixed Assets, Real Gross Depreciable Property, Real Owner's Equity, Real Total Sales.

8. When the real growth rate of any the following variables was at the upper 1% of its empirical distribution I set to missing all the financial statement variables: Real Total Net Fixed Assets, Real Gross Depreciable Property, Real Total Sales.

Appendix C

Cleaning the Survey Data for the Quantification

The wording of the survey question is so that it asks about sales expectations for the next three months as a whole. This means that expectations posted in the ending months of a year are also concerned with sales in the beginning months of the following year. Similarly, the survey questions about realized sales asks about sales in the previous three months, so that responses at the beginning of the year may include sales developments of the ending months in the previous year. For this reason I make adjustments to the submitted responses on realizations and forecasts in the concerning months, which are standard treatment of survey data in the literature. These adjustments ensure that the annualized weighted fractions do not extend on the following year. For forecasts, I multiply the survey variable with 2/3 in November and with 1/3 in December, as only two thirds and one thirds respectively, of the period over which expectations are recorded, belongs to the current calendar year. For realizations a similar argument applies and I set the responses in January to missing and use this observation with weight 1 in the final month of the preceding year. I further multiply recorded responses by 1/3 in February,

and 2/3 in March. The intuition is that e.g. the response submitted in beginning to mid-February will cover sales realizations that concern November to January and hence only one out of three months included in the response is concerned with the current year. The underlying assumption for this adjustment is that the respondents attach the same weight to the three months covered in their response. This is a standard assumption in the survey literature and implicitly assumed for example in Bachmann et al. (2013) and Massenot and Pettinicchi (2018).

As I already noted in Appendix A for Chapter 2, I have taken the following cleaning steps as well.

- I set to missing all the survey variables of the firm-month observations that were imputed.
- I have set to missing all firm-month observations in one particular year if I have less than three monthly survey responses of this firm within the year.

Appendix D

Detailed Mathematical Derivations

D.1 Derivation of Equation (3.4)

This section shows how equation (3.4) can be derived using equations (3.2) and (3.1). First, I take expectations of equation (3.1), which becomes

$$\mathbb{E}\big[x_{im}^{e,+}|\mathscr{F}_{i,y-1}\big] = \alpha + \gamma_1 x_{iy}^e + \mathbb{E}_{i,y-1}\nu_{im}^+, \quad \text{and} \quad \mathbb{E}\big[x_{im}^{e,-}|\mathscr{F}_{i,y-1}\big] = -\beta + \gamma_2 x_{iy}^e + \mathbb{E}_{i,y-1}\nu_{im}^-. \tag{D.1}$$

Then, I substitute equation (D.1) into (3.2)

$$x_{iy}^e = \mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}^+\big[\alpha + \gamma_1 x_{iy}^e + \nu_{im}^+\big] + \mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}^-\big[-\beta + \gamma_2 x_{iy}^e + \nu_{im}^-\big].$$

Then, using the definition for W_{im}^+ and W_{im}^-, I get

$$\begin{aligned} x_{iy}^e &= \big[\alpha + \gamma_1 x_{iy}^e\big]\mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=1]} + \mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=1]}\nu_{im}^+ \\ &+ \big[-\beta + \gamma_2 x_{iy}^e\big]\mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=-1]} + \mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=-1]}\nu_{im}^- \end{aligned} \tag{D.2}$$

To simplify the notation, I define

$$P_{iy} \triangleq \sum_{m \in y} W_{im} \mathbb{1}_{[XS^e_{im}=1]}, \qquad \text{and} \qquad N_{iy} \triangleq \sum_{m \in y} W_{im} \mathbb{1}_{[XS^e_{im}=-1]},$$

where P_{iy} (N_{iy}) denotes the weighted share of months within a year that indicate a rise (fall) in expected sales. Next, I assume that $\mathbb{E}_{i,y-1} P_{iy} = P_{iy}$ and $\mathbb{E}_{i,y-1} N_{iy} = N_{iy}$, which implies that, during year y, firm i makes as many positive/negative monthly forecasts as it was expecting to make at the end of the preceding year $y-1$. Note that this assumption does not imply perfect foresight whatsoever and is consistent with the costly information updates that are illustrated in equation (3.1). This allows me to rearrange equation (D.2) and solve for x^e_{iy}:

$$x^e_{iy} = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \qquad \text{with} \qquad \xi_{iy} = \frac{\mathbb{E}_{i,y-1} \sum_{m \in y} \left(W^+_{im} \nu^+_{im} + W^-_{im} \nu^-_{im} \right)}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}},$$

which is equation (3.4) in Section 3.2.

D.2 Proofs Related to the Estimation of Equation (3.6)

Statement 1. If $\mathbb{E}\left[x^{fe}_{iy} \middle| x^e_{iy}\right] = 0$, then $\mathbb{E}\left[x^{fe}_{iy} \middle| \mathcal{H}\left(x^e_{iy}\right)\right] = 0$ for any Borel measurable function $\mathcal{H}$. Therefore, $\mathbb{E}\left[x^{fe}_{iy} \middle| \tilde{x}^e_{iy}\right] = 0$ ($\tilde{x}^e_{iy}$ is defined in equation (3.7) in the main text).

Proof. First, note that the underlying mathematical form of this (and any) conditional expectation is $\mathbb{E}\left[x^{fe}_{iy} \middle| x^e_{iy}\right] = \mathbb{E}[x^{fe}_{iy} | \sigma(x^e_{iy})]$, where $\sigma(x^e_{iy})$ is the minimal sigma-algebra generated by x^e_{iy}. Intuitively, all the information that the forecast of the firm x^e_{iy} can convey. Then from the Doob-Dynkin Lemma (see Proposition 3 in Rao and

Swift (2006)) I know that $\sigma\big(\mathcal{H}(x^e_{iy})\big) \subset \sigma(x^e_{iy})$ for any Borel measurable function $\mathcal{H}$. As a result, from the general form of the Law of Iterated Expectations, I get $\mathbb{E}\big[x^{fe}_{iy}\big|\mathcal{H}(x^e_{iy})\big] = \mathbb{E}\Big[\mathbb{E}\big[x^{fe}_{iy}\big|x^e_{iy}\big]\Big|\mathcal{H}(x^e_{iy})\Big] = 0$. Next, I know that $\tilde{x}^e_{iy}$ is a Borel measurable function of XS^e_{im} for $m \in y$.[1] Also, XS^e_{im} is a Borel measurable function of x^e_{iy} (from ID1).[2] Overall, I have that $\sigma(\tilde{x}^e_{iy}) \subset \sigma(\{XS^e_{im}\}_{m\in y}) \subset \sigma(x^e_{iy})$, for $m \in y$. Therefore, $\mathbb{E}\big[x^{fe}_{iy}\big|\tilde{x}^e_{iy}\big] = 0$. This completes the proof.

Statement 2. If $\mathbb{E}\big[\xi_{iy}\big|\{XS^e_{im}\}_{m\in y}\big] = 0$, then $\mathbb{E}\big[\xi_{iy}\big|\mathcal{H}\big(\{XS^e_{im}\}_{m\in y}\big)\big] = 0$ for any Borel measurable function $\mathcal{H}$. Therefore, $\mathbb{E}\big[\xi_{iy}\big|\tilde{x}^e_{iy}\big] = 0$.

Proof. From the Doob-Dynkin Lemma (see Proposition 3 in Rao and Swift (2006)) and t he Law of Iterated Expectations I obtain the first part that $\mathbb{E}\big[\xi_{iy}\big|\{XS^e_{im}\}_{m\in y}\big] = 0$ implies $\mathbb{E}\big[\xi_{iy}\big|\mathcal{H}\big(\{XS^e_{im}\}_{m\in y}\big)\big] = 0$ for any Borel measurable function $\mathcal{H}$ — the proof is the same as that of Statement 1. From the proof of Statement 1 I also know that $\sigma(\tilde{x}^e_{iy}) \subset \sigma(\{XS^e_{im}\}_{m\in y})$. As a result, $\mathbb{E}\big[\xi_{iy}\big|\tilde{x}^e_{iy}\big] = \mathbb{E}\Big[\mathbb{E}\big[\xi_{iy}\big|\{XS^e_{im}\}_{m\in y}\big]\Big|\tilde{x}^e_{iy}\Big] = 0$. This completes the proof.

Statement 3. The error term $\tilde{\xi}_{iy}$ in equation (3.10) is mean-independent of the explanatory variables.

[1]This follows from the fact that $\tilde{x}^e_{iy}$ is a composition of the following three Borel functions: the numerator, the denominator and a function of type $1/f(\cdot)$. The latter, $1/f(\cdot)$, although not continuous it is still Borel measurable. The numerator and the denominator are Borel measurable, because they are continuous functions of XS^e_{im}: they are linear (continuous) functions of P_{iy} and N_{iy} which are also linear functions (continuous) of XS^e_{im}.

[2]This is true because XS^e_{im} is a composition of Borel measurable functions. In ID1, the quantitative monthly forecast, x^e_{im}, is a linear (continuous) function of the x^e_{iy}, hence Borel measurable. Depending on the value of x^e_{im}, then, XS^e_{im} takes the discrete values $\{-1, 0, +1\}$. I can see XS^e_{im} as a composition of indicator functions of x^e_{im}. Indicator functions are Borel measurable.

I provided a way to approximate the unobserved firm heterogeneity, and I derived the final estimable equation (3.10). For equation (3.10), by the same principles as for Statements 1 and 2, it suffices to prove that $\mathbb{E}[\tilde{\xi}_{iy}|\{XS^e_{im}\}_{m\in y}] = 0$. Then, $\tilde{\xi}_{iy}$ is also mean-independent of all the right hand side variables of equation (3.10). This means that the NLS error $\tilde{\xi}_{iy}$ is also mean independent of the rational function on the right hand side of (3.10), which satisfies Davidson and MacKinnon (2004)'s condition for consistency (equation (6.29)). Indeed, from equation (3.11)

$$\begin{aligned}
\mathbb{E}\big[\tilde{\xi}_{iy}|\{XS^e_{im}\}_{m\in y}\big] &= \mathbb{E}\big[x^{fe}_{iy}|\{XS^e_{im}\}_{m\in y}\big] + \mathbb{E}\left[\frac{\omega_i + \vartheta_{iy}}{1-\gamma_1 P_{iy} - \gamma_2 N_{iy}}\middle|\{XS^e_{im}\}_{m\in y}\right] \\
&= 0 + \frac{1}{1-\gamma_1 P_{iy} - \gamma_2 N_{iy}}\mathbb{E}\big[\omega_i + \vartheta_{iy}\big|\{XS^e_{im}\}_{m\in y}\big] \\
&= 0,
\end{aligned}$$

where the terms P_{iy} and N_{iy} 'go outside' the conditional expectation as they are functions of XS^e_{im}, $m \in y$, and therefore $\sigma(\{XS^e_{im}\}_{m\in y})$-measurable. This follows from the Doob-Dynkin Lemma and the standard properties of the conditional expectations. From Statement 1 I have that $\mathbb{E}\big[x^{fe}_{iy}|\{XS^e_{im}\}_{m\in y}\big] = 0$. Note that $\{XS^e_{im}\}_{m\in y} \subset \{XS^e_{im}\}_{m=1,2,\ldots,T_i}$ which implies that $\sigma\big(\{XS^e_{im}\}_{m\in y}\big) \subset \sigma\big(\{XS^e_{im}\}_{m=1,2,\ldots,T_i}\big)$. Therefore, from ID2, ID3 and the Law of Iterated Expectations I have that $\mathbb{E}\big[\omega_i+\vartheta_{iy}\big|\{XS^e_{im}\}_{m\in y}\big] = 0$. This completes the proof.

Appendix E

Robustness

This section shows results based on an alternative weighting scheme used in equation (3.3). In particular, while my baseline weighting controls for seasonalities within the year, I consider as an alternative that all observations are weighted equally per year.

Table 1.E reports results of the estimation of equation (3.10) using the alternative weights. Column (1) shows estimation results for the boom period up to 2008 and column (2) for the following recession. Parameter estimates are very close to the ones in the baseline case shown in Table 3.1. From Table 2.E it is evident that this close resemblance also results in almost identical forecasts. This table shows the distribution of the difference between individual firm-year forecasts based on the baseline weighting and forecasts based on the alternative weighting scheme.

Table 1.E: NLS Estimation of Equation (3.10) with alternative weighting.

	(1)	(2)
Coefficients	Dependent Variable: x_{iy}	
α	0.197**	0.105**
β	0.158*	0.243***
γ_1	-0.426	-0.479
γ_2	-0.236	0.0492
Firm-Year Observations	2,471	1,397
R^2	0.043	0.057
Period	$y \leq 2008$	$y > 2008$

Fixed effects proxies of equation (3.10) are omitted – but are included in the estimation – to maintain a simple representation. I use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance. Column (1) shows results for the boom period up to 2008 and column (2) for the following recession.

Table 2.E: Distribution of the difference between the baseline forecasts and forecasts based on alternative weighting.

Min	5%	10%	25%	Median	Mean	75%	90%	95%	Max
-0.008	-0.003	-0.002	-0.001	0	0	0.001	0.002	0.003	0.011

Appendix F

Alternative Quantification Techniques

Ordered response models — probit or logit — are alternatives to the NLS based method outlined in Section 3.2 to quantify sales growth forecasts.

The ordered response models assume that there is an unobserved latent variable $XS_{im}^{e,*}$ which defines the outcome of the observed survey response, XS_{im}^{e}, as follows

$$XS_{im}^{e} = -1 \text{ if } XS_{im}^{e,*} \leq a_1,$$

$$XS_{im}^{e} = 0 \text{ if } a_1 < XS_{im}^{e,*} \leq a_2,$$

$$XS_{im}^{e} = +1 \text{ if } XS_{im}^{e,*} > a_2,$$

with $a_1, a_2 \in \mathbb{R}$ being the unobserved threshold parameters. Now assume that $XS_{im}^{e,*}$ is linearly determined by a vector of explanatory variables, $XS_{im}^{e,*} = \delta X_{im}^{XS} + \psi_i + e_{im}$, with ψ_i being the unobserved firm heterogeneity and e_{im} the idiosyncratic error term. The assumed distribution of e_{im} determines whether the model will be standard normal (probit) or logistic (logit). The explanatory variables X_{im}^{XS} can be from both the survey and the financial statements. I can eliminate the unobserved heterogeneity ψ_i using the

Mundlak (1978) approximation, that is the cross-time firm-specific averages of all the panel dependent variables $\psi_i \approx \frac{1}{T_i}\sum_m^{T_i} X_{im}^{XS}$, where T_i is the number of months each firm i is present in the sample.

After accounting for unobserved firm heterogeneity in the ordered response models, I can get (maximum likelihood) consistent and unbiased estimations of $\hat{\delta}$ and compute the estimated latent variable values, $\hat{XS}_{im}^{e,*}$. These will be the quantified values of the survey variable, $\hat{XS}_{im}^{e}$ That is $\hat{XS}_{im}^{e} = \hat{XS}_{im}^{e,*} = \hat{\delta} X_{im}^{XS}$. The estimated $\hat{XS}_{im}^{e}$, are the quantified value of the firm's monthly response conditional on X_{im}^{XS}.[1] Finally, I can derive their annualized quantified values using the weighted average $\hat{x}_{iy}^{e} \triangleq \sum_{m\in y} W_{im}[\hat{XS}_{im}^{e}]$, using the weights given in equation (3.3).

Table 1.F reports the estimation results of the ordered probit and logit models. The variables that I have used in the vector of explanatory variables, X_{im}^{XS}, are (i) $\overline{XS}_m^e = (N_m)^{-1}\sum_i XS_{im}^e$, where N_m is the number of firms that responded in month m. This will capture aggregate time-specific effects and aggregate information. (ii) the growth rate of sales in the preceding year, $x_{i,y-1}$, from the financial statements (iii) $ORDS_{im}$ which is a categorical variable from the survey indicating the level of orders.[2] Because the survey questions are almost entirely focused on the supply side, I could not find any other questions that one can reasonably expect to be correlated with sales forecasts.

[1] An alternative would be to obtain the probability estimates for each possible response, $XS_{im} = -1/0/+1$, and then compute the mean response. But that would require to use $\frac{1}{T_i}\sum_m^{T_i} X_{im}^{XS}$ for the mean response, because the estimated cut-off values, a_1, a_2, are conditional on all explanatory variables, including the fixed effects specification. The problem with using $\frac{1}{T_i}\sum_m^{T_i} X_{im}^{XS}$ for the estimation is that I would introduce information to the firm's forecasts that were not available to the firm at the time of the forecast.

[2] It is based on question B.1 *'Your total orders outstanding (from either domestic or foreign markets) you deem, for this period of the year, to be high/normal/low.'*

Table 1.F: Ordered Probit and Logit Estimations of firm-month survey responses on sales growth forecasts

	Probit			Logit	
	(1)	(2)		(2)	(3)
Period	y≤2008	y>2008	Period	y≤2008	y>2008
	Dep. Variable: XS^e_{im}			Dep. Variable: XS^e_{im}	
$\overline{XS}^e_m$	1.631***	1.565***	$\overline{XS}^e_m$	2.784***	2.666***
$x_{i,y-1}$	0.0836**	0.0106	$x_{i,y-1}$	0.141**	0.0176
$ORDS_{im}$	0.358***	0.372***	$ORDS_{im}$	0.615***	0.645***
a_1	-0.958***	-1.211***	a_1	-1.620***	-2.023***
a_2	0.583***	0.366***	a_2	0.994***	0.620***
Observations	13,554	8,740	Observations	13,554	8,740
Pseudo-R^2	0.0575	0.0750	Pseudo-R^2	0.0561	0.0751

Fixed effects specification are omitted – but are included in the estimation – to maintain a simple representation. $\overline{XS}^e_{im}$ is the cross-sectional monthly average of the sales forecast reported based on survey question D.2, $x_{i,y-1}$ is the growth rate of sales in the preceding year, from the from the financial statements, $ORDS_{im}$ indicates the level of orders based on survey question B.1. I use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance.

Appendix G

Accuracy of the Quantification Methodology — A Monte Carlo Exercise

In this Appendix, I describe how artificial data is generated and subsequently used to evaluate the precision of my methodology to quantify qualitative forecasts. I first document details of the data generating process and its calibration. Finally, I discuss results that stress the robustness of the evidence shown in Table 3.6.

Generating Artificial Data

The following outlines how I generate artificial data on firm's (continuous) sales growth, z_{iy}, as well as corresponding qualitative expectations, ZS^e_{iy}, and quantitative expectations, z^e_{iy}. The realized sales growth and the qualitative expectations are then used as inputs to the quantification methodology in Section 3.2 to generate estimates for quantified sales growth expectations, $\hat{z}^e_{iy}$. This allows me to evaluate the accuracy between these estimates, $\hat{z}^e_{iy}$, and the actual underlying expectations, z^e_{iy}.

My dataset on Greek firms' sales growth is an unbalanced panel with 799 firms, 4,104 firm-year observations and 25,764 firm-month observations that spans 18 years. The final artificial datasets that I generate exactly matches this structure. I further take into account that the first eleven years in my sample were a boom period and the last seven years a severe bust. I start with generating a balanced panel that spans 20 years, where the first two years are used to inform lagged values. I now document how each of the three artificial variables is generated.

First, I generate artificial data for firm's sales growth, z_{iy}, based on an AR(1) process. I use the MA(∞) representation

$$z_{iy} = \sum_{l=0}^{y-1} \theta^l (\varepsilon_{i,y-l} + \varpi_i), \quad \text{for } y > 1; \qquad \text{and} \qquad z_{iy} = \varepsilon_{i0} + \varpi_i \quad \text{for } y = 1.$$

This is guided by the evidence in Section J.1 (Table 1.J) that this process explains the data well.[1] The innovations $\varepsilon_{iy} \sim N\big((1-\theta)\mu, (1-\theta^2)\sigma^2\big)$ are i.i.d. and $\varpi_i \sim N(0, \sigma^2_{\varpi_i})$ is unobserved firm heterogeneity.

Second, I generate firms' annual quantitative sales growth forecasts based on the process

$$z^e_{iy} = (1-\theta)\mu + \theta z_{i,y-1} + \varpi^e_i + \varepsilon^e_{i,y-1},$$

where ϖ^e_i is the unobserved firm heterogeneity, which can be seen as firm-specific degree of optimism or pessimism. The innovations $\varepsilon^e_{iy} \sim N(0, \sigma^2_{\varepsilon^e_{iy}})$ are i.i.d. and capture any additional information the firm might include in its forecast. There is no means of inferring the underlying process for expectation formation from the data. However, since realized sales growth in the data is well explained by an AR(1) process, it seems

[1] Using an AR(2) process to generate the artificial data does not materially affect the performance of my quantification methodology. Results are discussed below and shown in Table 2.G.

likely that such a process is also used by firms to form expectations.

Third, I generate the qualitative monthly expectations, ZS^e_{im}. These expectations need to correspond to the annual quantitative forecasts z^e_{iy}. For this reason, I first generate firms' monthly quantitative forecasts, z^e_{im}, and map these into qualitative categories (decline/unchanged/increase) in a second step. Firms' monthly quantitative expectations, conditional on their forecast for the whole year, are generated as

$$z^e_{im} = \mu + \gamma z^e_{iy} + \varepsilon^e_{im},$$

where $\varepsilon^e_{im} \sim N(0, \sigma^2_{\varepsilon^e_{im}})$ are i.i.d. and capture any additional information that the firm includes in its forecast. Note that this procedure to link the artificial annual and monthly observations derives closely from Pesaran (1987).

The only purpose for which the quantitative monthly expectations z^e_{im} have been generated, is to match these into three categories (*decline*/*unchanged*/*increase*) to derive qualitative monthly expectations, ZS^e_{im}. This mapping is constructed so that resulting proportions of observations in the three categories correspond to the proportion of *decline* responses, $C^-\%$, and the proportion of *increase* responses, $C^+\%$, in my survey data. In particular, I assign $ZS^e_{im} = 1$ for the largest $C^+\%$ of values in z^e_{im}; and $ZS^e_{im} = -1$ for the smallest $C^-\%$ of values in z^e_{im}. Since the percentage share of *unchanged* observations in the survey data equals 100-$C^+\%$-$C^-\%$, for the remaining observations in the middle of the distribution of z^e_{im} I set the corresponding $ZS^e_{im} = 0$.

Finally, for the three variables based on artificial data — z_{iy}, z^e_{iy} and ZS^e_{im} — I drop the appropriate observations so that I derive an unbalanced panel of artificial data that exactly corresponds to the structure of firm-year-month observations in my observable dataset.[2] I repeat the steps above to generate 1,000 random samples of artificial datasets.

[2]Prior to this, I have also dropped all observations of the first two years which had only been

Then, for each sample, I use z_{iy} and ZS^e_{im} as input to my quantification methodology and compare the resulting estimate for quantitative sales growth expectations, $\hat{z}^e_{iy}$, with the true underlying expectations, z^e_{iy}.

Calibration

To generate the artificial data I need to calibrate a number of parameters. This exercise is closely informed by my financial statements data on annual sales growth realizations and the survey data on monthly qualitative expectations. Based on the estimates reported in Table 1.J in Chapter 4, I set the autocorrelation coefficient in the AR(1) process for artificial sales growth, z_{iy}, to $\theta = -0.1$. The parameters μ and σ, that govern the moments of the corresponding innovations, are calibrated to match the respective moments in my sales growth data from the financial statements. Since particularly the mean differs across the boom and bust periods in my sample, I differentiate between these episodes and set $\mu = 0.077$ ($\mu = -0.059$) and $\sigma = 0.391$ ($\sigma = 0.401$) during the boom (bust) period.[3] The standard deviation of the unobserved firm heterogeneity, σ_{ϖ_i}, is set to 0.129 to match the standard deviation of the firm-specific cross-time average of sales growth in the financial statements data.

Since the artificially generated qualitative and quantitative expectations variables are linked, I jointly calibrate the remaining parameters that correspond to these variables to match a number of statistics in my data. I first discuss the parameters that govern the process for annual sales growth expectations. The firm specific optimism/pessimism, ϖ^e_i, should be related to the average firm-specific performance, ϖ_i. I scale $\varpi^e_i = 0.5 \cdot \varpi_i$ so that the standard deviation of the firm-specific average of the artificial monthly

employed to inform values of lagged variables.

[3] Apart from the mean μ, and the shares C^+ and C^-, the statistics used to calibrate the parameters in this section are very similar across boom and bust episodes which is why I refrain from a differentiation for these parameters.

qualitative expectations is close to the corresponding statistic in the observable dataset (0.431 vs. 0.422). The standard deviation of the innovation, $\sigma_{\varepsilon^e_{iy}} = 0.02$, is calibrated so that the standard deviation of the firm-year averages of the monthly qualitative expectations in the artificial data will be close to the one in the observable data (0.515 vs. 0.478).

Next, I turn to the remaining parameters required to generate the monthly expectations. The standard deviation of the innovations, $\sigma_{\varepsilon^e_{im}}$, is set to 0.05, based on the within-year variation of the monthly qualitative survey responses. I measure this variation as the arithmetic mean of the squared difference between the monthly survey responses and their firm-year average (0.211 in the artificial data vs. 0.259 in the survey responses). The parameter γ is calibrated to 0.8 so that the correlation between realized annual sales growth and the qualitative monthly expectation responses in the artificial data matches the corresponding correlation in my observable dataset.[4]

Table 1.G: Calibrated parameters to generate artificial data

Parameter	Value	Matched Moment from Financial Statements (FS) or Survey Data
μ	0.077 (-0.059)	Mean in boom (bust) period of sales growth from financial statements
σ	0.391 (0.401)	Standard deviation in boom (bust) period of sales growth from financial statements
θ	-0.1	Autocorrelation estimates (see Table 1.J) of sales growth from financial statements
σ_{ϖ_i}	0.129	Standard deviation of firm-specific cross-time average of sales growth in the FS
ϖ^e_i	$0.5\sigma_{\varpi_i}$	Scaled to match std. dev. of firm-specific average of monthly qual. survey expectations
$\sigma_{\varepsilon^e_{iy}}$	0.02	Std. dev. of the firm-year averages of the monthly qualitative survey expectations
$\sigma_{\varepsilon^e_{im}}$	0.05	Mean of squared difference between monthly survey responses and their firm-year average
γ	0.8	Correlation: annual sales growth from FS and qualitative monthly survey expectations
C^+	38% (24%)	Percentage share of positive monthly responses in the survey data during boom (bust)
C^-	11% (24%)	Percentage share of negative monthly responses in the survey data during boom (bust)

All calibrated parameters and the moments I target are summarized in Table 1.G. My calibration strategy carefully ensures close correspondence of the artificially generated data with my observable dataset. This is achieved by matching statistics that concern,

[4]In particular, I run the regression $XS^e_{im} = \beta_0 + \beta_1 x_{iy} + \phi_i + \eta_{im}$ where ϕ_i controls for firm fixed effects and η_{im} is an idiosyncratic error.

amongst others, relations between qualitative survey expectations and quantitative realizations, as well as monthly and annual data. I now evaluate the appropriateness of the calibration and the assumptions on underlying processes by evaluating how well the artificial data conforms to statistics in the observable data that are not targeted. I document three such statistics. First, for the error of the regression of monthly qualitative forecasts on annual sales growth realizations, the unobserved firm heterogeneity accounts for 35% of its variance in the artificial data vs. 33% in the dataset that comprises information from the survey and the financial statements.[5] Second, the coefficient of the regression of annualized survey responses on sales growth realizations is 0.169 in the artificial data vs. 0.193 in the observed data.[6] Third, in the error term of the latter regression (ZS^e_{iy} on z_{iy}) unobserved firm heterogeneity accounts for 56% of its variance vs. 58% in the observed data. The close correspondence between artificial and observed data in all three statistics is reassuring about the adequacy of my calibration. The second statistic particularly corroborates my calibration of γ, while the first and third statistics support my calibration of the variance of the unobserved firm heterogeneity.

Alternative Data Generating Process

Table 3.6 in the main body demonstrates, based on artificial data, a close correspondence between the estimated and the true quantitative forecast errors. The artificial data on sales growth has been generated based on the above AR(1) process. I now demonstrate robustness to an alternative Data Generating Process. I relax the AR(1) assumption and generate sales growth based on the AR(2) process $z_{iy} = 1.2\mu - 0.2z_{i,y-1} - 0.1z_{i,y-1} + \varepsilon_{iy} + \varpi_i$. Note that the sales growth expectations are still generated based on the process

[5] Using the notation for the artificial variables the regression is: $ZS^e_{im} = \beta_0 + \beta_1 z_{iy} + \phi_i + \eta_{im}$. The corresponding variables in my empirical dataset have been denoted XS^e_{im} on x_{iy} in the main body.

[6] Using the notation for the artificial variables the regression is $ZS^e_{iy} = \beta'_0 + \beta'_1 z_{iy} + \phi'_i + \eta'_{iy}$. Where ZS^e_{iy} is the firm-year arithmetic mean of the monthly survey responses.

Table 2.G: Distribution of the difference between the estimated quantitative forecast error and the true quantitative forecast error — alternative data generation

5%	10%	25%	Median	Mean	75%	90%	95%
-0.085	-0.065	-0.033	0.000	-0.002	0.031	0.058	0.075
(0.011)	(0.010)	(0.009)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)

I report the average across 1,000 sets of artificial data of the descriptive statistics. Standard deviations across the 1,000 sets for these statistics are reported in parenthesis. Sales growth realizations are generated based on an AR(2) process.

shown earlier, which introduces predictability and autocorrelation in the forecast errors. Table 2.G shows the distribution of the difference between the estimated quantitative forecast error and the true quantitative forecast error when the artificial data is generated based on the AR(2) process. Overall, results are robust to this change, and the estimated forecast errors still correspond closely to the underlying true forecast errors.

Appendix H

Robustness Checks of the Predictability of forecast Errors

First, in Table 1.H, I show alternative estimations of the predictability equation without the thresholds; i.e. linear equation (4.1). In Column (1), the GMM system is just identified without over-identifying restrictions. In column (2) I estimated equation (4.1) using the biased LSDV, following the standard practice of the literature.

Second, in Table 2.H, I show alternative estimations of the nonlinear equation (4.2) with the threshold. In column (1), I estimate the threshold regression with the biased LSDV classifying the major forecast errors based on the baseline threshold estimate of 26%. In column (2) of Table 2.H I use my baseline estimation strategy but with fewer lags, that is without any over-identifying restrictions. Column (3) is with the original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$. In column (4), I estimate the threshold regression with the biased LSDV classifying the major forecast errors based on the baseline threshold estimate of column (3).

Overall, as I discuss in detail in Section 4.1.3 my baseline results are robust to using

Table 1.H: Predictability of firms' forecast errors of sales growth – Robustness Checks for the Specification without Threshold.

	(1)	(2)
Estimation	FOT	LSDV
Stand. Errors	2-step, Windmeijer (2005) corrected	Robust
Lags as Instruments	2	N.A.
	Dependent Variable: Sales Growth Forecast Error, x_{iy}^{fe}	
$x_{i,y-1}$	-0.158***	-0.221***
$\overline{x}_{IND,y}$	0.827***	0.830***
Constant		-0.0142***
Observations	2,805	3,559
# of Firms	590	754
Over-identified	No	N/A
Hansen p-value	N.A.	N.A.
m2 test p-value	0.882	N.A.

Table shows alternative estimations of equation (4.1) without the threshold. Column (1) is estimated with the Arellano and Bover (1995) FOT GMM; column (2) with the LSDV. In column (1), the instruments are with only one lag dated in $y-2$ and collapsed. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate annual effects with the NACE two-digit industry, IND, year average of sales growth from the entire sample of the financial statements, $\overline{x}_{IND,y}$. x_{iy}^{fe} is the forecast error of sales growth for year y; $x_{i,y-1}$ is the lagged realized sales growth. *** indicates statistical significance at the 1% level, respectively.

shorter lag length as instruments. Moreover, my threshold estimates are also robust to using the original Seo and Shin (2016) threshold estimator with FD GMM. Finally, I provided evidence that using the original estimator with my data does not sufficiently correct the Nickell (1981) bias in the coefficients.

Table 2.H: Predictability of firms' sales growth forecast errors – Robustness Checks for the Threshold Specifications.

	(1)	(2)	(3)	(4)
Estimation	LSDV	FOT	FD	LSDV
Stand. Errors	Robust	2-step, Windmeijer corrected		Robust
Lags as Instruments	N.A.	2	2-6	N.A.
Estimated Threshold q	^{P}26%	28%	27%	^{P}27%
	Dependent Variable: Sales Growth Forecast Error, x_{iy}^{fe}			
$x_{i,y-1} * (1 - MAJ_{i,y-1}^q)$	-0.162*	-0.0936	-0.0941	-0.166*
$x_{i,y-1} * MAJ_{i,y-1}^q$	-0.236***	-0.157***	-0.198**	-0.235***
$MAJ_{i,y-1}^q$	-0.000848	-0.0279	-0.0224	-0.00643
$\overline{x}_{IND,y}$	0.826***	0.824***	0.852***	0.825***
Constant	-0.0159**	–	–	-0.0129*
Observations	2,643	2,069	1,915	2,643
# of Firms	574	432	423	574
Over-identified	N.A.	No	Yes	Yes
Hansen p-value	N.A.	N.A.	0.901	N.A.
m2 tes pt-value	N.A.	0.976	0.656	N.A.

Instruments in all specifications are collapsed; ^{P} indicates pre-estimated threshold cut-off value. The table shows alternative estimations of equation (4.2). Columns (1) and (4) are estimated using the biased LSDV for pre-estimated threshold cut-off values. For Columns (1) I used the estimated threshold from the baseline specification of Table 4.1; for (4) the threshold is estimated in column (3). Column (2) is the adapted Dynamic Panel Threshold estimator using the Arellano and Bover (1995) FOT GMM without any over-identifying restrictions; instruments lagged at $y-2$ and collapsed. Column (3) is with the original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate annual effects with the NACE two-digit industry, IND, year average of sales growth from the entire sample of the financial statements, $\overline{x}_{IND,y}$. x_{iy}^{fe} is the forecast error of sales growth for year y; $x_{i,y-1}$ is the lagged realized sales growth. MAJ_{iy} takes value one when the forecast error lies at the lower or upper q% of its empirical pool distribution. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level, respectively.

Appendix I

Robustness Checks for the autocorrelation of the forecast Errors

First, in Table 1.I, I show alternative estimations of the autocorrelation equation without the thresholds; i.e. linear equation (4.3). In Column (1), the GMM system is just identified without over-identifying restrictions. In column (2) I estimated equation (4.3) using the biased LSDV, following the standard practice of the literature.

Second, in Table 2.I, I show alternative estimations of the nonlinear equation (4.4) with the threshold. In column (1), I estimate the threshold regression with the biased LSDV classifying the major forecast errors based on the baseline threshold estimate of 26%. In column (2) of Table 2.I I use my baseline estimation strategy but with fewer lags, that is without any over-identifying restrictions. Column (3) is with the original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$. In column (4), I estimate the threshold regression with the biased LSDV classifying the major forecast errors based on the baseline threshold

Table 1.I: Autocorrelation of firms' forecast errors of sales growth – Robustness the Specification without Threshold.

	(1)	(2)
Estimation	FOT	LSDV
Stand. Errors	2-step, Windmeijer (2005) corrected	Robust
Lags as Instruments	2	N.A.
	Dependent Variable: Sales Growth Forecast Error, x_{iy}^{fe}	
$x_{i,y-1}^{fe}$	-0.163***	-0.238***
$\overline{x}_{IND,y}$	0.817***	0.809***
Constant	–	-0.0206***
Observations	2,069	2,643
# of Firms	432	574
Over-identified	No	N.A.
Hansen p-value	N.A.	N.A.
m2 test p-value	0.900	–

Table shows alternative estimations of equation (4.3) without the threshold. Column (1) is estimated with the Arellano and Bover (1995) FOT GMM estimator; column (2) shows estimates based on the LSDV. Column (1) uses collapsed instruments with lags in $y-2$ (not over-identified). The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate annual effects with the NACE two-digit industry, IND, year average of sales growth from the entire sample of the financial statements, $\overline{x}_{IND,y}$. x_{iy}^{fe} is the forecast error of sales growth for year y. *** indicate statistical significance at the 1% level, respectively.

estimate of column (3).

Overall, as I discuss in detail in Section 4.2.3 my baseline results are robust to using shorter lag length as instruments. Moreover, my threshold estimates are also robust to using the original Seo and Shin (2016) threshold estimator with FD GMM. Finally, I provided evidence that using the original estimator with my data does not sufficiently correct the Nickell (1981) bias in the coefficients.

Table 2.I: Autocorrelation of firms' forecast errors on sales growth – Robustness for the Threshold Estimation.

	(1)	(2)	(3)	(4)
Estimation	LSDV	FOT	FD	LSDV
Stand. Errors	Robust	2-step, Windmeijer corrected		Robust
Lags as Instruments	N.A.	2	2-6	N.A.
Estimated Threshold q	^{P}26%	20%	26%	^{P}26%
	Dependent Variable: Sales Growth Forecast Error, x_{iy}^{fe}			
$x_{i,y-1}^{fe} * (1 - MAJ_{i,y-1}^{q})$	-0.145	0.0494	0.304*	-0.145
$x_{i,y-1}^{fe} * MAJ_{i,y-1}^{q}$	-0.240***	-0.168***	-0.217***	-0.240***
$MAJ_{i,y-1}^{q}$	-0.00256	-0.0234	-0.0338	-0.00256
$\overline{x}_{IND,y}$	0.808***	0.813***	0.867***	0.808***
Constant	-0.0179***			-0.0179***
Observations	2,643	2,069	1,915	2,643
# of Firms	574	432	423	574
Over-identified	N.A.	No	Yes	N.A.
Hansen p-value	N.A.	N.A	0.955	N.A.
m2 test p-value	N.A.	0.943	0.535	N.A.

Instruments in all specifications are collapsed; ^{P} indicates pre-estimated threshold cut-off value. The table shows alternative estimations of equation (4.4). Columns (1) and (4) are estimated using the biased LSDV for pre-estimated threshold cut-off values. For Columns (1) I used the estimated threshold from the baseline specification of Table 4.2; for (4) the threshold is estimated in column (3). Column (2) is the adapted Dynamic Panel Threshold estimator using the Arellano and Bover (1995) FOT GMM without any over-identifying restrictions; instruments lagged at $y-2$ and collapsed. Column (3) is with the original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate annual effects with the NACE two-digit industry, IND, year average of sales growth from the entire sample of the financial statements, $\overline{x}_{IND,y}$. x_{iy}^{fe} is the forecast error of sales growth for year y. MAJ_{iy} takes value one when the forecast error lies at the lower or upper q% of its empirical pool distribution. ***, ** and * indicates statistical significance at the 1%, 5% and 10% level, respectively.

Appendix J

Additional Empirical Results

J.1 Autocorrelation of Sales Growth

In Table 1.J I report estimates for the autocorrelation of sales growth. In the first column, to eliminate firm fixed effects, I use the biased LSDV estimator. In the other fmy columns, I use the Arellano and Bover (1995) Two Step Forward Orthogonal Deviations GMM (FOT). I use distinct number of lags (for instruments) for robustness (see Roodman (2009), Caselli and Tesei (2016)). Additionally, because of the small number of firms (relatively to the moment conditions) I collapse the instruments and I use the Windmeijer (2005) corrected standard errors (Roodman (2009), Caselli and Tesei (2016)). Finally, for the realizations, I use the first differences as instruments as the instruments in levels indicated autocorrelation in the error. Table 1.J shows that annual real sales growth from the financial statements has a negative autocorrelation, and the estimated coefficient is robust to different lag lengths. Barrero (2019) also find evidence of negative serial autocorrelation of realized sales growth for the US data. Moreover, the autocorrelation coefficient of the FOT estimator is higher than that of the LSDV. This is to be expected as the latter is negatively biased for samples with finite time dimension

(see e.g. Pesaran (2015)).

Table 1.J: Autocorrelation of firms' realized sales growth

	(1)	(2)	(3)	(4)
Estimation	LSDV	FOT		
Stand. Errors	Robust	2-step, Windmeijer corrected		
Lags as Instruments	N.A.	2-11	2-6	2-4
	Dependent Variable: Sales Growth, x_{iy}			
$x_{i,y-1}$	-0.122***	-0.0995***	-0.103***	-0.0997***
Constant	0.260***	–	–	–
Observations	15,211	13,994	13,994	13,994
# of Firms	1,217	1,214	1,214	1,214
Over-identified	N.A.	Yes	Yes	No
Hansen p-value	N.A.	0.251	0.0369	N.A.
m2 test p-value	N.A.	0.553	0.617	0.549

Column (1) is with the standard fixed effects (LSDV); (2), (3) and (4) are the Arellano and Bover (1995) 2-Step Forward Orthogonal Deviations GMM (FOT). y fixed effects are included in all estimations, but are omitted. In (2)-(4), I use distinct number of lags (for instruments) for robustness, all are collapsed. The instruments are lagged first differences of the right hand side variable dated as indicated. The Arellano-Bond p-value (m2 test) shows no serial correlation of order 2 in the errors. x_{iy} is the sales growth observed from the financial statements. ***, ** and * indicates statistical significance at the 1%, 5% and 10% level, respectively.

Appendix K

Equation Derivations

K.1 Derivation of Equation (4.13)

To derive equation (4.13) for the optimal choice of attention, I begin from the original problem

$$\max_{\lambda} \Big[\mathbb{E}U(\lambda) - C(\lambda)\Big]. \tag{K.1}$$

and I follow Gabaix (2014). I take the Taylor expansion of $U(\lambda)$ around the rational expectations solution, $\lambda = 1$,[1]

$$U(\lambda) - U(1) = \frac{\partial U}{\partial \lambda}\bigg|_{\lambda=1} \big(\lambda - 1\big) + \frac{1}{2}\frac{\partial^2 U}{\partial \lambda^2}\bigg|_{\lambda=1} \big(\lambda - 1\big)^2 + o(\lambda^3), \tag{K.2}$$

where $o(\lambda^3) = 0$, because the utility is quadratic, so higher order derivatives with respect to λ are zero. $U(\lambda)$ is given by equation (4.11), so that for the derivatives in equation (K.2) I need to calculate $\partial x^e_{y+1}(\lambda)/\partial\lambda$. Before I proceed, I introduce some useful notation. The utility function has the general form: $U(A, B) = -\frac{1}{2}(A - B)^2$. Then, I

[1]Even though the utility function is quadratic, I cannot directly analytically solve equation (K.1), because of the presence of term $x^e_{y+1}(\lambda)$ which is unknown without knowing the choice for λ. However, with the Taylor expansion around $\lambda = 1$, this term reduces to $x^e_{y+1}(1)$ which is the known rational expectations solution.

can define the trivial derivatives $U_1 \triangleq \partial U/\partial A$, $U_2 \triangleq \partial U/\partial B$, $U_{11} \triangleq \partial^2 U/\partial A^2 = -1$, $U_{22} \triangleq \partial^2 U/\partial B^2 = -1$ and $U_{12} \triangleq \partial^2 U/\partial A \partial B = 1$.

Recall that $x^e_{y+1}(\lambda) \triangleq x^e_{y+1}\big(\lambda s_y\big) = \arg\max_{x_{y+1}} U(x_{y+1}, \lambda s_y)$. The first order condition implies $U_1\big(x^e_{y+1}(\lambda), \lambda s_y\big) = 0$. Therefore, I can use the implicit function theorem on the first order condition and obtain

$$\frac{\partial x^e_{y+1}(\lambda)}{\partial \lambda s_y} = -\frac{U_{12}}{U_{11}} = 1, \quad \forall \lambda.$$

Subsequently:

$$\frac{\partial x^e_{y+1}(\lambda)}{\partial \lambda} = \frac{\partial x^e_{y+1}(\lambda)}{\partial \lambda s_y}\frac{\partial \lambda s_y}{\partial \lambda} = -\frac{U_{12}}{U_{11}} s_y = s_y, \quad \forall \lambda.$$

I can now calculate the partial derivatives of the Taylor polynomial (K.2). First, for the first order term:

$$\frac{\partial}{\partial \lambda} U\left(x^e_{y+1}(\lambda), \lambda s_y\right) = U_1 \frac{\partial x^e_{y+1}(\lambda)}{\partial \lambda} + U_2 \frac{\partial \lambda s_y}{\partial \lambda} = U_2 s_y, \quad \forall \lambda,$$

because $U_1 = 0$ at the optimum (recall that I are working with the indirect utility). Next, for the second order term:

$$\frac{\partial^2}{\partial \lambda^2} U\left(x^e_{y+1}(\lambda), \lambda s_y\right) = U_{21} \frac{\partial x^e_{y+1}(\lambda)}{\partial \lambda} s_y + U_{22} \frac{\partial \lambda s_y}{\partial \lambda} s_y = -s_y^2, \quad \forall \lambda,$$

because, the cross-partial derivatives of the indirect utility are zero at the optimum, $U_{21} = 0$, and $U_{22} = -1$.

Substituting these results of the Taylor expansion into the maximization problem of equation (K.1), I obtain

$$\max_{\lambda}\left\{-\mathbb{E}\left[\frac{1}{2}s_y^2(\lambda-1)^2\right]-C(\lambda,c_y)\right\}.$$

This result follows from the fact that $U_2\big|_{\lambda=1}=0$ and $U(1)=0$.

Finally, using the fact that $\mathbb{E}s_y=\mathbb{E}x_y=0$ and that $\mathbb{E}\epsilon_y x_y=0$, $\forall y$, I have that $\mathbb{E}s_y^2=\sigma_s^2=\sigma_x^2+\sigma_\epsilon^2$. This results in equation (4.13) in the main body.

Appendix L

Test for ARCH(2) in the Survey-Based Forecast Errors

I test for a second order ARCH process in the survey forecast errors by including one additional lag in equation (5.2), that is $\left(QS^{fe}_{i,m-6} - \overline{QS^{fe}_{m-6}}\right)^2$. I estimate this ARCH(2) model using the Arellano and Bover (1995) Forward Orthogonal Transformations, using only 1 lag as instrument (collapsed) for each of the independent variables, with Windmeijer (2005) corrected standard errors. I discuss the details of my estimation strategy in the main body, in Section 5.2.1. As I show in Table 1.L, the estimated coefficient of this additional lag is not statistically significant. This validates my assumption that the survey based forecast errors follow an ARCH process of order 2.

Table 1.L: Estimated ARCH Model or Order 2 for the Survey-based Forecast Errors on Production.

Estimation	FOT
Stand. Errors	2-step, Windmeijer corrected
Lags as Instruments	$m-4$ and $m-7$
Dependent Variable: $\left(QS^{fe}_{i,m} - \overline{QS^{fe}_{m}}\right)^2$	
$\left(QS^{fe}_{i,m-3} - \overline{QS^{fe}_{m-3}}\right)^2$	0.0601***
$\left(QS^{fe}_{i,m-6} - \overline{QS^{fe}_{m-6}}\right)^2$	0.0240
$\hat{\psi}_m$	0.962***
Constant	N.A.
Observations	9,387
# of Firms	539
Over-identified	No
Hansen p-value	N.A.
m2 test p-value	N.A.

Table shows estimates of equation (5.2). Instruments are collapsed and with lags dated at $m-4$. The Arellano-Bond p-value (m2 test) is not available because the model has lagged values of order 2. $\hat{\psi}_m$ is the proxy of the aggregate monthly effects and is equal to the monthly cross-sectional arithmetic mean of the dependent variable. The survey-based forecast error on production is defined as $QS_{i,m+3} - QS^{e}_{im} = QS^{fe}_{im}$. *** indicates statistical significance at the 1%.

Appendix M

Robustness Checks on ARCH Estimation.

In this section I discuss some robustness checks concerning the General Method of Moments (GMM) estimates of the ARCH model. First, I should note that it is standard practice to include more lags as instruments in the GMM estimators of the Dynamic Panels. This is recommended as it improves efficiency. However, higher order lags can be weak instruments leading to biased estimates. In Table 1.M, in Columns (1) and (2) I show estimates of the variance persistence coefficient of equation (5.2) using the Arellano and Bover (1995) Two-Step GMM estimator with two different lag lengths as instruments; lags dated from $m-4$ to $m-5$ in column (1) and from $m-4$ to $m-9$ in column (2), collapsed in both cases. I observe that the Hansen p-value is lower than 5% (0.0172 and 0.0056 respectively) which is a clear indication of weak instruments. This justifies my choice of using only one lag as instrument without over-identifying the GMM system. Moreover, the estimated persistence coefficient in column (3) can serve as a benchmark for the negative Nickell (1981) bias in LSDV estimates of Dynamic Panel equations. I observe that my baseline estimate (0.0695 in Table 5.3) is quite higher that

that in column (3) of 1.M (0.0281), indicating that I have corrected the bias. What is more, the estimates of the persistence in columns (1) and (2) with the weak instruments are lower than my baseline one and closer the LSDV biased one: 0.0569 and 0.046 versus 0.0695 in my baseline estimate. This is probably the result of weak identification and it further justifies my choice of instruments in my baseline specification.

Table 1.M: Estimated ARCH Model for the Survey-based Forecast Error Variance of Production – Robustness Checks.

	(1)	(2)	(3)
Estimation	FOT		LSDV
Stand. Errors	2-step, Windmeijer corrected		Robust
Lags as Instruments	$m-4$ to $m-5$	$m-4$ to $m-9$	N.A.
Dependent Variable: Demeaned Squared Survey-based Forecast Error, $\left(QS^{fe}_{im} - \overline{QS^{fe}_{m}}\right)^2$			
$\left(QS^{fe}_{i,m-3} - \overline{QS^{fe}_{m-3}}\right)^2$	0.0569***	0.0460***	0.0281**
$\hat{\psi}_m$	0.987***	0.924***	1.012***
Constant	N.A.	N.A.	-0.0247
Observations	14,019	14,019	14,825
# of Firms	678	678	806
Over-identified	Yes	Yes	N.A.
Hansen p-value	0.0172	0.0056	N.A.
m2 test p-value	0.711	0.632	N.A.

Table shows estimates of equation (5.2). Columns (1) and (2) are with the Arellano and Bover (1995) Two-Step GMM estimator; column (3) with the ordinary LSDV estimator. Instruments are collapsed and with lags dated from $m-4$ to $m-5$ in (1) and from $m-4$ to $m-9$ in (2). The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate monthly effects with the monthly cross-sectional arithmetic mean of the dependent variable, I denote this as $\hat{\psi}_m$. The survey-based forecast error on production is defined as $QS_{i,m+3} - QS^{e}_{im} = QS^{fe}_{im}$. *** indicates statistical significance at the 1%.

In Table 2.M, I estimate the persistence of the squared forecast error without demeaning it. That is

$$(Q^{fe}_{i,m})^2 = FV_{im} + \psi_y + \psi_i + \eta_{im} = \phi(Q^{fe}_{i,m-3})^2 + \psi_m + \psi_i + \eta_{im}$$

The estimates are remarkably close to my baseline ones documenting robustness: 0.0708 in Table 2.M versus 0.0695 of the baseline estimate in Table 5.3.

Table 2.M: Estimated ARCH Model for the Variance of the Survey-based Forecast Errors on Production – Robustness to alternative specification.

Estimation	FOT
Stand. Errors	2-step, Windmeijer corrected
Lags as Instruments	$m-4$
Dependent Variable: Squared Survey-based Forecast Error, $\left(QS_{im}^{fe}\right)^2$	
$(QS_{i,m-3}^{fe})^2$	0.0708***
$\overline{(QS_m^{fe})^2}$	0.990***
Constant	N.A.
Observations	13,873
# of Firms	677
Over-identified	No
Hansen p-value	N.A.
m2 test p-value	0.993

Table shows estimates of equation (5.2). Instruments are collapsed and with lags dated at $m-4$. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. I proxy the aggregate monthly effects with the monthly cross-sectional arithmetic mean of the dependent variable, $(\overline{(QS_m^{fe})^2}$. The survey-based forecast error on production is defined as $QS_{i,m+3} - QS_{im}^{e} = QS_{im}^{fe}$. *** indicates statistical significance at the 1%.

Bibliography

Altig, D., Barrero, J. M., Bloom, N., Davis, S. J., Meyer, B. H., and Parker, N. (2019). Surveying business uncertainty. *National Bureau of Economic Research*.

Anderson, T. W. and Hsiao, C. (1981). Estimation of dynamic models with error components. *Journal of the American statistical Association*, 76(375):598–606.

Anderson Jr, O. (1952). The business test of the IFO-Institute for economic research, Munich, and its theoretical model. *Revue de l'Institut International de Statistique*, 20(1):1–17.

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297.

Arellano, M. and Bover, O. (1995). Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics*, 68(1):29–51.

Asimakopoulos, S. and Karavias, Y. (2016). The impact of government size on economic growth: A threshold analysis. *Economics Letters*, 139:65–68.

Bachmann, R., Born, B., Elstner, S., and Grimme, C. (2019). Time-varying business volatility and the price setting of firms. *Journal of Monetary Economics*, 101:82–99.

Bachmann, R. and Elstner, S. (2015). Firm optimism and pessimism. *European Economic Review*, 79:297–325.

Bachmann, R., Elstner, S., and Sims, E. R. (2013). Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, 5(2):217–249.

Bachmann, R. and Zorn, P. (2020). What drives aggregate investment? evidence from german survey data. *Journal of Economic Dynamics and Control*, page 103873.

Barrero, J. M. (2019). The micro and macro of managerial beliefs. *Stanford Institute for Economic Policy Research — Working Paper No. 19-010*.

Bartelsman, E. J., Caballero, R. J., and Lyons, R. K. (1994). Customer-and supplier-driven externalities. *The American Economic Review*, 84(4):1075–1084.

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

Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28(2):153–76.

Bloom, N., Bunn, P., Chen, S., Mizen, P., Smietanka, P., and Thwaites, G. (2019). The impact of brexit on uk firms. *NBER Working Paper 26218*.

Bloom, N., Davis, S. J., Foster, L., Lucking, B., Ohlmacher, S., and Saporta Eksten, I. (2017). Business-level expectations and uncertainty. *Mimeo*.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, 86(3):1031–1065.

Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016). Stereotypes. *The Quarterly Journal of Economics*, 131(4):1753–1794.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *The Journal of Finance*, 73(1):199–227.

Botsis, A., Görtz, C., and Sakellaris, P. (2020). Quantifying qualitative survey data: New insights on the (ir)rationality of firms' forecasts. *Mimeo, May 2020*.

Capistrán, C. and Timmermann, A. (2009). Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking*, 41(2-3):365–396.

Caselli, F. and Tesei, A. (2016). Resource windfalls, political regimes, and political stability. *Review of Economics and Statistics*, 98(3):573–590.

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

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *The American Economic Review*, 105(8):2644–2678.

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2015). How do firms form their expectations? new survey evidence. Technical report, National Bureau of Economic Research.

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). How do firms form their expectations? New survey evidence. *American Economic Review*, 108(9):2671–2713.

Cover, T. and Thomas, J. A. (2006). *Elements of Information Theory, 2nd Edition*. John Wiley and Sons.

Datta, D. and Du, W. (2012). Nonparametric HAC estimation for time series data with missing observations. *FRB International Finance Discussion Paper*, (1060).

Davidson, R. and MacKinnon, J. G. (2004). *Econometric theory and methods*, volume 5. Oxford University Press New York.

De Leeuw, F. and McKelvey, M. J. (1984). Price expectations of business firms: Bias in the short and long run. *The American Economic Review*, 74(1):99–110.

DGECFIN (2016). The joint harmonised eu programme of business and consumer surveys: User guide. *European Commission: Directorate-General for Economic and Financial Affairs*.

Dhami, S. (2016). *The foundations of behavioral economic analysis*. Oxford University Press.

ELSTAT (2002). Statistical classification of the branches of economic activity: Stakod 2003. Technical report, Hellenic Statistical Authority.

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

Enders, Z., Hünnekes, F., and Müller, G. J. (2019b). Monetary policy announcements and expectations: Evidence from German firms. *Journal of Monetary Economics*, 108:45–63.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, pages 987–1007.

EUROSTAT (2008a). Correspondence table nace rev. 1.1 - nace rev. 2. Technical report, EUROSTAT: The Statistical Office of the European Union.

EUROSTAT (2008b). Nace rev. 2: Statistical classification of economic activities in the european community. Technical report, EUROSTAT: The Statistical Office of the European Union.

Fan, J. and Gijbels, I. (1996). *Local Polynomial Modelling and Its Applications*. Chapman & Hall.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710.

Gennaioli, N., Ma, Y., and Shleifer, A. (2016). Expectations and investment. *NBER Macroeconomics Annual*, 30(1):379–431.

Gilchrist, S., Sim, J. W., and Zakrajšek, E. (2014). Uncertainty, financial frictions, and investment dynamics. Technical report, National Bureau of Economic Research.

Gorbachev, O. (2011). Did household consumption become more volatile? *American Economic Review*, 101(5):2248–70.

Greenwood, R. and Shleifer, A. (2014). Expectations of returns and expected returns. *The Review of Financial Studies*, 27(3):714–746.

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

Kosova, R. (2010). Do foreign firms crowd out domestic firms? Evidence from the Czech Republic. *The Review of Economics and Statistics*, 92(4):861–881.

Kozeniauskas, N., Orlik, A., and Veldkamp, L. (2018). What are uncertainty shocks? *Journal of Monetary Economics*, 100:1–15.

Lovell, M. C. (1986). Tests of the rational expectations hypothesis. *The American Economic Review*, 76(1):110–124.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of economic theory*, 4(2):103–124.

Lucas, R. E. (1973). Some international evidence on output-inflation tradeoffs. *The American Economic Review*, 63(3):326–334.

Lui, S., Mitchell, J., and Weale, M. (2011). Qualitative business surveys: signal or noise? *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 174(2):327–348.

Macaulay, A. (2019). Cyclical attention to saving. *University of Oxford Mimeo*.

Mackowiak, B., Matejka, F., and Wiederholt, M. (2018). Rational inattention: A disciplined behavioral model. Technical report.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER macroeconomics annual*, 18:209–248.

Massenot, B. and Pettinicchi, Y. (2018). Can firms see into the future? Survey evidence from Germany. *Journal of Economic Behavior & Organization*, 145:66–79.

Mitra, K. (2005). Is more data better? *Journal of Economic Behavior & Organization*, 56(2):263–272.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, pages 69–85.

Nerlove, M. (1983). Expectations, plans, and realizations in theory and practice. *Econometrica*, 51(5):1251–1279.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, pages 1417–1426.

OECD (2020). *OECD Composite Leading Indicators: Turning Points of Reference Series and Component Series*.

Pesaran, M. H. (1985). Formation of inflation expectations in british manufacturing industries. *The Economic Journal*, 95(380):948–975.

Pesaran, M. H. (1987). *The limits to rational expectations*. Blackwell Publishers.

Pesaran, M. H. (2015). *Time series and panel data econometrics*. Oxford University Press.

Pesaran, M. H. and Weale, M. (2006). Survey expectations. *Handbook of economic forecasting*, 1:715–776.

Polemis, M. L. and Stengos, T. (2019). Does competition prevent industrial pollution? evidence from a panel threshold model. *Business Strategy and the Environment*, 28(1):98–110.

Rao, M. M. and Swift, R. J. (2006). *Probability theory with applications*, volume 582. Springer Science & Business Media.

Roodman, D. (2009). A note on the theme of too many instruments. *Oxford Bulletin of Economics and Statistics*, 71(1):135–158.

Semykina, A. and Wooldridge, J. M. (2010). Estimating panel data models in the presence of endogeneity and selection. *Journal of Econometrics*, 157(2):375–380.

Seo, M. H. and Shin, Y. (2016). Dynamic panels with threshold effect and endogeneity. *Journal of Econometrics*, 195(2):169–186.

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

Smith, J. and McAleer, M. (1995). Alternative procedures for converting qualitative response data to quantitative expectations: an application to Australian manufacturing. *Journal of Applied Econometrics*, 10(2):165–185.

Souleles, N. S. (2004). Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the michigan consumer sentiment surveys. *Journal of Money, Credit and Banking*, pages 39–72.

Tanaka, M., Bloom, N., David, J. M., and Koga, M. (2020). Firm performance and macro forecast accuracy. *Journal of Monetary Economics*, 114:26–41.

Theil, H. (1952). On the time shape of economic microvariables and the Munich business test. *Revue de l'Institut International de Statistique*, 20(2):105–120.

Triguero, A. and Córcoles, D. (2013). Understanding innovation: An analysis of persistence for Spanish manufacturing firms. *Research Policy*, 42(2):340–352.

Van Beveren, I. (2012). Total factor productivity estimation: A practical review. *Journal of Economic Surveys*, 26(1):98–128.

Windmeijer, F. (2005). A finite smple correction for the variance of linear efficient two-step GMM estimators. *Journal of Econometrics*, 126(1):25–51.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data.* MIT press.