ECONOMIC STATISTICS CENTRE OF EXCELLENCE

A COLLABORATION WITH Office for National Statistics

# Quantifying Qualitative Survey Data: New Insights on the (Ir)Rationality of Firms' Forecasts

**Alexandros Botsis, Christoph Görtz and Plutarchos Sakellaris**

ESCoE Discussion Paper 2021-14

October 2021

ISSN 2515-4664

DISCUSSION PAPER

Quantifying Qualitative Survey Data: New Insights on the (Ir)Rationality of Firms' Forecasts
Alexandros Botsis, Christoph Görtz and Plutarchos Sakellaris
ESCoE Discussion Paper No. 2021-14
October 2021

**Abstract**

Using a novel dataset that combines firms' qualitative survey-based sales forecasts with their quantitative balance-sheet data on realized sales, we document that only major forecast errors (those in the two tails of the distribution) are predictable and display autocorrelation. This result is a particular violation of the Full Information Rational Expectations hypothesis that requires explanation. In contrast, minor forecast errors are neither predictable nor autocorrelated. To arrive at this finding, we develop a novel methodology to quantify qualitative survey data on forecasts. It is generally applicable when quantitative information, e.g. from firm balance sheets, is available on the realization of the forecasted variable. To explain our empirical result, we provide a model of rational inattention. When operating in market environments where information processing is more costly, firms optimally limit their degree of attention to information. This results in larger absolute forecast errors that become predictable and autocorrelated.

Keywords: Expectations, Firm Data, Forecast Errors, Panel Threshold Models, Rational Inattention, Survey Data.

*JEL classification:* C53, C83, D22, D84, E32

Christoph Görtz, University of Birmingham, c.g.gortz@bham.ac.uk.

Published by:
Economic Statistics Centre of Excellence
National Institute of Economic and Social Research
2 Dean Trench St
London SW1P 3HE
United Kingdom
www.escoe.ac.uk

# Quantifying Qualitative Survey Data: New Insights on the (Ir)Rationality of Firms' Forecasts*

Alexandros Botsis
University of Nottingham

Christoph Görtz
University of Birmingham

Plutarchos Sakellaris
Athens University of Economics and Business

**Abstract**

Using a novel dataset that combines firms' qualitative survey-based sales forecasts with their quantitative balance-sheet data on realized sales, we document that only major forecast errors (those in the two tails of the distribution) are predictable and display autocorrelation. This result is a particular violation of the Full Information Rational Expectations hypothesis that requires explanation. In contrast, minor forecast errors are neither predictable nor autocorrelated. To arrive at this finding, we develop a novel methodology to quantify qualitative survey data on forecasts. It is generally applicable when quantitative information, e.g. from firm balance sheets, is available on the realization of the forecasted variable. To explain our empirical result, we provide a model of rational inattention. When operating in market environments where information processing is more costly, firms optimally limit their degree of attention to information. This results in larger absolute forecast errors that become predictable and autocorrelated.

*Keywords:* Expectations, Firm Data, Forecast Errors, Panel Threshold Models, Rational Inattention, Survey Data.

*JEL Classification:* C53, C83, D22, D84, E32.

---

* We thank Fabio Antoniou, Danilo Cascaldi-Garcia, Wanyu Chung, Yiannis Dendramis, Martin Ellison, Christos Genakos, Bruce Hansen, Menelaos Karanasos, Yiannis Karavias, Tryphon Kollintzas, Alistair Macaulay, Bartosz Mackowiak, Kaushik Mitra, Thanasis Stengos, Mirko Wiederholt and Klaus Wohlrabe for extremely useful comments and suggestions. We particularly thank Foteini Thomaidou from IOBE for detailed insights about the survey data. All remaining errors are our own. This work was produced using statistical data from ONS, and we are also thankful for their support in data provision. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates. Botsis: University of Nottingham, Department of Economics. Email: alex.botsis@nottingham.ac.uk. Görtz (corresponding author): University of Birmingham, Department of Economics, University House, B15 2TT, Birmingham, UK. Email: c.g.gortz@bham.ac.uk. Sakellaris: Athens University of Economics and Business, Department of Economics, Email: plutarch@aueb.gr.

# 1 Introduction

Do firms make errors in forecasting their future sales that are predictable and display autocorrelation? If so, what does that reveal about firm behavior and the way they form expectations? We address these questions using a novel panel dataset that contains qualitative survey data on manufacturing firms' own sales forecasts as well as corresponding balance-sheet data on realized sales. We document that only when firms make major forecast errors are these predictable and display autocorrelation. Hence, firms' behavior violates the Full Information Rational Expectations (FIRE) hypothesis. In contrast, minor forecast errors do not violate FIRE. These differences in firms' forecasting behavior have not been documented before. Major forecast errors are those that lie in the two tails of the distribution. Key to demonstrating this empirical result is our novel methodology that quantifies forecasts when survey-based data on expectations is only available in categorical form. To explain our empirical findings, we provide a model of rational inattention. When operating in market environments where information processing is more costly, firms optimally limit their degree of attention to new information. Limited attention results in major forecast errors that become autocorrelated and predictable.

Given the dynamic nature of firm activity, expectations play a major role in firms' economic behavior. Economic models that describe firm behavior are naturally dynamic and contain assumptions about expectations. Many papers have emphasized the importance of obtaining evidence on expectations formation that is independent of model assumptions (see Nerlove (1983) and Manski (2004) among others). This makes the use of survey data on expectations particularly useful. However, as survey-based measures for expectations are typically categorical, some important questions cannot be answered. For instance, whether firms make substantial errors in their forecasts and what are their statistical properties. Our paper provides a remedy to this obstacle, as we develop a novel methodology that converts categorical survey data on expectations to continuous quantities.

We develop a unique dataset by matching confidential information on firms' monthly qualitative forecasts on own sales together with annual quantitative balance-sheet information on sales. The dataset covers Greek firms in manufacturing for the period of 1998 to 2015. After generating quantified measures of expectations, we test and find that forecast errors are both predictable and autocorrelated. This is in accordance with existing results in the literature that concern either firm-level or aggregate variables (see Gennaioli et al. (2016), Massenot and Pettinicchi (2018), Bordalo et al. (2018a), Bordalo et al. (2018b), Tanaka et al. (2019)). These are clear violations of the FIRE hypothesis. Where we differ from previous papers, however, is that we show that, in our dataset, this rejection is entirely due to forecast errors in the upper and lower 26% of the distribution. Only these major forecast errors are predictable and autocorrelated. Forecast errors in the middle 48% of the distribution are still often economically significant, but are neither predictable nor autocorrelated. We derive this novel finding using a modified version of the Dynamic Panel Threshold estimator of Seo and Shin (2016). This estimator endogenously determines the 26% threshold, that distinguishes minor from major forecast errors, to fit the data best. Major forecast errors may lead firms to make suboptimal decisions, pointing to the possibility that policy design can be geared to helping firms avoid these.[1] Clearly, one cannot carry out such analysis with qualitative data and this points to the importance of our quantification exercise.

Our quantification methodology builds on the work of Pesaran (1987) and Smith and McAleer (1995) and extends it to retain the panel nature of the dataset. We use higher-frequency (monthly) qualitative survey data on expected sales growth together with lower-frequency (annual) quantitative data on realized sales growth to estimate quantified expected annual sales growth. In order to retain the panel nature of the dataset we need to overcome challenges such as unobserved firm heterogeneity and omitted variable problems for forecast errors. This requires identifying assumptions that allow us to derive two nonlinear equations. The first one relates observed quantitative annual sales growth to observable variables and the second one relates unobserved quantitative expected sales growth to observable variables. The key is that both of these relationships depend on the same parameters. Then, we estimate the common parameters from the first equation using Nonlinear Least Squares (NLS), and use these estimated parameters in the second equation to derive fitted values for quantitative expectations on sales. This methodology can be applied to a wide range of

[1]We document that the vast majority of firms, independent of their size, make minor as well as major forecast errors. Major and minor forecast errors (or short series thereof) tend to alternate.

applications and datasets and is not limited to quantifying firm forecasts. The only requirement is that the researcher has access to high-frequency categorical survey data on expectations together with (potentially lower-frequency) quantitative realizations of the corresponding variables.

We provide evidence of external validity and accuracy for our methodology in four ways. First, we show that our quantified estimates on sales growth expectations are fully consistent in terms of sign with the corresponding qualitative survey-based expectations. In a horse race, our methodology also substantially outperforms ordered response models which are potential alternatives for obtaining quantified predictions. Second, we construct a small dataset of UK manufacturing firms that contains monthly qualitative survey expectations and the corresponding annual realizations from balance sheets which allow us to use our methodology to derive estimates for annual forecast errors. Importantly, for each firm, the dataset also includes annual quantitative survey expectations which we employ as a benchmark. Comparing our estimated annual forecast errors with the directly observable benchmark forecast errors confirms the accuracy of our quantification methodology. Such an exercise can only be conducted using a dataset that includes quantitative forecasts, made by the same forecaster at a high and a lower frequency. In practice, this is challenging to do due to the rare availability of such data on quantitative firm-level expectations. In fact, this dearth of data highlights the need for and value of our quantification methodology, which allows researchers to utilize the large number of qualitative surveys to quantify expectations. Third, we perform a Monte Carlo exercise that provides a benchmark based on simulated data. We find forecast errors based on our methodology are highly accurate when compared with forecast errors based on the underlying artificial 'true' data. Although the three exercises above confirm that our methodology delivers rather accurate forecast errors, it is noteworthy that, in a fourth exercise, we additionally show that our empirical results on the autocorrelation and persistence properties of large forecast errors are not driven by the quantification methodology but are a feature in the underlying qualitative survey data.

We use a framework of rational inattention to interpret our empirical results on the predictability and autocorrelation of forecast errors. Our work links to a large literature on rational inattention to which we cannot do full justice here — the seminal work is Sims (2003) and Mackowiak et al. (2018)

provide a recent survey of the literature. In particular, we study the properties of forecast errors when firms endogenously choose the degree of attention to costly information. We show that when the cost of processing information is high, firms make major forecast errors which are predictable and negatively autocorrelated. In the absence of these costs, firms make fully informed and rational forecasts with minor forecast errors. Importantly, the framework can also rationalize our empirical negative coefficient estimates on persistence and autocorrelation for major forecast errors. These coefficients are negative in the model only if the autocorrelation of sales growth is negative, which is indeed also a feature in our data.

Quantifying forecast errors using qualitative survey data is a very important matter for many questions, but there has been little work on this and no generally accepted methodology. Theil (1952) and Anderson Jr (1952) developed the so called 'probability method'. It provides the theoretical grounds for the 'balance statistics' that are widely used for the published business and consumer sentiment indexes. Pesaran (1987) provides a useful analysis of the limitations of this approach (see also Pesaran and Weale (2006)). A very useful first step to overcoming such limitations is Bachmann and Elstner (2015). They first restrict their survey sample to firms that reported expected output to be unchanged over the following three months. Then, they classify non-zero percentage change of firm's reported utilization as a forecast error. This technique has some limitations compared to our quantification method. Our method does not only deliver continuous forecast errors but also expectations themselves. It further is not limited to the quantification of firms' production, but can be applied to any variable in principle, given the data requirements outlined above. Importantly, our method can be used on the entire sample rather than only on a potentially small subset of firms.

Important early work on the use and pitfalls of survey data to analyze how firms form expectations includes de Leeuw and McKelvey (1981) and Nerlove (1983). Our work is part of a now fast growing literature that uses information from surveys to understand firms' decision making. Enders et al. (2019a) use German data from the IFO Business Survey to study how firms' expectations about future production affect their current decisions on production and price setting. Tanaka et al. (2019) use novel Japanese data to study how firm characteristics affect their GDP forecasts. To the best of our knowledge, these two datasets are the only ones constructed so far to contain categorical firm survey data with corresponding quantitative data, e.g. from balance sheets or national accounts. We contribute to the survey literature by providing a novel dataset that combines responses to a rich firm-level survey with the corresponding balance sheet information for Greece. Our empirical results point to the importance of further work on merging existing quantitative datasets with qualitative survey data.[2] Applying our quantification methodology would then allow for a deeper understanding of how firms or households form expectations and their economic impact. There are many other contributions in the literature that use survey data to help our understanding of firm-level and aggregate variables. Enders et al. (2019b) for example use German data from the IFO Business Survey to study how monetary policy announcements affect firms' expectations. Bachmann and Zorn (2013) use this survey to understand the drivers of aggregate investment. Bloom et al. (2019) use survey responses to understand the causes and consequences of Brexit for the UK economy. Coibion et al. (2018) study how firms form expectations about macroeconomic conditions using novel survey evidence from New Zealand.

Our findings contribute to the broader literature on testing whether agents form expectations rationally. In addition to papers mentioned above, there are several key contributions in this literature. For example, Coibion and Gorodnichenko (2015) use data from professional forecasters to test the FIRE hypothesis. They find that agent's expectation formation violates the FIRE hypothesis and show, in line with the spirit of our model, that this is consistent with the presence of information rigidities. Coibion and Gorodnichenko (2012) use survey data from firms, households and professional forecasters and show that expectation formation is better aligned with models of noisy information, similar to our model, rather than with frameworks in which information is sticky.

The rest of the paper is organized as follows. Section 2 discusses the data. Section 3 lays out our methodology to quantify firms' forecasts and describes the characteristics of the estimated forecasts and the resulting forecast errors. This section also provides evidence of external validity and accuracy for our methodology. Section 4 discusses our empirical results on the predictability and

[2] A novel dataset that combines households' survey based inflation expectations with administrative data has recently been developed in Vellekoop and Wiederholt (2019).

autocorrelation of forecast errors. These empirical findings are interpreted in Section 5 in a model with rational inattention. Section 6 provides concluding remarks.

## 2 Data

Our dataset is constructed by merging two databases that cover Greek firm-level data. The first database includes annual information on firms' balance sheets and income statements. We obtain this data from ICAP S.A., a private consultancy firm, which collects and digitalizes this information from official publicly available records. The financial statements are compiled by certified auditors (chartered accountants) and are used, among other things, for reporting to tax authorities and investors, by commercial banks for credit decisions, and by the central bank for credit rating information. They are available to us at an annual frequency from 1998 to 2015 which determines our sample horizon. As such, our dataset includes two distinct episodes of the Greek economy, a long boom up to 2008 and the subsequent severe recession. We use firms' sales from the financial statements, which is deflated using the implicit gross value added deflator from Eurostat.[3]

The second database comprises firms' responses to a monthly survey conducted by the Foundation for Economic and Industrial Research (IOBE). This survey is used by IOBE to construct the much-followed business climate index for the Greek economy since 1985 and is part of the European Commission's business climate index for the European Union.[4] All survey questions concern current, past or expected future firm-level developments. The survey does not include any questions about aggregate macroeconomic or sector-level conditions. Since participation is confidential and voluntary, firms have no strategic interest in misreporting. Further details about the survey are provided in Appendix A.1.

The IOBE classifies firms in four broad sectors — manufacturing, construction, retail trade, and

[3]Nominal and real (2005 base year) value added for Greece is available from Table nama_10_a64.

[4]The survey is commissioned by the European Commission and conducted for the Greek economy in compliance with the guidelines of the European Commission's Directorate General for Economic and Financial Affairs (see DGECFIN (2017)). A corresponding survey is conducted for the European Commission for example for the United Kingdom by the Confederation of British Industry and for Germany by the IFO Institute.

services — and sends out surveys that include somewhat different questions across these sectors. We focus on the manufacturing sector as this sector's survey includes questions about anticipated and current sales developments. Responses to these two questions, and the fact that the question on current sales has a direct counterpart in the financial statements data, are key for the quantification of forecast errors.[5] The relevant (translated) questions in the survey are

> 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, we 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. For the remainder of the paper, variables in capital letters denote qualitative variables and lower case letters stand for quantitative variables.

Under a strict confidentiality agreement, we were given access to the un-anonymized survey data. Using the firm's unique tax identifier, we merged their survey responses with the respective balance sheet data. Details about the cleaning procedures for the two parts of our dataset are outlined in Appendix A.2 and A.3. Our cleaned and merged 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. Table 1 provides an overview of the firms in our sample. Our sample includes very small firms but also 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 Appendix A.4 we provide evidence that our sample is representative for the manufacturing sector

[5]The manufacturing sector is also the largest of these broad sectors as it includes 38% of survey observations and 36% of observations in the financial statements data.

and establish in several exercises the high quality of the survey responses. Appendix B.6 shows the distribution of survey expectations on sales growth and documents their evolution over time.

Table 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 |

# 3 Quantitative Forecast Errors

The forecast error on sales growth is defined as the difference between actual sales growth and its forecast for the corresponding period. Evaluating the size of firms' forecast errors hence requires quantitative data on sales growth forecasts and their subsequent realization. While the financial statements data provide an annual quantitative measure for the latter, quantitative data on firm's sales growth forecasts is not readily available.

Section 3.1 proposes a novel methodology that uses qualitative survey data on the direction of firm sales growth forecasts and quantitative data on realized sales growth from financial statements, to derive a quantitative estimate for firms' sales growth forecasts. In particular, to quantify the survey responses we extend the methodology by Pesaran (1987) and Smith and McAleer (1995) who aggregate qualitative firm observations cross-sectionally to derive quantitative time series. We extend their work and show how the panel dimension of our dataset can be retained by using the monthly survey observations to derive annual quantitative sales growth forecasts. Retaining the panel dimension comes with new challenges, such as dealing with unobserved heterogeneity and an omitted variable problem, and we show how to address these in our quantification framework.

Section 3.2 applies this methodology to our dataset and provides a first look at the characteristics and distributions of the estimated quantified forecast errors. In Section 3.3, we provide evidence of external validity and accuracy for our methodology. First, we show that our quantified estimates are quite accurate in terms of the sign of expected sales growth when compared to the qualitative survey-based data. In this aspect, our methodology also outperforms ordered response models in a horse race. Second, we provide evidence for the accuracy of our quantification methodology in terms of the magnitude of firm growth forecasts. This is challenging since we do not observe true quantitative firm growth forecasts in our data. Our solution is to conduct a Monte Carlo study using artificial datasets.

### 3.1 Quantifying Expected Sales Growth

Consider the expected annual growth rate of sales for firm $i$ in year $y$, $x^e_{iy} \triangleq \mathbb{E}\big[x_{iy}|\mathcal{F}_{i,y-1}\big]$, that is based on an information set $\mathcal{F}$ at the end of year $y-1$. Theoretically, we can decompose firm's expected annual sales growth into its monthly components. For this purpose, we define firm $i$'s expectation about average sales growth in the next three months as $x^e_{im} \triangleq \mathbb{E}\big[x_{i,\{m,m+1,m+2\}}|\mathcal{F}_{i,m-1\in y}\big]$, where $x_{i,\{m,m+1,m+2\}}$ is the average growth rate of sales for the following three-month period.[6] Note that this expectation is formed based on an information set at the end of month $m-1$. This quantitative monthly sales forecast is consistent, in terms of the structure and horizon of the information set, with the qualitative survey forecast $XS^e_{im}$. To re-iterate, the annual forecast $x^e_{iy}$ is based on the information available to the firm at the time of the forecast, this is at the end of year $y-1$. The monthly forecast is also based on the information at the time of the forecast, which is the end of month $m-1$. The monthly expected growth rates can in turn be separately expressed, for firm $i$ in month $m$, using their positive, $x^{e,+}_{im}$, and negative, $x^{e,-}_{im}$, components.[7] The aggregation of monthly growth rates to an annual frequency can be formalized as the following weighted average

$$x^e_{iy} = \mathbb{E}_{i,y-1} \sum_{m\in y} W^+_{im} x^{e,+}_{im} + \mathbb{E}_{i,y-1} \sum_{m\in y} W^-_{im} x^{e,-}_{im}, \tag{1}$$

[6] To be able to decompose the annual expected sales growth into monthly components, we define $x^e_{im}$ for November (December) to include expectations about the next two (one) months only. This is consistent with our 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 A.2. This scheme avoids double counting of months.

[7] Zero growth rates have no effect on the decomposition. We separate positive and negative components to allow below for possible differences in relationships with their annual counterparts.

where the weights are defined as $W^{+}_{im} = W_{im}\mathbb{1}_{[XS^{e}_{i,m}=+1]}$ and $W^{-}_{im} = W_{im}\mathbb{1}_{[XS^{e}_{i,m}=-1]}$ 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 outcome. It is defined as

$$W_{im} \triangleq \frac{w_{im}}{\sum_{m\in y} w_{im}}, \tag{2}$$

where $w_{im}$ is the ratio of the seasonally unadjusted over the seasonally adjusted real gross value added. Intuitively, when this ratio exceeds unity, 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 our practical implementation below, data availability limits the design of $w_{im}$ to be the same across all firms in the manufacturing sector at quarterly frequency.[8] The second component of the weights is dummy variables that take a value of unity if the expected sales growth rate is either positive, $\mathbb{1}_{[XS^{e}_{i,m}=+1]}$, or negative, $\mathbb{1}_{[XS^{e}_{i,m}=-1]}$. While we do not observe quantitative expectations of sales growth in equation (1) — $x^{e}_{iy}$, $x^{e,+}_{im}$ and $x^{e,-}_{im}$ — our dataset includes survey responses on the qualitative expected change in sales, $XS^{e}_{i,m}$. The aim of this section is to derive a quantitative estimate for annual expected sales growth, $x^{e}_{iy}$.

As a first step towards this, we follow Pesaran (1987) and assume that for each firm the monthly expected sales growth rates are linearly positively correlated with the corresponding annual expected sales growth.[9] We also allow for this linear correlation to be asymmetric, as in Smith and McAleer (1995), depending on whether the monthly expectation variable is positive or negative. This is the first identifying assumption (ID1) we make to quantify firms' forecast errors. It can be formalized

[8]We 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. We use value added since information on sales is not available at monthly or quarterly frequency. Appendix B.3 shows that quantitative sales forecasts are almost identical if an alternative weighting is applied that assigns equal weights across all observations per year.

[9]While this assumption is intuitive, we also show in Appendix B.4 that ID1 is sensible in the sense that the monthly qualitative survey forecasts are correlated with their annual quantitative counterparts.

as

$$x_{im}^{e,+} = \alpha + \gamma_1 x_{iy}^e + \nu_{im}^+, \qquad \text{and} \qquad x_{im}^{e,-} = -\beta + \gamma_2 x_{iy}^e + \nu_{im}^-, \qquad \text{[ID1]} \tag{3}$$

where the error terms, $\nu_{im}^+$ and $\nu_{im}^-$, are assumed to be independently distributed across firms.[10] We further allow for the coefficients $\alpha$, $\beta$, $\gamma_1$ and $\gamma_2$ to differ across boom and bust periods (1998-2008 and 2009-2015 in our sample). We will specify this at the end of this section, but abstain from accounting for this state dependence in the notation for now to ease the exposition.

Equations (3) are not formulated to conduct any inference, but merely to reflect that for each firm the annual expected growth rate should be correlated with the corresponding monthly components. In fact, this linear relationship in equations (3) can be used to eliminate the unobserved variables $x_{im}^{e,+}$ and $x_{im}^{e,-}$ from equation (1). Combining equations (1) and (3) yields (detailed derivations are shown in Appendix B.1)

$$x_{iy}^e = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \qquad \text{with} \qquad \xi_{iy} = \frac{\mathbb{E}_{i,y-1} \sum_{m \in y} \left( W_{im}^+ \nu_{im}^+ + W_{im}^- \nu_{im}^- \right)}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}, \tag{4}$$

where we define

$$P_{iy} \triangleq \sum_{m \in y} W_{im} \mathbb{1}_{[XS_{im}^e = 1]}, \qquad \text{and} \qquad N_{iy} \triangleq \sum_{m \in y} W_{im} \mathbb{1}_{[XS_{im}^e = -1]}, \tag{5}$$

to ease the notation. $P_{iy}$ ($N_{iy}$) denotes the weighted share of months per year that record a rise (decline) in expected sales of firm $i$. These qualitative variables are directly available from the survey data so that we can observe $P_{iy}$ and $N_{iy}$. However, we cannot estimate equation (4) since we do not observe quantitative expectations of annual sales growth, $x_{iy}^e$, in the data. In fact, deriving quantitative sales growth expectations was our goal in the first place. Instead, if we had estimates for the parameters and knowledge of the error term — and given that we observe $P_{iy}$ and $N_{iy}$ — we could use equation (4) to derive fitted values for $x_{iy}^e$. Indeed, the next steps of the derivation are undertaken to facilitate exactly this.

We know that for each firm $i$, realized sales growth in year $y$ is the sum of expected sales growth

[10] Potential monthly serial autocorrelation in these error terms is not of concern since we will aggregate them at a firm-year frequency.

in that year and a forecast error, $x_{iy} = x^e_{iy} + x^{fe}_{iy}$. Using this expression to replace $x^e_{iy}$ in equation (4) yields after rearranging

$$x_{iy} \quad = \quad \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + x^{fe}_{iy} + \xi_{iy}. \tag{6}$$

In principle, this equation can be estimated, as the financial statements data includes quantitative information about realized annual sales growth, $x_{iy}$. While the forecast error, $x^{fe}_{iy}$, is still unobserved, estimating equation (6) without this variable is simply an omitted variable problem that adds to the error term. For the remainder of this subsection, we discuss this omitted variable problem and deal with unobserved firm heterogeneity to obtain an expression of equation (6) that can be estimated.

**Omitted Variable Problem.** To ease the notational burden in this section, we use equation (4) to define the conditional expectation of the quantitative sales growth expectations as

$$\tilde{x}^e_{iy} \triangleq \mathbb{E}\big[x^e_{iy}\big|P_{iy}, N_{iy}\big] = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{7}$$

To obtain consistent estimates of the parameters in equation (6), we need the composite error term $x^{fe}_{iy} + \xi_{iy}$ to be mean independent of the non-linear function $\tilde{x}^e_{iy}$ (see Davidson and MacKinnon (2004)). We proceed now to show this. Note that the forecast error, $x^{fe}_{iy}$, is mean independent from the forecast $x^e_{iy}$.[11] Since $\mathbb{E}\big[x^{fe}_{iy}\big|x^e_{iy}\big] = 0$ holds, it also implies that $\mathbb{E}\big[x^{fe}_{iy}\big|\tilde{x}^e_{iy}\big] = 0$. We provide a proof of this statement in Appendix B.2. Intuitively, firms' expected sales growth, $x^e_{iy}$, cannot 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 established the forecast error's mean independence of $\tilde{x}^e_{iy}$, and in order to obtain consistent estimates of the parameters in equation (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

[11] 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.

$\mathbb{E}[\xi_{iy}|\{XS^e_{im}\}_{m\in y}] = 0$. In Appendix B.2 we provide a formal proof of this statement. This leaves us with the task to control for the unobserved firm heterogeneity that is likely to make $\xi_{iy}$ correlated with $\{XS^e_{im}\}_{m\in y}$. We turn to this next.

**Unobserved Firm Heterogeneity.** From equation (4), we know that the numerator of $\xi_{iy}$ depends on the error terms $\nu^+_{im}$, $\nu^-_{im}$, and on $XS^e_{im}$. We need to account for the effect of the unobserved heterogeneity hidden in this numerator. For that, our second identifying assumption (ID2) is to assume that the error term in equation (4) can be decomposed as

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

where $\psi_i$ captures the effect of unobserved firm heterogeneity on sales growth. $\vartheta_{iy}$ is an idiosyncratic error which is mean-independent of $XS^e_{im}$ for all $m$.[12] Note that the notation $\{XS^e_{im}\}_{m=1,\dots,T_i}$ denotes the entire history of months $m$ for variable $XS^e_{im}$, where $T_i$ is firm $i$'s total number of monthly observations.[13] The unobserved firm heterogeneity, $\psi_i$, 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 $\mathbb{E}[\psi_i + \vartheta_{iy}|\{XS^e_{im}\}_{m=1,\dots,T_i}] = \mathbb{E}[\psi_i|\{XS^e_{im}\}_{m=1,\dots,T_i}] \neq 0$ from assumption ID2. To control for unobserved heterogeneity, we need to approximate $\mathbb{E}[\psi_i|\{XS^e_{im}\}_{m=1,\dots,T_i}]$.[14]

A widely used approximation for this purpose is the one suggested in Mundlak (1978).[15] The

[12]ID2 is a standard assumption to deal with unobserved heterogeneity. Essentially, ID2 says that the composite error term in the numerator of equation (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.

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

[14]The structure of the non-linear equation (4) that we want to estimate does not allow us to derive an estimator for $\psi_i$ analytically, and we cannot use dummy variables either, because the cross-sectional dimension is very large and the effect of $\psi_i$ is time varying. Another possibility would be to linearize (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, we would need to use local polynomial fitting methods which are too complex, both algebraically and computationally, in our context with even two explanatory variables.

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

original Mundlak (1978) specification is linear, but in the following we additionally include a second-order term due to the non-linearity of equation (4). Therefore, our third identifying assumption is that the conditional expectation of the unobserved firm heterogeneity in the error term $\xi_{iy}$ is

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

where $\delta_1$ and $\delta_2$ are coefficients. This results in the following auxiliary regression for $\psi_i$

$$\psi_i \quad = \quad \delta_1\overline{XS^e_i} + \delta_2(\overline{XS^e_i})^2 + \omega_i, \tag{8}$$

where $\omega_i$ 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 $i$. We can now substitute equation (8) for $\psi_i$ in the numerator of $\xi_{iy}$, obtaining

$$\xi_{iy} = \frac{\delta_1\overline{XS^e_i} + \delta_2(\overline{XS^e_i})^2 + \omega_i + \vartheta_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{9}$$

**The Final Equation to be Estimated.** As we have provided a way to approximate the unobserved firm heterogeneity, we can now derive the final equation to be estimated. We substitute equation (9) into equation (6) and obtain

$$x_{iy} \quad = \quad \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1\overline{XS^e_i} + \delta_2(\overline{XS^e_i})^2}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \tilde{\xi}_{iy}, \tag{10}$$

where

$$\tilde{\xi}_{iy} = x^{fe}_{iy} + \frac{\omega_i + \vartheta_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{11}$$

Overall equation (10) is estimable because the error term $\tilde{\xi}_{iy}$ is mean-independent of the explanatory variables. We provide a formal proof of this statement in Appendix B.2. This addresses the issue of the unobserved heterogeneity in equation (6), so that we can obtain consistent estimates of the coefficients of interest, $\alpha$, $\beta$, $\gamma_1$ and $\gamma_2$.[16]

models, it is equivalent to the least squares dummy variable and the standard within estimator.

[16] The error term, $\tilde{\xi}_{iy}$, in equation (10) is likely to be heteroscedastic and autocorrelated within each firm. When we estimate such an equation, we will use the heteroscedasticity robust estimator for the standard errors which addresses

**Summary of the Quantification Method.** We have derived two nonlinear equations. Equation (10), relates observed quantitative annual sales to observable variables and the identifying assumptions ID2 and ID3 ensure that the coefficient estimates are consistent. Equation (4) relates unobserved quantitative expected sales growth to observable variables. The key is that both of these relationships depend on the same parameters. We estimate the parameters from equation (10) using Nonlinear Least Squares (NLS), and use these estimated parameters in equation (4) to derive fitted values for quantitative expectations on sales growth.

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

4. Use the NLS estimated coefficients of equation (10) to compute the fitted values, $\hat{x}_{iy}^e$, for quantified sales growth forecasts from equation (4).

Our parameter estimates of the NLS estimation of equation (10) are documented in Appendix B.3. 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 the following sections, we will drop the hat from the expression for the forecast error to ease notation. Our methodology to quantify forecasts and forecast errors is generally applicable to variables other than sales growth. It is applicable to any qualitative (survey based)

both problems — this robust estimator treats errors as clustered within cross-sectional units.

variable on future developments, as long as a quantitative corresponding variable on realization is available.

### 3.2 Descriptive Statistics on the Quantified Forecast Errors

The previous section outlined the methodology to derive annual quantitative forecast errors for sales growth using monthly qualitative survey data and annual quantitative data from financial statements. This section provides an overview of the characteristics of the estimated forecast errors.

Figure 1 shows the distribution of forecast errors. We report moments on this distribution in Table 2. The average forecast error in our sample is zero and slightly larger than the median (-0.03). This implies that the median forecast on sales growth is three percentage points more optimistic than the subsequent realization. Overall, a number of forecast errors made by firms are small (in absolute value), as these are centred close to zero, but still a significant number of forecast errors made are quite substantial. Since the remainder of the paper will be concerned with such major forecast errors, we now also provide some statistics about these. For this purpose, we classify the top and bottom 26% of forecast errors to be major, which is in line with the estimates for this threshold obtained in Section 4. At the bottom (top) 26 percentiles, firms expected sales growth to be 14.3 (8.6) percentage points higher (lower) than subsequently realized. Hence, also a large number of the remaining 48% of forecast errors in the center of the distribution, which we call minor, are still economically significant. Interestingly, Table 2 shows that the distribution of forecast errors is very stable across the boom and the severe depression periods in our sample — during both periods their shares are close to the 26% of the full sample which is imposed by construction.[17]

How are these major and minor forecast errors distributed over different dimensions of our sample? Panel A of Table 3 sorts the sample according to the share of major forecast errors in a firm's total

[17]Appendix B.6 documents that the share of major positive and negative forecast errors can vary somewhat in particular years — e.g. in 2009, the first year of the Greek depression, the share of negative forecast errors increased. Overall, the shares are rather stable though, also when averaging over fewer years than included in the boom and bust subperiods. Furthermore, the shares of major positive and negative forecast errors are also very similar across 2-digit sectors — results are available upon request.

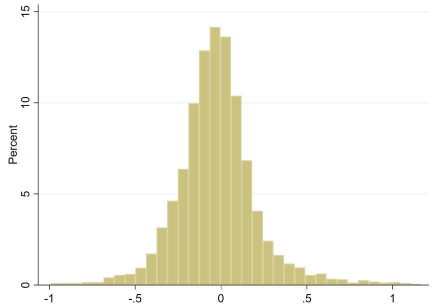


Figure 1: **Distribution of Annual Quantified Sales Growth Forecast Errors.** The 1% of forecast errors at the top of the distribution are omitted to ease visibility.

Table 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 |

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).

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 our sample for a relatively long time as they also account for the vast majority of the firm year observations (3,119 out of total 3,868). 179 (111) firms make exclusively major (minor) forecast errors, however these account only for 320 (194) firm-year observations in our sample and are hence quite short lived. Panel B of Table 3 shows that the share of major forecast errors in the total observations of a firm is relatively constant across different firm sizes. It varies between 48% and 55% across the firm size distribution where the larger firms make slightly fewer major forecast errors.

Table 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.

Table 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.

The evidence in Table 3 shows that, independent of their size, most firms make major as well as minor forecast errors. Table 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 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.

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 observations. Forecast errors are not highly persistent and both major and minor forecast errors tend to alternate. Additionally, this section documented that the share of major positive and negative forecast errors is stable across the boom and bust periods in our sample. Given that forecast errors do not exhibit distinctly different distributions across the boom and bust period, we focus our empirical analysis in Section 4 on the full sample.

## 3.3 External Validity and Accuracy of the Methodology

In this section, we conduct two types of exercises to demonstrate the external validity of our quantification methodology. In the first type, we use the qualitative firm forecast data from the survey as a benchmark and test whether our quantified estimates are accurate in terms of the sign of expected sales growth. We also perform a horse race with ordered response models. In the second type, we test the accuracy of our quantification methodology in terms of the magnitude of firm growth forecasts. We do so by conducting a Monte Carlo experiment using artificial datasets, and also by employing our methodology on a dataset of UK firms for which qualitative monthly and quantitative annual survey forecasts are directly available.

#### 3.3.1 Directional Consistency of Estimated Forecasts with the Survey Data.

We can use the observed survey data on the direction of expected sales growth to benchmark how well our quantified forecasts match the direction of expected sales growth. To facilitate the comparison of the monthly survey data with our annual forecast estimates, we 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 (2). The distributions of the raw monthly and annualized survey expectations are reported in Appendix B.6. 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 our estimates of quantified forecasts against the annualized survey-based qualitative forecasts, we split responses in each of these two variables into three categories — positive, zero or negative — and cross-tabulate the three directions. Panel A.1 in Table 5 reports how well our quantified forecasts match the direction of the annualized observable. 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 our 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 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, we also report in Table 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.[18]

Table 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) we used the direction of the sales growth predicted by the model at the firm-month level instead of the latent variable and then we 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.

These results provide a first indication of the quality of our forecast estimates. Next, we run a horse race with alternative ways to quantify sales growth forecasts — namely, ordered response models such as logit and probit. We outline the details of these alternatives in Appendix B.5. Panels B.1 and C.1 in Table 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 we 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

[18]Results are fully directionally consistent even if we 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.

based on logit or probit and the observable survey responses. 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, our exercise shows that forecasts based on our quantification methodology are fully consistent with the direction of sales growth implied by the qualitative survey responses. Furthermore, our estimates massively outperform alternatives based on ordered response models. This is strong evidence for the accuracy of our quantification methodology. We next turn to a Monte Carlo exercise that uses simulated data to infer how precisely our estimates match the magnitude of underlying true forecast errors.

### 3.3.2 Matching the Magnitude of Forecast Errors.

**Monte Carlo Experiment.** It is important to understand how well forecast errors based on our methodology match, in terms of magnitude, true quantitative forecast errors. In practice, this is challenging to do due to the unavailability of data on quantitative firm-level expectations. This dearth of data was, indeed, the key motivation for developing the quantification methodology proposed in this paper. The vast majority of surveys contain 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. To overcome this obstacle, we perform a Monte Carlo exercise that provides a benchmark based on simulated data. In particular, we simulate data on firm (continuous) annual sales growth realizations, as well as corresponding qualitative and quantitative expectations. We then use the data on realized sales growth and qualitative expectations as inputs to the quantification methodology of Section 3.1 and generate estimates for quantified sales growth expectations. Subsequently, we evaluate the accuracy of the estimated forecast errors in comparison to those based on the underlying artificial 'true' data.

We 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. Details about the data generation are provided in Appendix B.8. 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. We 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 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 only about one percentage point of sales growth. This is very small, particularly when recalling from Figure 1 and Table 2 that the absolute median forecast error in our data is three percentage points and the empirical distribution has non-negligible mass at forecast error values as large as 50 percentage points of sales growth. In general, the distribution of the difference between the artificial true and the artificial estimated forecast error is rather tight. For the $10^{th}$ ($90^{th}$) percentile the difference is 5 (7) percentage points; and even for the $5^{th}$ ($95^{th}$) percentile, it is still reasonably small at 7 (9) percentage points.

Table 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) |

We 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.

The close correspondence between the estimated and the true forecast error can also be illustrated in a scatter plot. Figure 2 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.

Anticipating our empirical results, the analysis in Section 4 endogenously establishes 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 therefore important that we check whether

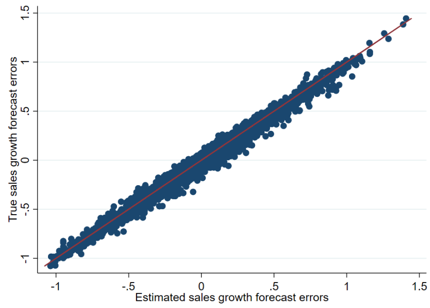


Figure 2: **Pairs of true and estimated sales growth forecast errors based on artificial data.** The figure shows all points in the dataset (we randomly selected one of the 1,000 draws for the datasets). The $45^o$ line is shown in red.

a forecast error based on our methodology is mapped into the same segment of the corresponding 'true' forecast error (both based on artificial data). We find this is the case for the vast majority, 94% (0.007 standard deviation across all across all 1,000 draws), of such pairs.

**Validating Forecast Error Accuracy in a Sample of UK Firms.** As explained above, the vast majority of firm-level surveys contain only qualitative questions. If surveys have quantitative features, these are typically limited to a small number of specific variables. While this highlights the importance of developing methodologies to quantify qualitative survey responses, it makes it difficult to validate our methodology against survey-based forecast errors. In principle, we can do so if a dataset comprises firm-level information on (1) monthly qualitative survey based forecasts for the three-month period ahead, (2) quantitative annual survey forecasts, and (3) annual realizations of the underlying variable. This information is available for a very limited sample of firms in the UK manufacturing sector. To the best of our knowledge, it is the only dataset that comprises all three types of data required to inspect the accuracy of our methodology. In particular, we consider quantitative one-year ahead forecasts on firm's own turnover growth from the Management and Expectations Survey which was conducted by the Office for National Statistics (ONS) in 2016. During the same year, the Confederation of British Industry (CBI) independently collected qualitative monthly survey forecasts on firm's output growth.[19] To obtain the annual realizations on turnover

[19]Details about the ONS survey can be found in Awano et al. (2018). The CBI's survey forecasts are similar to the ones from the IOBE for Greece as they are also used for the business climate of the Directorate-General for Economic

Table 7: Distribution of the difference between the estimated quantitative forecast error and the observed quantitative forecast error in a sample of firms in the UK manufacturing sector.

| 5% | 10% | 25% | Median | Mean | 75% | 90% | 95% |
|---|---|---|---|---|---|---|---|
| -0.135 | -0.128 | -0.052 | -0.001 | -0.004 | 0.052 | 0.098 | 0.129 |

growth we match the survey data with the Financial Statements from Bureau Van Dijk's FAME dataset.[20] Since the ONS and CBI surveys are conducted independently, the resulting matched sample is very small. It consists of 173 firm-month observations for qualitative survey forecasts on output growth, and 47 observations for annual quantitative forecasts on turnover growth and the corresponding realizations.

We implement our quantification methodology as follows. In the interest of statistical power, we fit the non-linear equation (10) to the realized turnover growth from FAME for a sample of 2,502 firm-year observations, for the period from 2000 until 2016. We then compute the forecast errors according to the methodology outlined in Section 3.1 for the 47 firms for which qualitative monthly survey forecasts are available. We compare these forecast error estimates with the quantitative forecast errors from the ONS survey. The distribution of the differences between the estimated and survey-based forecast errors is summarized in Table 7. Both, the mean and median of this distribution are very close to zero. Given that the ONS survey-based forecast errors have a mean of zero and a standard deviation of 0.31, the overall distribution for differences in forecast errors shown in Table 7 is rather tight. This is striking also because the monthly survey question is concerned with output growth and the annual survey question with turnover growth, which are closely related, but may not be perceived by respondents as exactly equal.

Overall, we have shown in the first exercise in this section that our estimates for quantified sales growth forecasts are fully consistent with the qualitative information contained in the underlying survey data. The Monte Carlo exercise has further demonstrated that our quantitative forecast error

and Financial Affairs (see DGECFIN (2017)).

[20]We thank Nick Bloom, Paul Mizen, Rebecca Riley and Michael Mahony for sharing the survey data and linking tables.

estimates are highly accurate and it is reassuring that this has also been confirmed using firm-level survey data from the UK manufacturing sector. Next, we proceed to use these quantified forecast errors in studying firm expectation formation.

# 4 Predictability and Autocorrelation of Forecast Errors

In this section, we study how the size of firms' forecast errors can affect results on their predictability and serial correlation. Crucial for this investigation is the quantification of forecast errors in the previous section. We start our analysis with the predictability of forecast errors.

## 4.1 Predictability of 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, we 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 behavioural biases or noisy signals may affect firms' weight attached to the lagged value, i.e. $z_t^e = \rho \Lambda z_{t-1}$.[21] 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 over-weight (under-weight) the lagged values of the predicted variable.

We estimate the extrapolation bias, $\varphi \triangleq \rho(1-\Lambda)$, using the following equation

$$x_{iy}^{fe} = \varphi x_{i,y-1} + \Psi_y + \Psi_i + \eta_{iy}, \tag{12}$$

where $\Psi_i$ and $\Psi_y$ control for unobserved firm heterogeneity and aggregate annual effects, respectively, and $\eta_{iy}$ is an idiosyncratic error. If $\varphi$ is statistically significant, firms' forecasts extrapolate incorrectly. To evaluate the effects of major forecast errors on sales growth, we further estimate the following

[21] Here, $\Lambda > 1$ could capture behavioural biases as for example in Bordalo et al. (2018b). A parameter $\Lambda < 1$ is in line with noisy signals as in Gabaix (2014).

threshold regression which allows for shifts in the extrapolation bias

$$x_{iy}^{fe} = \varphi_1 x_{i,y-1} * (1 - FEL_{i,y-1}^q) + \varphi_2 x_{i,y-1} * FEL_{i,y-1}^q + \varphi_3 FEL_{i,y-1}^q + \Psi_y + \Psi_i + \eta_{iy}, \tag{13}$$

where $FEL_{iy}^q$ 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, we call all forecast errors in the center of the distribution minor forecast errors. The extrapolation bias for minor forecast errors is $\varphi_1$, whereas following a major forecast error, that bias is $\varphi_2$. 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. $\varphi_3$ indicates whether the occurrence of a major forecast error has any effect on the forecast error in the following period.

We estimate equation (13) using a slightly adapted version of the Dynamic Panel Threshold estimator of Seo and Shin (2016). The original estimator is widely used in applications with thresholds (see e.g. Asimakopoulos and Karavias (2016) and Polemis and Stengos (2019)) and consists of two steps. The first step involves estimating equation (13) 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.[22] The original Threshold estimator of Seo and Shin (2016) uses Arellano and Bond (1991) First-Difference GMM (FD) for this estimation. The first-differencing results in loss of observations which is a substantial problem in severely unbalanced panels such as ours (see e.g. Roodman (2009) and Gorbachev (2011)). Instead we use the Arellano and Bover (1995) Forward Orthogonal Transformation (FOT) GMM to estimate the equation, which is the only difference to the original Seo and Shin (2016) Threshold estimator. The FOT subtracts from each observation the firm-specific arithmetic mean of its future values to eliminate the firm fixed effects, and hence avoids the loss of observations through first-differencing.

In the second step, in line with the original Seo and Shin (2016) estimator, for all values of $q\%$ we find the one that minimizes the objective function which then determines the final estimates for

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

$\varphi_1$, $\varphi_2$ and $\varphi_3$, as well. For our baseline estimation of equation (13) we use the Arellano and Bover (1995) FOT GMM estimator with collapsed instruments, limited lag length and Windmeijer (2005) corrected standard errors as it is standard in the literature (see e.g. Gorbachev (2011) and Caselli and Tesei (2016)).[23] Our approach significantly limits the risks of data loss by the use of FOT, and the risk of over-identification bias by avoiding the 'proliferation of instruments' in our sample through collapsed instruments and limited lag length (Roodman (2009)).[24] Finally, we also estimate the linear equation (12) with the same methodology, Arellano and Bover (1995) Forward Orthogonal Transformations GMM.

In Table 8, we document the results from the estimation of equations (12) and (13). Column (1) reports results from the simple linear equation and documents a highly significant negative estimate for the coefficient on $x_{i,y-1}$. This violation of the efficiency property of the FIRE hypothesis, and particularly the result that firms underweight the lagged variable, is in line with findings in the literature, which document that firms' forecast errors are predictable by past realizations (see e.g. Gennaioli et al. (2016), Massenot and Pettinicchi (2018) and Bordalo et al. (2018b)).

Column (2) in Table 8 reports estimation results for equation (13). The threshold for major forecast errors is estimated to include those observations at the top and bottom $q = 26\%$ of the distribution. These forecast errors are substantial and economically significant. A forecast error at 26% (74%) of the distribution implies that sales growth was expected to be 14.3 (8.6) percentage points higher (lower) than the subsequent realization.[25] Importantly, the coefficient of the lagged

[23]We use the Windmeijer (2005) corrected standard errors as standard errors might otherwise be biased downwards since we have a large number of instruments compared to the number of firms for inference.

[24]If a GMM system is excessively over-identified, the estimated coefficients are biased on the direction of the Nickell (1981) bias and the Hansen statistic is also biased. To limit the over-identification bias, we collapse the instruments, and only use five lags of instruments length. Our instruments are the lagged right hand side variables. We discuss our choice of instruments in Appendix C.1, where we also provide evidence that our results are robust to using fewer lags.

[25]We attempted to estimate the threshold equation with asymmetric thresholds for upper and lower cut-off %. 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 variable that we have on the RHS, we 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 RHS

realization is now only statistically significant when it is interacted with $FEL^q_{i,y-1}$. Our results in column (2) of Table 8 show that firms form biased predictions on sales growth and violate the FIRE hypothesis only following major forecast errors. In fact, only major forecast errors are negatively correlated with past sales growth. Unless firms make these major forecast errors, their predictions are more in line with the FIRE hypothesis as the estimate on the interaction with $(1 - FEL^q_{i,y-1})$ is not statistically different from zero. The Hansen p-value and the Arellano-Bond test of serial correlation of order two (m2 test) are both substantially larger than 0.1 and hence strongly reject the null that the specification is weak. This indicates that the non-linearity indeed exists and our specification is valid. Appendix C.1 shows our results are robust to using fewer lags. We also document why the original Seo and Shin (2016) FD GMM would be unsuitable with our data.

Table 8: 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, $x^{fe}_{iy}$ | |
| $x_{i,y-1}$ | -0.161*** | – |
| $x_{i,y-1} * (1 - FEL^q_{i,y-1})$ | – | -0.0583 |
| $x_{i,y-1} * FEL^q_{i,y-1}$ | – | -0.146** |
| $FEL^q_{i,y-1}$ | – | 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 (12) without the threshold. Column (2) is the Dynamic Panel Threshold estimator of Seo and Shin (2016) using the Arellano and Bover (1995) FOT GMM for equation (13). 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. We 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, $\overline{x}_{IND,y}$. $x^{fe}_{iy}$ is the forecast error of sales growth for year $y$; $x_{i,y-1}$ is lagged realized sales growth. $FEL_{iy}$ 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 Autocorrelation of Forecast Errors

Under the full information rational expectations hypothesis, forecast errors should be neither predictable by past realizations nor serially correlated. In this section, we turn to the latter and show

variables.

that again our findings crucially depend on the size of forecast errors.

To examine the autocorrelation of the quantified forecast errors of sales growth, we estimate the equation

$$x_{iy}^{fe} = \varrho x_{i,y-1}^{fe} + \Psi_y + \Psi_i + \eta_{iy}, \tag{14}$$

where $\varrho$ is the autocorrelation coefficient, $\Psi_i$ and $\Psi_y$ control for unobserved firm heterogeneity and year fixed effects, and $\eta_{iy}$ is an idiosyncratic error.

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

$$x_{iy}^{fe} = \varrho_1 x_{i,y-1}^{fe} * (1 - FEL_{i,y-1}^{q}) + \varrho_2 x_{i,y-1}^{fe} * FEL_{i,y-1}^{q} + \varrho_3 FEL_{i,y-1}^{q} + \Psi_y + \Psi_i + \eta_{iy}, \tag{15}$$

where $FEL_{iy}^{q}$ 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 $\varrho_1$, while following a major forecast error, forecast errors are autocorrelated with coefficient $\varrho_2$. If only $\varrho_2$ is statistically significant for the estimated threshold $q\%$, then forecast errors show persistence only following a major forecast error.

We estimate equations (14) and (15) using the exact same estimators and specifications as for the corresponding equations (12) and (13) on forecast error persistence. Note though that the threshold value in equation (15) is estimated endogenously and independently of the estimation of equation (13). Table 9 shows the estimation results for the former two equations. Column (1) reports that based on the simple linear equation, forecast errors are negatively autocorrelated.[26] This violates the FIRE hypothesis as firms fail to incorporate all new information to their forecasts, for example because they may be inattentive to new information. While estimates of this simple regression are indicative, we found in the previous section that only major forecast errors are predictable, very much in contrast to the result for minor forecast errors. For this reason we estimate the threshold

[26] In the literature (see e.g. Tanaka et al. (2019)) unobserved firm heterogeneity is usually dealt with using the within estimator. To avoid possible bias when the panel dimension is small (Nickell (1981)) we use the unbiased, consistent and efficient Arellano and Bover (1995) Forward Orthogonal Transformations (FOT) GMM.

equation (15) and report results in column (2). The results are consistent with the ones in the previous section. Only major forecast errors at the tails of the distribution are autocorrelated and violate the FIRE hypothesis. The estimate on the coefficient of $x^{fe}_{i,y-1} * FEL^{q}_{i,y-1}$ is highly significant and suggests a negative autocorrelation of forecast errors. In contrast, the estimate on the coefficient of $x^{fe}_{i,y-1} * (1 - FEL^{q}_{i,y-1})$ is not significantly different from zero so that minor forecast errors are not autocorrelated. It is reassuring that the cut-off $q = 26\%$ for the threshold is exactly the same as the one, independently estimated, for the equation on forecast error predictability. For the non-linear threshold model, both the Hansen p-value and the Arellano-Bond test of serial correlation of order two (m2 test) strongly reject the null that the specification is weak. Similarly to the predictability equation, this indicates that our model specification is valid, i.e. the non-linearity indeed exists. In Appendix C.2 we show that our results are robust in multiple dimensions, including the use of different lag length.

Table 9: 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, $x^{fe}_{iy}$ | |
| $x^{fe}_{i,y-1}$ | -0.164*** | – |
| $x^{fe}_{i,y-1} * (1 - FEL^{q}_{i,y-1})$ | – | 0.213 |
| $x^{fe}_{i,y-1} * FEL^{q}_{i,y-1}$ | – | -0.167** |
| $FEL^{q}_{i,y-1}$ | – | 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 (14) without the threshold. Column (2) is the Dynamic Panel Threshold estimator of Seo and Shin (2016) using the Arellano and Bover (1995) FOT GMM for equation (15). 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. We 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, $\overline{x}_{IND,y}$. $x^{fe}_{iy}$ is the forecast error of sales growth for year $y$; $x_{i,y-1}$ is the lagged realized sales growth. $FEL_{iy}$ 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.

**Autocorrelation and Predictability of Qualitative Forecast Errors in the Survey Data.**

Even though we have shown in Section 3.3 the accuracy of the forecast error estimates in a number of ways, one may wonder whether the empirical findings we obtained using our quantification methodology are also a feature of the qualitative survey data. Below, we provide evidence that this is the case — our findings also hold in the directly observable survey-based forecast errors which suggests that our results are not driven by the imputation of the quantified forecast errors.

We construct monthly forecast errors from qualitative expectations data in a manner employed by Bachmann et al. (2013) and Massenot and Pettinicchi (2018). We then document that large forecast errors are autocorrelated and predictable only in years that our quantification methodology flags as involving a major forecast error. Since the survey responses are qualitative, binary choice models permit us 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, we can construct monthly forecast errors with values $XS_{im}^{fe} = \{-2, -1, 0, +1, +2\}$.[27] For the outcomes labelled $-2$ and $+2$, the firm's forecast completely mispredicted the direction of change of its sales and those errors are likely large. We label them $XS_{im}^{Lfe}$, and analyze below whether their occurrence is predictable or autocorrelated.

We identify years in which firms made a major forecast error from the annual quantified values based on the threshold regressions contained in Tables 8 and 9. The months that belong to these years are assigned $FEL_{iy} = 1$. To test the predictability and the autocorrelation of large survey forecast errors, we use probit models that directly correspond to the continuous regressions (13) and (15) above. Details about the estimations are contained in Appendix C.3. Table 10 shows that only during years when annual quantified forecast errors are classified as major (i.e. $FEL_{iy} = 1$) are survey-based forecast errors predictable and autocorrelated. The extrapolation bias (in Panel A) and the persistence coefficient (in Panel B) in years with minor forecast errors (i.e. $(1 - FEL_{iy}) = 1$) are insignificant. These results based on the monthly qualitative survey forecasts are consistent with the ones we obtained from the annual quantified forecasts.[28]

[27] Compare the survey questions A.2 and D.2 in Section 2 where *increased/unchanged/decreased* responses are labelled as -1/0/+1. Essentially, these responses indicate the direction of change of sales, expected and realized.

[28] Even though equations (13) and (15) are analogous to the multinomial probits of Table 10, the signs of the coefficients are not comparable. The reason is that the multinomial probit models the probability of a (major) forecast error and not its size. The statistical significance of the coefficients of the dependent variable is what shows the

Table 10: Predictability and Persistence of firms' forecast errors of sales growth in the qualitative survey data. Probit Estimates.

| **Panel A:** Predictability | | **Panel B:** Autocorrelation | |
|---|---|---|---|
| $XS_{im} * (1 - FEL_{iy})$ | -0.0398 | $XS^{Lfe}_{i,m-3} * (1 - FEL_{iy})$ | -0.0554 |
| $XS_{im} * FEL_{iy}$ | -0.133*** | $XS^{Lfe}_{i,m-3} * FEL_{iy}$ | 0.326*** |

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. $FEL_{iy}$ takes value one when the annual quantified forecast error lies in the lower or upper 26% of its empirical pool distribution. The estimation includes fixed year effects, proxies for firm fixed effects and proxies for initial conditions, but are omitted here to simplify the exposition. The detailed estimations are in Table 15.C in Appendix C.3. *** indicates statistical significance at the 1% level.

Overall, we have documented in this section that with respect to autocorrelation and predictability, only for major forecast errors do firms violate the FIRE hypothesis. For smaller (in absolute value) forecast errors we find firms' forecasts are more in line with the FIRE hypothesis. One explanation for the violation of this hypothesis can be that when firms make major forecast errors the underlying forecasts are based on a limited information set, while smaller absolute forecast errors are based on (nearly) full information sets. In the next section, we rationalize our empirical findings in a model where the quality of firm's forecasts on sales growth depends on the potentially costly level of attention to information on current sales.

# 5 Model of a Firm with Rational Inattention

In this section, we 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 rational inattention models in 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. We subsequently show that a simple model with limited attention to information and variations in the cost for attentiveness can rationalize the empirical findings of Section 4.

violation of the FIRE in both cases.

## 5.1 Forecasts in a Simple Signal-Extraction Framework

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 $y$. We abstain from a subscript $i$ for the remainder of this section to ease notation. We assume sales growth follows an AR(1) process,

$$x_y = \rho_0 + \rho x_{y-1} + u_y, \tag{16}$$

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, we assume henceforth that $\mu = 0$. Finally, we assume that the shock, $u_y$, and the noise term, $\epsilon_y$, 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 $x_y$.[29] Then the optimal forecast, $x_{y+1}^e$, is[30]

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

The first order condition yields $x_{y+1}^e = \mathbb{E}\big[x_{y+1}|s_y\big]$ and using the fact that $x_{y+1}$ follows the AR(1) process (16), we obtain

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

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 $\epsilon_y$ 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 $\mathbb{E}\big[x_y|s_y\big]$,

$$x_{y+1}^e = \rho\mathbb{E}\big[x_y|s_y\big] = \rho\lambda_0 + \lambda\rho s_y, \qquad \text{where} \qquad \lambda = \frac{Cov(x_y, s_y)}{V(s_y)}, \quad \text{and} \quad \lambda_0 = (1-\lambda)\mu = 0. \tag{17}$$

[29] 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.

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

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

$$\lambda = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\epsilon^2}. \tag{18}$$

Equation (18) shows that in the presence of noise, $\sigma_\epsilon^2 > 0$, $\lambda$ is strictly between 0 and 1. This has implications on the optimal forecast (17), which is, when applying the definition for the signal,

$$x_{y+1}^e = \lambda\rho s_y = \lambda\rho x_y + \lambda\rho\epsilon_y. \tag{19}$$

Equation (19) 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 (19) 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, we will endogenize this choice. Then the firm can choose its level of attention to information by determining the parameter $\lambda$; which is equivalent to choosing the information content in the signal by varying the noise variance $\sigma_\epsilon^2$. 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 $0 < \lambda < 1$. We 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, $\lambda$.

## 5.2 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 $\lambda$ 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 latest noisy signal — as before this signal is the sum of actual sales growth and the noise.[31] The full information optimal forecast ($\lambda = 1$) would be $x^e_{y+1} = \rho s_y$ which is in line with the underlying AR(1) process (16) 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. We 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 $\lambda$. If we substitute the optimal forecast, $x^e_{y+1}(\lambda)$, into the utility function, we obtain the indirect utility function

$$U(\lambda) = W\left(x^e_{y+1}(\lambda), \lambda s_y\right), \tag{20}$$

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

$$C(\lambda, c_y) = c_y K(\lambda), \tag{21}$$

[31] As in the simple model above, also in this extended model the firm only relies on current information to make a forecast. The underlying assumption is that processing information on past signals or past relizations is just as costly per unit as processing current information. Hence, it would always be optimal for the firm to rely on the most up to date information for the forecast. This assumption is similar to one made in Mackowiak and Wiederholt (2009) — who assume that past realizations of the state variable are never observed — and eases the exposition of our model.

where $K(\lambda)$ is a continuous increasing and convex function in $\lambda$. 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.[32] 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 $\lambda^*$, and conditional on this choice, it obtains in a second step the optimal forecast for sales growth.[33] We 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 (20) and the cost function (21). 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 D.1) that the firm obtains the optimal level of attention, $\lambda^*$ by solving the following intratemporal problem

$$\lambda_y^* \triangleq \arg\max_{\lambda} \Big[ -\frac{1}{2}\sigma_s^2(1-\lambda)^2 - c_y K(\lambda)\Big], \tag{22}$$

where $\sigma_s^2$ 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[34]

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

where $K'(\cdot)$ denotes the first derivative of $K(\cdot)$. Our results that follow in this section below do not require us to specify a particular functional form for $K(\cdot)$.[35] However, to briefly provide intuition about how the optimal level of inattention depends on the information cost and the variance of the

[32]We 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 (23) below. The only requirement on the positive upper bound on $c_y$ is that it is specified to ensure that $\lambda > 0$.

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

[34]Note, it is satisfied only for $c_y > 0$ and $0 < \lambda^* < 1$. When $c_y = 0$, then optimal attention is equal to 1.

[35]In fact, our assumptions on $K(\lambda)$ are consistent with several specific functional forms used in the literature. For example $K(\lambda) = \frac{1}{2}log_2((1-\lambda)^{-1})$, in the tradition of Sims (2003), would be a micro-founded functional form based on the Shannon entropy. Alternative functional forms could be based e.g. on Caplin and Dean (2015) or Gabaix (2014).

signal, we specify $K(\lambda) = \lambda^a$ where $a \geq 1$. Then the first order condition has, for the cases $a = 1$ and $a = 2$, the following simple analytical solutions

$$\begin{aligned}
\lambda_y^* &= \frac{\sigma_s^2 - c_y}{\sigma_s^2}, \qquad \text{for} \quad a = 1. \\
\lambda_y^* &= \frac{\sigma_s^2}{\sigma_s^2 + 2c_y}, \qquad \text{for} \quad a = 2.
\end{aligned} \tag{23}$$

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 $a$, we assumed an upper bound for $c_y$ that guarantees $0 < \lambda_y^* < 1$.

Having chosen the optimal level of attention via (22), the firm's optimal forecast is given by

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

so that the optimal forecast is

$$x_{y+1}^e = \lambda_y^* \rho s_y. \tag{24}$$

As in the simple signal extraction problem above, the forecast understates the persistence of the signal on 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, we extended the simple framework of Section 5.1 so that the firm may pay limited attention to information. This will be key to explaining our empirical facts on predictability and autocorrelation of forecast errors, which we will show next.

**The Size of Forecast Errors, their Predictability and Autocorrelation.** Next, we 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 (16) and the optimal forecast (24), the ex-post forecast error in the framework with rational inattention is given by

$$x_{y+1}^{fe} \triangleq x_{y+1} - x_{y+1}^e = (1 - \lambda_y^*)\rho x_y - \lambda_y^* \rho \epsilon_y + u_{y+1}, \tag{25}$$

where we used that $s_y = x_y + \epsilon_y$. We will use this equation to derive three results from our 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, $\epsilon_y$, is zero for $\lambda^* = 1$ as implied by equation (18). A positive cost, $c_y > 0$, reduces $\lambda_y^*$ 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 $\lambda^* = 1$.[36] 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 our empirical results in Section 4.1 on differences between major and minor forecast errors. In this section, we document that the estimated coefficient on past sales growth in equation (13) — which corresponds to $(1-\lambda_y^*)\rho$ in model equation (25) — is significantly different from zero for large absolute forecast errors. Hence, for these major forecast errors it implies $0 < \lambda_y^* < 1$.[37] For minor forecast errors though, we find the coefficient estimate is not statistically significant anymore which implies that $\lambda_y^*$ is (close to) unity.

[36] 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.

[37] It will become clear in the discussion of Result 2 why the coefficient estimate for large forecast errors is negative in equation (13).

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 (25) 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 $\lambda_y^*$ reduces from unity to positive values strictly lower than one. For $\lambda_y^* = 1$, the forecast error as given in equation (25) 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 $(1 - \lambda_y^*)\rho x_y$. The forecast error can only be negatively correlated with lagged sales growth if the coefficient $(1 - \lambda_y^*)\rho$ in equation (25) is negative, which only is the case if $\rho < 0$. In Appendix B.7 we provide empirical evidence from our 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 (13) in Section 4.1. In this section, we document the predictability of major forecast errors as well as a negative relation between these major forecast errors and lagged sales growth. We 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 (25) and using that $x_y^e = \lambda_{y-1}^* \rho s_{y-1}$ as well as the definition of the signal, we 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 $\lambda_y^* = 1$. We know from the discussion of 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 our empirical findings. The estimation results of equation (15) in Section 4.2 show that major forecast errors are negatively autocorrelated. We further document in that section that the autocorrelation of minor forecast errors is not significantly different from zero.

Overall, our 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 our theoretical model are consistent with our empirical results documented in Section 4. The model can also rationalize our negative estimates of the coefficients on persistence and autocorrelation for major forecast errors. We show that the negative sign of these estimates is the result of the negative autocorrelation of sales growth in our data.

The above has shown that, despite its simplicity, our model is able to rationalize our main empirical findings. Key for the model results to hold are variations in firm's optimal level of rational inattention, $\lambda_y^*$, 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. Our dataset can provide some first guidance. The empirical evidence in Section 3.2 documents that major forecast errors are not specific to a particular time, sector 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 our case, be less likely to capture macroeconomic or low-frequency shocks, but may be linked to high-frequency effects.[38] A variety of reasons, and a combination of these, can be behind the latter, for example changes in the specific market environment, regulatory

[38] 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 ours. 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) or Acharya and Wee (2020).

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 our aim was to develop a parsimonious model to rationalize our empirical results, a model with a fully endogenous cost function goes beyond the scope of this paper and we leave it for future research.

# 6 Conclusion

In this paper we 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, we have developed a novel methodology to quantify qualitative survey data on firm forecasts. This methodology is applicable generally when quantitative information is available on the realization of the forecasted variable. As an example, all European Union countries run a survey of firm expectations similar in structure to the Greek one that we use here. This survey data could be combined with firm balance sheet data to produce quantified expectations estimates using our methodology. In order to interpret our empirical results that show that the Full Information Rational Expectations hypothesis is violated, we also provide a model of rational inattention. 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.

Some questions emerge naturally from these findings. For example, under which circumstances do firms make major forecast errors and how do these affect firm decisions? Our unique dataset together with our novel methodology to quantify forecast errors is highly suitable to answer such questions. In a companion paper (Botsis et al. (2020)), we analyze the causes of major forecast errors and their effects on firm production, investment, and financing decisions. Our aim is to explore the underlying market environments that result in different degrees of limited attention. This could help in the appropriate calibration of the information cost in rational inattention models.

Provided that major forecast errors lead firms to make suboptimal decisions, a question that arises is whether policy design can be geared to helping firms avoid these. Such policy would likely aim to limit uncertainty and stabilize expectations. This could involve a combination of transparency and stable rules. Clearly, this requires analysis with appropriate models and is a useful direction for further research.

# References


Acharya, S. and Wee, S. L. (2020). Rational inattention in hiring decisions. *American Economic Journal: Macroeconomics*, 12(1):1–40.

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.

Awano, G., Bloom, N., Dolby, T., Mizen, P., Riley, R., Senga, T., Van, J., Reenen, J. V., and Wales, P. (2018). A firm-level perspective on micro-and macro-level uncertainty. *ESCoE Discussion Pape*, 10.

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. (2013). What drives aggregate investment? Evidence from German survey data. *CESifo Working Paper Series*, (4218).

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., Bunn, P., Chen, S., Mizen, P., Smietanka, P., and Thwaites, G. (2019). Brexit as a misallocation shock. *University of Nottingham Mimeo.*

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2018a). Over-reaction in macroeconomic expectations. *National Bureau of Economic Research.*

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

Botsis, A., Görtz, C., and Sakellaris, P. (2020). Firm's decisions and large forecast errors. *University of Birmingham Mimeo.*

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.

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. (2018). How do firms form their expectations? New survey evidence. *American Economic Review*, 108(9):2671–2713.

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. (1981). Price expectations of business firms. *Brookings Papers on Economic Activity*, 1981(1):299–314.

DGECFIN (2017). The joint harmonised eu programme of business and consumer surveys: User guide. *European Commission: Directorate-General for Economic and Financial Affairs.*

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.

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.

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

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.

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.

Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Manski, C. F. (2004). Measuring expectations. *Econometrica*, 72(5):1329–1376.

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

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.

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

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.

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.

Tanaka, M., Bloom, N., David, J. M., and Koga, M. (2019). Firm performance and macro forecast accuracy. *Journal of Monetary Economics*.

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.

Vellekoop, N. and Wiederholt, M. (2019). Inflation expectations and choices of households. *SAFE Working Paper*.

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

# Appendix (for Online Publication)

# A Data

In the following we provide further details about the conduct of the IOBE survey (Section A.1) and the cleaning procedures on the IOBE and ICAP data (Sections A.2 and A.3). Section A.4 provides additional information about our matched sample and discusses representativeness and quality of survey responses.

## A.1 Details on the Survey Data

The firm-level survey data are collected every month by IOBE. IOBE sends surveys to a sample of firms included in the ICAP firm directory. This directory covers more than 75% of the economy's output. The sample of surveyed firms is chosen to represent the distribution of firm sizes in terms of gross value added in each 2-digit sector. Every 4-5 years it is replenished by removing those firms who never replied and those who have stopped replying. These are replaced with new firms, following the same sampling principles, while the firms that have been responsive are retained in the sample. According to IOBE researchers, the response rate is somewhat smaller than 20% which is in line with response rates of surveys conducted for the European Commission in other countries.

IOBE send surveys by mail and email between the 22th and the 25th of each month – surveys refer to the following month. More than 80% of firms that reply do so by the 15th of the month the survey refers to, and more than 95% reply by the 20th. Responses that arrive well past the month they refer to, are dropped by the IOBE as it is unclear to which month responses refer. Less than 10% of responses are received by email. The vast majority of surveys are completed on paper and returned by mail in a prepaid envelope. The IOBE requests that surveys are completed by managers or a person who has complete knowledge of the entire activity of the surveyed firm.

Surveys are conducted monthly with the exception of August. In August the majority of firms are closed as managers and employees take their annual leave. For this reason, there are no surveys sent out at the end of July to record the responses for August. IOBE uses imputation methods to produce data for August and for monthly non-responses.[39] We will remove imputed observations in the cleaning section A.2.

## A.2 Cleaning the Survey Data

The wording of the survey question is so that it asks about sales expectations for the next three months. This means expectations that include the last two months of a year would also be concerned with sales in the first one or two 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 months in the previous year. For this reason we make adjustments to the submitted responses on realizations and forecasts in the concerning months, which are standard treatment of survey data in the literature. For forecasts, we 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 we set the responses in January to missing and use this observation with weight 1 in the final month of the preceding year. We 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 relizations 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 our 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).

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

Finally, we have set to missing all firm-month observations in one particular year if we have less

[39] This is standard practice of survey providers. Lui et al. (2011) for example report that for the UK business climate survey, the Confederation of British Industry (who administer the survey) also implements imputation techniques for missing data.

than three monthly survey responses of this firm within the year. This was necessary because our quantification 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 us with 1,093 firms in the manufacturing sector that provide survey responses.

## A.3 Cleaning the Financial Statements Data

We have financial statements data available from ICAP. In the following we 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 we 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 $TotalNetAssets_{i,y} \neq TotalLiabilitiess_{i,y}$, we replaced their values with those from an alternative Balance Sheet data-base of Hellastat S.A.[40,41] We confirmed that for the replaced values of $TotalNetAssets_{i,y}$ and $TotalLiabilitiess_{i,y}$ 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, we 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, we replaced their values with

[40] 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 we can assume that the two data-bases do not include the same errors.

[41] 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 our ICAP data, but includes a less detailed break-down of financial statement variables.

those from Hellastat. Then we 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, we set to missing all the financial statement variables of these firm-year observations.

3. We 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. We 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 $y$. We 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 we 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 we used the implicit deflator of capital stocks from Eurostat Table nama_10_nfa_st.

7. In the final cleaning steps, we 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 we 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 we set to missing all the financial statement variables: Real Total Net Fixed Assets, Real Gross Depreciable Property, Real Total Sales.

## A.4 The Matched Sample and Quality of Survey Responses

We match firms' financial statements data with the corresponding survey responses using the firm's unique tax identifier. As described in Section A.2, our cleaned survey data comprised 1,093 firms. We could match 73.1% of these firms (76.7% of the firm-month observations), so that the sample for which we have both survey and financial statement data comprises of — after the cleaning procedures described above — 799 firms in the manufacturing sector with 25,764 monthly responses from the survey on the two questions A.2 and D.2 and 4,104 annual balance sheet observations on sales. This section first establishes that our sample is representative for the manufacturing sector. Then we evaluate the quality of survey responses.

**Representativeness.** We evaluate representativeness of our sample in a number of ways using data from the survey and the financial statements.

First, we report a time-series correlation of 0.95 between the official IOBE business sentiment index for the manufacturing sector and a recalculated sentiment index based on our manufacturing sector dataset.[42] This high correlation shows that our dataset is still highly representative when responses are aggregated, even though we abstain from using the imputed survey responses and we dropped observations if firms responded fewer than three times in a calendar year. Second,

[42]The monthly sentiment index for the manufacturing sector is computed as $\frac{QS_{im}+QS^{e}_{im}-INV_{im}}{3}$, where $INV_{im}$ corresponds to the question *'The level of finished goods inventories you deem it is...'* with the possible responses being above/at/below normal levels and coded as $+1/0/-1$, respectively; and $QS_{im}$ corresponds to the survey question *'For the preceding 3 months you assess that your production did...'*, $QS^{e}_{im}$ corresponds to the question *'For the next 3 months you foresee that your production will...'*, and the possible responses are rise/no change/fall, coded as $+1/0/-1$, respectively.

we report a correlation of 0.64 between the average real growth rate of output in our sample as reported in the financial statements and the corresponding manufacturing sector output growth from Eurostat.[43] We perform this comparison using output since Eurostat only publishes sales for the Greek manufacturing sector from 2008. Third, to further examine the representativeness of our final sample we study the share of each 2-digit sector in the total manufacturing sector sales. We compare the contributions based on our sample with the ones from the official Eurostat data. Table 1.A exemplifies these statistics for two years — 2009 and 2012 — and we observe that most of the shares based on our dataset are close to the ones reported by Eurostat with few exceptions of over- and under- representativeness.

Table 1.A: 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 our 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 our 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.

**Quality of Survey Responses.** In this section we first establish that the survey responses

[43]Output from the financial statements is the sum of sales plus the contemporaneous first difference of final goods inventories. We 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. We use the simple arithmetic mean of the firm-year observations to obtain the average growth rate of our sample. The manufacturing growth rate of real output from Eurostat for Greece is from Table nama_10_a64.

are consistent across different questions and then, we show they are consistent with data from the financial statements.

In the spirit of Coibion et al. (2015) we use a regression-based approach to evaluate the consistency of the survey responses across questions. We conduct two exercises to establish consistency that will jointly cover around two thirds of the survey questions. Turning to the first exercise, economic intuition suggests that if a firm expects excess future production capacity relative to sales, it is more likely to (i) report higher than normal inventory levels (ii) expect a drop in the sales (iii) expect it will have to decrease employment (iv) have lower capacity utilization that would allow it to increase production if need be. To confirm that this economic intuition holds in our data we estimate the following linear equation:

$$D3_{im} = \beta_0 + \boldsymbol{\beta}\Big[INV_{im}, XS^e_{im}, L^e_{im}, U_{im}\Big]' + \psi_i + \psi_y + \eta_{im}, \tag{26}$$

where the vector $\boldsymbol{\beta} = [\beta_1, \beta_2, \beta_3, \beta_4]$, $\psi_i$ and $\psi_y$ control for firm and year fixed effects respectively, and $\eta_{im}$ 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.[44]

We estimate equation (26) twice: first, by eliminating $\psi_i$ using standard fixed effects tools and second, by substituting NACE sector dummies for $\psi_i$. In Panel A of Table 2.A we report the results from estimating equation (26). We observe that the signs of the variables under examination are as expected based on the economic intuition outlined above and that all estimates are statistically significant at the 1% level. The relatively low $R^2$ indicates that there are other factors that explain

[44] 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'*. $XS^e_{im}$ refers to question D.2 outlined in the main body. $L^e_{im}$, 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.

expected movements in production capacities. However, for the purpose of verifying the consistency of survey answers we are only interested in the directional relationship between variables.

In the second exercise, we focus on production. When we observe an increase in production, economic intuition indicates one factor behind this could be a rise in capacity utilization. We check this by estimating the following linear equation:

$$QS_{im} = \beta_0 + \beta_1[U_{i,m-1} - U_{i,m-3}] + \psi_i + \psi_y + \eta_{im}, \tag{27}$$

where $U_{im}$ 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, $\psi_i$ and $\psi_y$ control for firm and year fixed effects respectively, and $\eta_{im}$ is the idiosyncratic error.[45] As previously, we estimate equation (27) in two ways: firstly, we eliminate $\psi_i$ using standard fixed effects tools and secondly, we substitute NACE sector dummies for $\psi_i$. Results are reported in Panel B of Table 2.A. These are in line with economic intuition: an increase in production is positively and significantly correlated with a reported three-month increase in capacity utilization (from $m-3$ to $m-1$) over the same time horizon.

Table 2.A: 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 pool OLS (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; $SS^e_{im}$ 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.

Having substantiated the consistency of survey responses across questions, we now turn to evaluating their consistency with the information in the financial statements. Annual sales growth of

[45] $QS_{im}$ corresponds to question A.1: *'During the previous 3 months, your production, has increased/remained unchanged/decreased.'*

firm $i$ in the income statements, $x_{iy}$, should be positively correlated with the survey question A.2 concerning the evolution of current sales, $XS_{im}$.

We examine this by estimating the following linear equations

$$XS_{im} = \beta_0 + \beta_1 x_{iy} + \psi_i + \psi_y + \eta_{im}, \tag{28}$$

where $\psi_i$ and $\psi_y$ control for firm and year fixed effects respectively and $\eta_{im}$ is the idiosyncratic error. As previously, we estimate equation (28) in two ways, using standard fixed effects tools or NACE sector dummies. To estimate this regression, the firm-year observations from financial statements are treated as the same for each month in a particular year. We can do so as the survey data is qualitative while the data from financial statements is quantitative and we are simply interested in a correlation between the two. In Table 3.A, we observe that the monthly responses are positively and highly significantly correlated with the growth rates from the financial statements. In other words, qualitative survey responses on changes in current sales and production are on average consistent with their quantitative counterparts reported in the financial statements.

Table 3.A: 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 ***.

Overall, based on the results in Tables 2.A and 3.A, we find that survey responses are consistent, both with each other within the questionnaire, but also with the information in the financial statements. In addition, the fact that survey responses are positively correlated with the corresponding financial statement variables is consistent with the information from IOBE that surveys are completed by executives who have a complete overview about the firm's activities. We can draw this conclusion, because the financial statements are published after the respondents fill in the survey.

# B Quantification of Forecast Errors

## B.1 Derivation of Equation (4)

This section shows how equation (4) can be derived using equations (1) and (3). First, we take expectations of equation (3), which becomes

$$\mathbb{E}\big[x_{im}^{e,+}|\mathcal{F}_{i,y-1}\big] = \alpha + \gamma_1 x_{iy}^e + \mathbb{E}_{i,y-1}\nu_{im}^+, \quad \text{and} \quad \mathbb{E}\big[x_{im}^{e,-}|\mathcal{F}_{i,y-1}\big] = -\beta + \gamma_2 x_{iy}^e + \mathbb{E}_{i,y-1}\nu_{im}^-. \tag{29}$$

Then, we substitute equation (29) into (1)

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

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

$$\begin{aligned} x_{iy}^e \;=\; & \big[\alpha + \gamma_1 x_{iy}^e\big]\mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=1]} + \mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=1]}\nu_{im}^+ \\ +\; & \big[-\beta + \gamma_2 x_{iy}^e\big]\mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=-1]} + \mathbb{E}_{i,y-1}\sum_{m\in y} W_{im}\mathbb{1}_{[XS_{im}^e=-1]}\nu_{im}^-. \end{aligned} \tag{30}$$

To simplify the notation, we define

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

where $P_{iy}$ ($N_{iy}$) denotes the weighted share of months within a year that indicate a rise (fall) in expected sales. Next, we 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 sales growth forecasts as was expected at the end of year $y-1$.[46] The assumption here is akin to saying that firms make their annual budget and related sales forecasts at the end of the year using information up to end of December. During the year they may update monthly forecasts and the budget, but not whether they expect positive or negative development in sales growth during the month. This allows us to rearrange equation (30) to solve for $x_{iy}^e$:

$$x_{iy}^e = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \qquad \text{with} \qquad \xi_{iy} = \frac{\mathbb{E}_{i,y-1}\sum_{m\in y}\big(W_{im}^+\nu_{im}^+ + W_{im}^-\nu_{im}^-\big)}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}},$$

which is equation (4) in Section 3.1.

[46] This assumption does not drive our results on the predictability and autocorrelation, because we show in Section 4.2 that this same behavior is also present in the directly observable survey based forecast errors.

## B.2 Proofs Related to the Estimation of Equation (6)

**Statement 1.** If $\mathbb{E}\big[x_{iy}^{fe}\big|x_{iy}^{e}\big] = 0$, then $\mathbb{E}\big[x_{iy}^{fe}\big|\mathcal{H}\big(x_{iy}^{e}\big)\big] = 0$ for any Borel measurable function $\mathcal{H}$. Therefore, $\mathbb{E}\big[x_{iy}^{fe}\big|\tilde{x}_{iy}^{e}\big] = 0$ ($\tilde{x}_{iy}^{e}$ is defined in equation (7) in the main text).

*Proof.* Firstly, note that the underlying mathematical form of this (and any) conditional expectation is $\mathbb{E}\big[x_{iy}^{fe}\big|x_{iy}^{e}\big] = \mathbb{E}[x_{iy}^{fe}|\sigma(x_{iy}^{e})]$, where $\sigma(x_{iy}^{e})$ is the minimal sigma-algebra generated by $x_{iy}^{e}$. Intuitively, all the information that $x_{iy}^{e}$ can convey. Then from the Doob-Dynkin Lemma (see Proposition 3 in Rao and Swift (2006)) we know that $\sigma\big(\mathcal{H}\big(x_{iy}^{e}\big)\big) \subset \sigma(x_{iy}^{e})$ for any Borel measurable function $\mathcal{H}$. As a result, from the general form of the Law of Iterated Expectations, we get $\mathbb{E}\big[x_{iy}^{fe}\big|\mathcal{H}\big(x_{iy}^{e}\big)\big] = \mathbb{E}\Bigg[\mathbb{E}\big[x_{iy}^{fe}\big|x_{iy}^{e}\big]\Bigg|\mathcal{H}\big(x_{iy}^{e}\big)\Bigg] = 0$. Next, we know that $\tilde{x}_{iy}^{e}$ is a Borel measurable function of $XS_{im}^{e}$ for $m \in y$.[47] Also, $XS_{im}^{e}$ is a Borel measurable function of $x_{iy}^{e}$ (from ID1).[48] Overall, we have that $\sigma(\tilde{x}_{iy}^{e}) \subset \sigma(\{XS_{im}^{e}\}_{m\in y}) \subset \sigma(x_{iy}^{e})$, for $m \in y$. Therefore, $\mathbb{E}\big[x_{iy}^{fe}\big|\tilde{x}_{iy}^{e}\big] = 0$. This completes the proof.

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

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

[47] This follows from the fact that $\tilde{x}_{iy}^{e}$ 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 $XS_{im}^{e}$: they are linear (continuous) functions of $P_{iy}$ and $N_{iy}$ which are also linear functions (continuous) of $XS_{im}^{e}$.

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

**Statement 3.** The error term $\tilde{\xi}_{iy}$ in equation (10) is mean-independent of the explanatory variables.

We provided a way to approximate the unobserved firm heterogeneity, and we derived the final estimable equation (10). For equation (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 (10). This means that the NLS error $\tilde{\xi}_{iy}$ is also mean independent of the rational function on the right hand side of (10), which satisfies Davidson and MacKinnon (2004)'s condition for consistency (equation (6.29)). Indeed, from equation (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 we have that $\mathbb{E}\big[x^{fe}_{iy}|\{XS^e_{im}\}_{m\in y}\big] = 0$. Note that $\{XS^e_{im}\}_{m\in y} \subset \{XS^e_{im}\}_{m=1,2,\dots,T_i}$ which implies that $\sigma\big(\{XS^e_{im}\}_{m\in y}\big) \subset \sigma\big(\{XS^e_{im}\}_{m=1,2,\dots,T_i}\big)$. Therefore, from ID2, ID3 and the Law of Iterated Expectations we have that $\mathbb{E}\big[\omega_i + \vartheta_{iy}\big|\{XS^e_{im}\}_{m\in y}\big] = 0$. This completes the proof.

## B.3 Estimation Results and Robustness

**Nonlinear Least Squares Estimation.** Table 4.B reports the results of the NLS estimation of equation (10). Column (1) shows estimation results for the boom period up to 2008 and column (2) for the following recession. As a reminder, $\alpha$ and $-\beta$ are the constant terms in the positive and negative continuous monthly forecasts of ID1. We observe that the constant of the positive monthly forecasts is larger during the boom than in the bust which is consistent with our 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.

Table 4.B: NLS Estimation of Equation (10).

| | (1) | (2) |
|---|---|---|
| Coefficients | Dependent Variable: $x_{iy}$ | |
| $\alpha$ | 0.190** | 0.104** |
| $\beta$ | 0.151* | 0.238*** |
| $\gamma_1$ | -0.366 | -0.446 |
| $\gamma_2$ | -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 (10) are omitted – but are included in the estimation – to maintain a simple representation. We 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.

**Alternative Weighting Scheme.** This section shows results based on an alternative weighting scheme used in equation (2). In particular, while our baseline weighting controls for seasonalities within the year, we consider as an alternative that all observations are weighted equally per year.

Table 5.B reports results of the estimation of equation (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 4.B. From Table 6.B it is evident that this close resemblance also results in almost identical forecasts. The 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 5.B: NLS Estimation of Equation (10) with alternative weighting.

| | (1) | (2) |
|---|---|---|
| Coefficients | Dependent Variable: $x_{iy}$ | |
| $\alpha$ | 0.197** | 0.105** |
| $\beta$ | 0.158* | 0.243*** |
| $\gamma_1$ | -0.426 | -0.479 |
| $\gamma_2$ | -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 (10) are omitted – but are included in the estimation – to maintain a simple representation. We 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 6.B: 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 |

## B.4 Relation between Monthly and Annual Survey Forecasts

We can test ID1 if a dataset includes qualitative monthly and the corresponding quantitative annual firm-level survey expectations. Availability of this information in one dataset is seldom, yet we have constructed data with these features for UK manufacturing firms in Section 3.3.2. In this dataset we find support for ID1 as the monthly forecasts are positively correlated (0.43 at 1% significance) with their annual counterparts. Additionally, given the qualitative nature of the survey forecasts, we estimate an ordered probit model of the survey forecasts on the observable quantitative forecasts. This ordered probit model also verifies ID1 (Wald chi-squared statistic for the whole ordered probit model with 1 degree of freedom equals 37.46).

## B.5 Alternative Quantification Techniques

Ordered response models — probit or logit — are alternatives to the NLS based method outlined in Section 3.1 to quantify sales growth forecasts.

For the ordered response models, we 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 $\psi_i$ 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. We can eliminate the unobserved heterogeneity $\psi_i$ 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, we 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 $X_{im}^{XS}$.[49] Finally, we 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 (2).

Table 7.B reports the estimation results of the ordered probit and logit models. The variables that we 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.[50]

## B.6 Statistics on Forecasts and Forecast Errors

**Statistics on Survey Forecasts.** This section provides an overview about the information on sales forecasts in the survey. The left subplot of Figure 1.B shows the distribution of monthly responses to survey question D.2 on firms' expected sales during the next three months. These possible re-

[49] 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 we would introduce information to the firm's forecasts that were not available to the firm at the time of the forecast.

[50] 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 7.B: 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. We use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance.

sponses, increase/no change/decline, are coded as $+1/0/1$, respectively. The right subplot of Figure 1.B shows the distribution of annualized survey forecasts based on the same question. We annualize the monthly survey responses by computing a weighted yearly average $\sum_{m\in y} W_{im}[XS^e_{im}]$, where the weights are based on equation (2). The right subplot of Figure 2.B documents the number of survey responses on sales expectations (survey question D.2) per year. The number of responses is relatively constant across our sample. Towards the end of the sample it is somewhat lower. The reason is that responses are digitized only about 2 years after they have been received. At the time we obtained the data not all responses at the end of the sample had been digitized. The left subplot of Figure 2.B 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 optimistic (pessimistic) responses is higher in the first (second) half of our sample, consistent with the strong boom that ended in 2008 and the following severe depression.

**Statistics on Quantified Forecast Errors.** Figure 3.B shows 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 substantially across years, e.g. during 2009, the first year of the Greek crisis resulted in a relatively high share of major negative forecast errors.

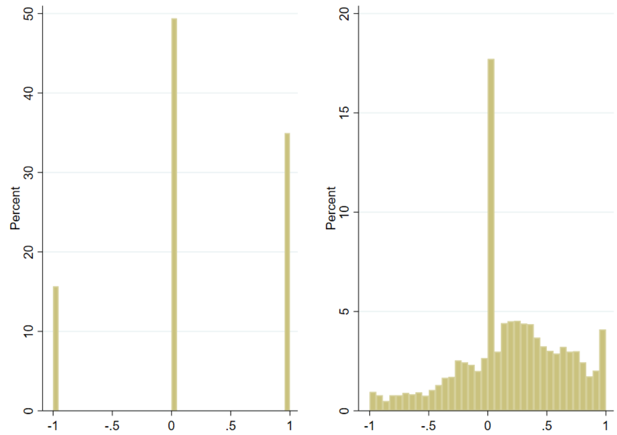


Figure 1.B: **Distribution of Sales Forecasts based on Qualitative Survey Data.** The figure on the left shows the distribution of firm-month sales forecasts based on survey question D.2. The figure on the right shows the distribution of the survey based firm-year sales forecasts when the monthly survey responses are annualized using a yearly weighted average.

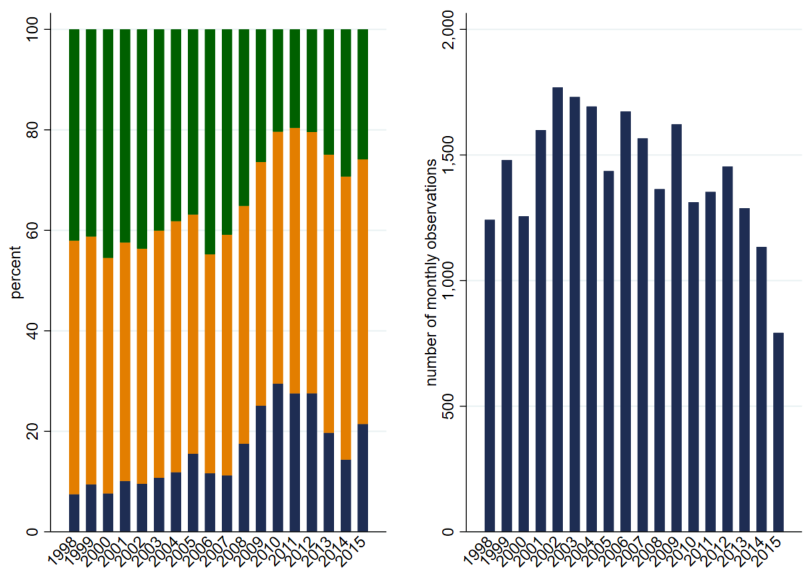


Figure 2.B: **Qualitative Survey Responses on Expected Sales Growth over Time (Survey Question D.2).** The figure on the left shows the responses indicating an increase/unchanged/decrease in green/orange/blue as share of total monthly observations per year. The figure on the right shows the total number of monthly survey responses per year distribution of firm-month sales forecasts based on survey question D.2.

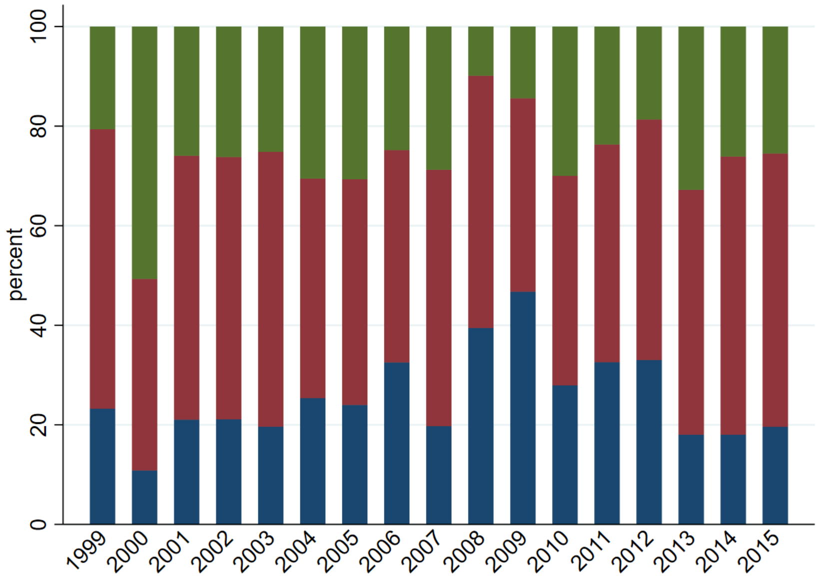


Figure 3.B: **Distribution of Quantified Sales Growth Forecast Errors across years.** Blue (green) indicates the share of major negative (major positive) forecast errors and red stands for the share of minor forecast errors. Major forecast errors are defined for the purpose of this figure as the 26% of forecast errors at the top and bottom of the distribution.

## B.7 Autocorrelation of Sales Growth

In Table 8.B we report estimates for the autocorrelation of sales growth. In the first column, to eliminate firm fixed effects, we use the biased LSDV estimator. In the other four columns, we use the Arellano and Bover (1995) Two Step Forward Orthogonal Deviations GMM (FOT). We 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) we collapse the instruments and we use the Windmeijer (2005) corrected standard errors (Roodman (2009), Caselli and Tesei (2016)). Finally, for the realizations, we use the first differences as instruments as the instruments in levels indicated serial autocorrelation in the error. Table 8.B shows that annual real sales growth from the financial statements has a negative autocorrelation, and the estimated coefficient is robust to different lag lengths. 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)). Overall, we find the result of negative autocorrelation in sales growth is very robust (also across subsamples, results are available upon request). This is consistent with evidence on other datasets in the literature, see e.g. Barrero (2019).

Table 8.B: 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), (4) and (5) 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)-(5), we 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.

## B.8 Accuracy of the Quantification Methodology: Monte Carlo Exercise

In this section, we describe how artificial data is generated and subsequently used to evaluate the precision of our methodology to quantify qualitative forecasts. We first document details of the data generating process and its calibration. Finally, we discuss results that stress the robustness of the evidence shown in Table 6.

**Generating Artificial Data.** The following outlines how we 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 3.1 to generate estimates for quantified sales growth expectations, $\hat{z}^e_{iy}$. This allows us to evaluate the accuracy between these estimates, $\hat{z}^e_{iy}$, and the actual underlying expectations, $z^e_{iy}$.

Our 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 we generate exactly matches this structure. We further take into account that the first eleven years in our sample were a boom period and the last seven years a severe bust. We start with generating a balanced panel that spans 20 years, where the first two years are used to inform lagged values. We now document how each of the three artificial variables is generated.

First, we generate artificial data for firm's sales growth, $z_{iy}$, based on an AR(1) process. We use the MA($\infty$) 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 B.7 (Table 8.B) that this process explains the data well.[51] 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, we 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 $\varpi^e_i$ 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 likely that such a process is also used by firms to form expectations.

Third, we generate the qualitative monthly expectations, $ZS^e_{im}$. These expectations need to correspond to the annual quantitative forecasts $z^e_{iy}$. For this reason, we 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

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

[51] Using an AR(2) process to generate the artificial data does not materially affect the performance of our quantification methodology. Results are discussed below and shown in Table 10.B.

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 our survey data. In particular, we 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 $z^e_{im}$ we set the corresponding $ZS^e_{im} = 0$.

Finally, for the three variables based on artificial data — $z_{iy}$, $z^e_{iy}$ and $ZS^e_{im}$ — we drop the appropriate observations so that we derive an unbalanced panel of artificial data that exactly corresponds to the structure of firm-year-month observations in our observable dataset.[52] We repeat the steps above to generate 1,000 random samples of artificial datasets. Then, for each sample, we use $z_{iy}$ and $ZS^e_{im}$ as input to our 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 we need to calibrate a number of parameters. This exercise is closely informed by our financial statements data on annual sales growth realizations and the survey data on monthly qualitative expectations. Based on the estimates reported in Table 8.B, we set the autocorrelation coefficient in the AR(1) process for artificial sales growth, $z_{iy}$, to $\theta = -0.1$. The parameters $\mu$ and $\sigma$, that govern the moments of the corresponding innovations, are calibrated to match the respective moments in our sales growth data from the financial statements. Since particularly the mean differs across the boom and bust periods in our sample, we differentiate between these episodes and set $\mu = 0.077$ ($\mu = -0.059$) and $\sigma = 0.391$ ($\sigma = 0.401$) during the boom

[52] Prior to this, we have also dropped all observations of the first two years which had only been employed to inform values of lagged variables.

(bust) period.[53] The standard deviation of the unobserved firm heterogeneity, $\sigma_{\varpi_i}$, 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, we jointly calibrate the remaining parameters that correspond to these variables to match a number of statistics in our data. We first discuss the parameters that govern the process for annual sales growth expectations. The firm specific optimism/pessimism, $\varpi_i^e$, should be related to the average firm-specific performance, $\varpi_i$. We scale $\varpi_i^e = 0.5 \cdot \varpi_i$ so that the standard deviation of the firm-specific average of the artificial monthly qualitative expectations is close to the corresponding statistic in the observable dataset (0.431 vs. 0.422). The standard deviation of the innovation, $\sigma_{\varepsilon_{iy}^e} = 0.02$, 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, we turn to the remaining parameters required to generate the monthly expectations. The standard deviation of the innovations, $\sigma_{\varepsilon_{im}^e}$, is set to 0.05, based on the within-year variation of the monthly qualitative survey responses. We 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 $\gamma$ 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 our observable dataset.[54]

All calibrated parameters and the moments we target are summarized in Table 9.B. Our calibration strategy carefully ensures close correspondence of the artificially generated data with our observable dataset. This is achieved by matching statistics that concern, amongst others, relations between qualitative survey expectations and quantitative realizations, as well as monthly and annual data. We now evaluate the appropriateness of the calibration and the assumptions on underlying

[53] Apart from the mean $\mu$, 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 we refrain from a differentiation for these parameters.

[54] In particular, we run the regression $XS_{im}^e = \beta_0 + \beta_1 x_{iy} + \phi_i + \eta_{im}$ where $\phi_i$ controls for firm fixed effects and $\eta_{im}$ is an idiosyncratic error.

Table 9.B: Calibrated parameters to generate artificial data

| Parameter | Value | Matched Moment from Financial Statements (FS) or Survey Data |
|---|---|---|
| $\mu$ | 0.077 (-0.059) | Mean in boom (bust) period of sales growth from financial statements |
| $\sigma$ | 0.391 (0.401) | Standard deviation in boom (bust) period of sales growth from financial statements |
| $\theta$ | -0.1 | Autocorrelation estimates (see Table 8.B) of sales growth from financial statements |
| $\sigma_{\varpi_i}$ | 0.129 | Standard deviation of firm-specific cross-time average of sales growth in the FS |
| $\varpi_i^e$ | $0.5\sigma_{\varpi_i}$ | Scaled to match std. dev. of firm-specific average of monthly qual. survey expectations |
| $\sigma_{\varepsilon_{iy}^e}$ | 0.02 | Std. dev. of the firm-year averages of the monthly qualitative survey expectations |
| $\sigma_{\varepsilon_{im}^e}$ | 0.05 | Mean of squared difference between monthly survey responses and their firm-year average |
| $\gamma$ | 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) |

processes by evaluating how well the artificial data conforms to statistics in the observable data that are not targeted. We 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.[55] 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.[56] Third, in the error term of the latter regression ($ZS_{iy}^e$ 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 our calibration. The second statistic particularly corroborates our calibration of $\gamma$, while the first and third statistics support our calibration of the variance of the unobserved firm heterogeneity.

**Alternative Data Generating Process.** Table 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. We now demonstrate robustness to an alternative Data Generating Process. We relax the AR(1) assumption

[55]Using the notation for the artificial variables the regression is: $ZS_{im}^e = \beta_0 + \beta_1 z_{iy} + \phi_i + \eta_{im}$. The corresponding variables in our empirical dataset have been denoted $XS_{im}^e$ on $x_{iy}$ in the main body.

[56]Using the notation for the artificial variables the regression is $ZS_{iy}^e = \beta_0' + \beta_1' z_{iy} + \phi_i' + \eta_{iy}'$. Where $ZS_{iy}^e$ is the firm-year arithmetic mean of the monthly survey responses.

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 as shown above which introduces predictability and autocorrelation in the forecast errors. Table 10.B 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.

Table 10.B: 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) |

We 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.

# C Forecast Error Predictability and Autocorrelation

This appendix includes additional results that corroborate the robustness of the results on predictability and autocorrelation of forecast errors shown in the main body and justifies our choice of baseline estimation strategy. Amongst other things, we show that our results on threshold regressions are robust to using fewer lags as recommended by Roodman (2009).[57] Our results also show that the estimates of the original Seo and Shin (2016) FD GMM are close to the biased LSDV estimates. This further justifies our choice for FOT.

[57] This is particularly important since the Hansen p-values in our GMM estimates can be on the high side. Roodman (2009) suggests that high Hansen p-values might signal biased estimates as a result of 'instrument proliferation' (or proliferation of over-identifying restrictions). We want to stress that we have taken every care to avoid instrument proliferation: our instruments are collapsed and with limited lags. Our robustness checks indicate that our results are robust to using fewer lags. We also show that the GMM estimated coefficients are much higher than the biased LSDV ones. As a result, the Hansen p-value as well as the estimates that we get are not subject to bias resulting from instrument proliferation.

## C.1 Robustness on Forecast Error Predictability

This section provides additional evidence related to the results on forecast error predictability in Section 4.1.

Table 11.C summarizes results of alternative estimations of the predictability without the threshold, equation (12). Column (1) is estimated with the Arellano and Bover (1995) FOT GMM, but without additional lags, that is without any over-identifying restrictions. The results in column (1) also have a high Hansen p-value, and the coefficient of $x_{i,y-1}$ is very close to the one in our baseline estimate (-0.158 here vs. -0.161 in Table 8). The proximity of the two coefficients clearly demonstrates that our baseline estimates are not biased as a result of instrument proliferation. Column (2) shows estimation results with the LSDV which serves as a benchmark for the dynamic panel bias. Given that the LSDV estimates suffer from the negative Dynamic Panel bias (Nickell (1981)), and our estimates using Dynamic Panel Data methods are higher than the LSDV ones, we can conclude that our baseline specification corrects this bias.

Table 11.C: 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, $x^{fe}_{iy}$ | |
| $x_{i,y-1}$ | -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 (12) 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. We 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, $\overline{x}_{IND,y}$. $x^{fe}_{iy}$ 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.

Table 12.C provides robustness on the estimation of equation (13) which includes the threshold. Columns (1) and (4) are estimated using the biased LSDV for pre-estimated threshold cut-off values. In column (1), we document a coefficient estimate of $-0.236$ for $x_{i,y-1}FEL^{q}_{i,y-1}$, compared to the

−0.146 in our baseline estimation. This indicates that we corrected the dynamic panel bias in the baseline specification. Column (2) shows that our baseline specification without over-identifying restrictions delivers estimated coefficients (−0.157 vs. −0.146 in our baseline estimation) and a threshold cut-off (28% vs. 26% in baseline) close to the corresponding figures based on our baseline setup. Our estimates in Section 4.1 are robust to having no over-identifying restrictions, which indicates that our choice of instruments is valid. In column (3), with the FD GMM, the estimated threshold cut-off is very close to the ones obtained with our modified methodology: 27% vs. 26% in our baseline estimation, and 28% when the GMM system is just identified. This implies that our threshold estimate is robust to using the original Seo and Shin (2016) estimator. However, the coefficient estimate (−0.198) is close to the biased LSDV ones (−0.235) indicating the presence of bias and justifying our choice of FOT for our baseline results.

Overall, evidence in this section corroborates our baseline result and choice of estimation methodology. The predictability and autocorrelation coefficients become non-zero following a major forecast error and this finding is robust to the lag length of instruments. With our data, the original Seo and Shin (2016) estimation with the Arellano and Bond (1991) FD delivers biased coefficient estimates which justifies our choice of Arellano and Bover (1995) FOT estimator as a baseline.

## C.2 Robustness on Forecast Error Autocorrelation

In this subsection we show that our results on the autocorrelation of sales growth forecast errors in Section 4.2 also hold using alternative estimations techniques for the threshold regression.

Table 13.C, summarizes results of alternative estimations of the predictability equation (14) without the threshold. Column (1) is estimated with the Arellano and Bover (1995) FOT GMM and only one lag in instruments (no over-identifying restrictions). The results in column (1) also have a high Hansen p-value, and the coefficient of $x^{fe}_{i,y-1}$ is very close to the one in our baseline estimate (-0.163 here vs. -0.164 in Table 9). The proximity of the two coefficients clearly demonstrates that our baseline estimates are not biased as a result of instrument proliferation. Column (2) shows estimates with the LSDV which serves as a benchmark for the dynamic panel bias. Given that the LSDV

Table 12.C: 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, $x^{fe}_{iy}$ | | | |
| $x_{i,y-1} * (1 - FEL^q_{i,y-1})$ | -0.162* | -0.0936 | -0.0941 | -0.166* |
| $x_{i,y-1} * FEL^q_{i,y-1}$ | -0.236*** | -0.157*** | -0.198** | -0.235*** |
| $FEL^q_{i,y-1}$ | -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 (13). Columns (1) and (4) are estimated using the biased LSDV for pre-estimated threshold cut-off values. For Columns (1) we used the estimated threshold from the baseline specification of Table 8; 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. We 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, $\overline{x}_{IND,y}$. $x^{fe}_{iy}$ is the forecast error of sales growth for year $y$; $x_{i,y-1}$ is the lagged realized sales growth. $FEL_{iy}$ 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.

estimates suffer from the negative Dynamic Panel bias (Nickell (1981)), and our estimates using Dynamic Panel Data methods are higher than the LSDV ones, we can conclude that our baseline specification corrects this bias.

Table 13.C: 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, $x^{fe}_{iy}$ | |
| $x^{fe}_{i,y-1}$ | -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 (14) 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. We 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, $\overline{x}_{IND,y}$. $x^{fe}_{iy}$ is the forecast error of sales growth for year $y$. *** indicate statistical significance at the 1% level, respectively.

Table 14.C provides robustness on the estimation of equation (15) which includes the threshold. Columns (1) and (4) are estimated using the biased LSDV for pre-estimated threshold cut-off values. In column (1), we document a coefficient estimate of $-0.24$ for $x^{fe}_{i,y-1}FEL^{q}_{i,y-1}$, compared to the $-0.167$ in our baseline estimation. This indicates that we corrected the dynamic panel bias in the baseline specification. Column (2) shows that our baseline specification without over-identifying restrictions delivers estimated coefficients ($-0.168$ vs. $-0.167$ in our baseline estimation) and a threshold cut-off (20% vs. 26% in baseline) not too far from the corresponding figures based on our baseline setup. Our estimates in Section 4.2 are robust to having no over-identifying restrictions, which indicates that our choice of instruments is valid. In column (3), with the FD GMM, the estimated threshold cut-off is exactly the same as the one obtained with the baseline estimation (26%). However, the coefficient estimate ($-0.217$) is close to the biased LSDV ones ($-0.240$) indicating the presence of bias and justifying our choice of FOT for our baseline results.

Overall, evidence in this section corroborates our baseline result and choice of estimation methodology. The autocorrelation coefficient becomes non-zero following a major forecast error and our finding is robust to the lag length of instruments. With our data, the original Seo and Shin (2016) estimation with the Arellano and Bond (1991) FD delivers biased coefficient estimates which justifies our choice of Arellano and Bover (1995) FOT estimator as a baseline.

## C.3 Autocorrelation and Predictability in Survey-Based Forecast Errors

This section documents the details on the probit models used to evaluate whether large survey-based monthly forecast errors are autocorrelated and predictable only in years that our quantification methodology flags as involving a major forecast error.

We test for predictability using the following probit model which is directly comparable to the continuous regressions (13)

$$\begin{aligned}\mathbb{P}\Big\{XS^{Lfe}_{im}=1\big|XS_{im},FEL_{iy}\Big\} \;&=\; G\Big(\varphi_1' XS_{im}*(1-FEL_{iy})+\varphi_2' XS_{im}*FEL_{iy}\\ &\quad +\varphi_3' FEL_{iy}+\Psi_y'+\Psi_i'\Big)\end{aligned} \tag{31}$$

Table 14.C: 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, $x_{iy}^{fe}$ | | | |
| $x_{i,y-1}^{fe} * (1 - FEL_{i,y-1}^{q})$ | -0.145 | 0.0494 | 0.304* | -0.145 |
| $x_{i,y-1}^{fe} * FEL_{i,y-1}^{q}$ | -0.240*** | -0.168*** | -0.217*** | -0.240*** |
| $FEL_{i,y-1}^{q}$ | -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 (15). Columns (1) and (4) are estimated using the biased LSDV for pre-estimated threshold cut-off values. For Columns (1) we used the estimated threshold from the baseline specification of Table 9; 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. We 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 $y$. $FEL_{iy}$ 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.

where $\Psi_i'$ and $\Psi_y'$ control for firm and year fixed effects, $FEL_{iy}$ takes the value 1 when there is a major forecast error (lower or upper 26% of the annual quantified forecast error distribution). We subtract the expectational survey responses from the corresponding realization ones, and we construct monthly forecast errors with values $XS_{im}^{fe} = \{-2, -1, 0, +1, +2\}$ (following Bachmann et al. (2013) and Massenot and Pettinicchi (2018)). $XS_{im}^{Lfe} = 1$ is a binary variable taking value one when the survey-based forecast error is $XS_{im}^{fe} = \pm 2$. $XS_{im}$ is the survey-based realization. The extrapolation bias in years with minor forecast errors is $\varphi_1'$, whereas during a major forecast error, firm's bias is $\varphi_2'$. $\varphi_3'$ indicates whether the occurrence of a major forecast error has any effect on the survey forecast errors of that year. $\mathbb{P}\{\cdot|\cdot\}$ is the conditional probability and $G()$ is the standard normal distribution which gives us the probit model.

Analogous to the predictability test we have the following dynamic probit for the persistence of the survey forecast errors, directly comparable to the continuous version, equation (15),

$$\mathbb{P}\Big\{XS^{Lfe}_{im} = \pm 2 \big| XS^{Lfe}_{i,m-3}, FEL_{iy}\Big\} = G\Big(\rho'_1 XS^{Lfe}_{i,m-3} * (1 - FEL_{iy}) + \rho'_2 XS^{Lfe}_{i,m-3} * FEL_{iy} + \rho'_3 FEL_{iy} + \Psi'_y + \Psi'_i\Big). \tag{32}$$

The persistence in years with minor forecast errors is given by $\rho'_1$, while during a major forecast error, survey forecast errors show persistence with coefficient $\rho'_2$. If only $\rho'_2$ is statistically significant for the estimated threshold 26%, then forecast errors show persistence only following a major forecast error.[58]

Table 15.C: Predictability and Persistence of firms' forecast errors of sales growth in the qualitative survey data. Probit Estimates.

| Panel A: Predictability | | Panel B: Autocorrelation | |
|---|---|---|---|
| $XS_{im} * (1 - FEL_{iy})$ | -0.0398 | $XS^{Lfe}_{i,m-3} * (1 - FEL_{iy})$ | -0.0554 |
| $XS_{im} * FEL_{iy}$ | -0.133*** | $XS^{Lfe}_{i,m-3} * FEL_{iy}$ | 0.326*** |
| $FEL_{iy}$ | 0.407* | $FEL_{iy}$ | 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. $FEL_{iy}$ 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.

[58] To proxy for the firm fixed effects we follow Wooldridge (2010). In equation (31), we use the firm specific cross-time average of the right hand side variables $XS_{im}$ and $FEL_{iy}$, and in (32), the firm-specific average of $FEL_{iy}$. To address the initial conditions problem we also include the observation of the dependent variable for each firm, $XS^{Lfe}_{i,0}$, on the right hand side (see Wooldridge (2010)).

# D Model Derivations

## D.1 Derivation of Equation (22)

To derive equation (22) for the optimal choice of attention, we begin from the original problem,

$$\max_{\lambda}\Big[\mathbb{E}U(\lambda)-C(\lambda)\Big]. \tag{33}$$

and we follow Gabaix (2014). We take the Taylor expansion of $U(\lambda)$ around the rational expectations solution, $\lambda=1$,[59]

$$U(\lambda)-U(1)=\frac{\partial U}{\partial\lambda}\bigg|_{\lambda=1}\big(\lambda-1\big)+\frac{1}{2}\frac{\partial^2 U}{\partial\lambda^2}\bigg|_{\lambda=1}\big(\lambda-1\big)^2+o(\lambda^3), \tag{34}$$

where $o(\lambda^3)=0$, because the utility is quadratic, so higher order derivatives with respect to $\lambda$ are zero. $U(\lambda)$ is given by equation (20), so that for the derivatives in equation (34) we need to calculate $\partial x^e_{y+1}(\lambda)/\partial\lambda$. Before we proceed, we introduce some useful notation. Our utility has the general form: $U(A,B)=-\frac{1}{2}(A-B)^2$. Then, we can define the trivial derivatives $U_1\triangleq\partial U/\partial A$, $U_2\triangleq\partial U/\partial B$, $U_{11}\triangleq\partial^2U/\partial A^2=-1$, $U_{22}\triangleq\partial^2U/\partial B^2=-1$ and $U_{12}\triangleq\partial^2U/\partial A\partial B=1$.

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, we 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:

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

We can now calculate the partial derivatives of the Taylor polynomial (34). Firstly, for the first

[59] Even though the utility function is quadratic, we cannot directly analytically solve equation (33), because of the presence of term $x^e_{y+1}(\lambda)$ which is unknown without knowing the choice for $\lambda$. However, with the Taylor expansion around $\lambda=1$, this term reduces to $x^e_{y+1}(1)$ which is the known rational expectations solution.

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 we 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 (33), we 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 $\mathbb{E}\epsilon_y x_y = 0$, $\forall y$, we have that $\mathbb{E}s_y^2 = \sigma_s^2 = \sigma_x^2 + \sigma_\epsilon^2$. This results in equation (22) in the main body.

## References in the Appendix


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.

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.

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

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

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

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

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.

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

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.

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

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

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

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

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

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.

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

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.