cesifo WORKING PAPERS

8148
2020

March 2020

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

Alexandros Botsis, Christoph Görtz, Plutarchos Sakellaris

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 8148*

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

Abstract

Using a novel dataset that contains qualitative firm survey data on sales forecasts as well as balance-sheet data on realized sales, we document that only major forecast errors 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 result, we develop a novel methodology to quantify qualitative survey data on firm forecasts. It is generally applicable when quantitative information, e.g. from balance sheets, is available on the realization of the forecasted variable. Finally, we provide a model of rational inattention that explains our empirical results. Firms optimally limit their degree of attention to information when operating in market environments where information processing is more costly. This results in major forecast errors that are predictable and autocorrelated.

JEL-Codes: C530, C830, D220, D840, E320.

Keywords: survey data, firm data, expectations, forecast errors, rational inattention.

Alexandros Botsis
University of Birmingham
Department of Economics
United Kingdom – B15 2TT, Birmingham
AXB1175@student.bham.ac.uk

*Christoph Görtz**
University of Birmingham
Department of Economics
United Kingdom – B15 2TT, Birmingham
c.g.gortz@bham.ac.uk

Plutarchos Sakellaris
Athens University of Economics and Business
Department of Economics
Athens / Greece
plutarch@aueb.gr

*corresponding author

February 2020
We thank Fabio Antoniou, Yiannis Dendramis, Christos Genakos, Bruce Hansen, Yiannis Karavias, Tryphon Kollintzas, Alistair Macaulay, Bartosz Mackowiak, Kaushik Mitra, Thanasis Stengos 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.

1 Introduction

Do firms make substantial errors in forecasting their future sales? Are these errors predictable and do they display autocorrelation? 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. In answering the above questions, 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 to quantify forecasts when the survey data on expectations is only available in categorical form. A requirement for this methodology is the parallel availability of quantitative data on firm realized actions from accounting or administrative sources. We provide a model of rational inattention that explains our empirical findings. Firms optimally limit their degree of attention to new information when operating in market environments where information processing is particularly costly. Limited attention results in substantial forecast errors that are autocorrelated and predictable.

Given the dynamic nature of firm activity, expectations play a major role in firms' economic behavior. Economic models that explain 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, many important questions, such as the aforementioned ones, cannot be answered with many survey-based measures for expectations, as they are typically categorical. 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. (2018), 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 18% of the distribution. Only these major forecast errors are predictable and autocorrelated. Forecast errors in the middle 64% 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 18% 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 and actual sales growth together with lower-frequency (annual) quantitative data on actual 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 the need to derive proxies for forecast errors. This requires some identifying assumptions that allow us to derive two nonlinear equations. The first one relates observed quantitative annual sales to observable variables and the second one relates unobserved quantitative expected sales 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

[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.

to derive fitted values for quantitative expectations on sales. We would like to stress that this methodology can be applied to any firm variable as long as there is a quantitative counterpart to a high-frequency categorical survey variable. We compare our methodology to two alternative ones for obtaining quantified predictions based on ordered logit and probit. We find that our method performs best, in the sense that it is highly consistent with the direction indicated by survey-based annualized qualitative forecasts.

To interpret our empirical results, we use a framework of rational inattention to study how agents make optimal forecasts when attention to information is costly.[2] In particular, we study the behavior of forecast errors when firms' sales growth follows an auto-regressive process and firms endogenously choose the degree of inattention to information. We show that information cost spikes lead to major forecast errors that are predictable and negatively autocorrelated. In the absence of these cost spikes, firms make fully informed and rational forecasts with minor forecast errors.

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 requirement 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

[2]The seminal work here is Sims (2003). Mackowiak et al. (2018) provide a recent survey of the literature on rational inattention.

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 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. 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 also 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 paper is organized as follows. Section 2 discusses the data. Section 3 lays out our methodology to quantify firms' forecasts. It further provides an overview of the characteristics of the estimated forecasts and the resulting forecast errors. Section 4 discusses our empirical results on the predictability and autocorrelation of forecast errors. In Section 5 introduces a model with rational inattention that can explain these empirical results. 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. 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 the 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 services — and sends out surveys that include somewhat different questions across these sectors. We

[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.

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 aggregate sales, has increased/remained unchanged/decreased.*

> Question D.2: *During the next 3 months, you expect your aggregate sales to increase/remain unchanged/decrease.*

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 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.5 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	4,089	163	75	285
Real Sales (in 2005 Euros)	5,825	6,710,000,000	29,100,000	7,201,930	179,000,000
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 shows how we can use the qualitative information from the survey and the quantitative data from the 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 subsequently discusses the quality of our estimate for sales growth forecasts and provides an overview about the implied quantified forecast errors.

3.1 Quantifying Expected Sales Growth

Theoretically, we can decompose the expected annual growth rate of sales for firm i in year y, x^e_{iy}, into its monthly components. 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.[6] The aggregation of monthly growth rates to an annual frequency can be formalized as the following weighted average

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

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 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.[7] The second component of the weights are dummy variables which 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 any of the variables about quantitative expectations on 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

[6]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.

[7]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.

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. We also allow for this linear correlation to be asymmetric, as in Smith and McAleer (1995), depending on whether the monthly variable is positive or negative. This is the first identifying assumption (ID1) we make to quantify firms' forecast errors. It can be formalized as

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

where the error terms, ν^+_{im} and ν^-_{im}, are assumed to be normally and independently distributed across equations, time and firms. We further allow for the coefficients α, β, γ_1 and γ_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 their corresponding monthly components. In fact, this linear relationship in equations (3) can be used to algebraically eliminate the unobserved variables $x^{e,+}_{im}$ and $x^{e,-}_{im}$ from equation (1). Combining equations (1) and (3) yields (detailed derivations are shown in Appendix B.1)

$$x^e_{iy} = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \qquad \text{with} \qquad \xi_{iy} = \frac{\sum_{m\in y} W^+_{im}\nu^+_{im} + \sum_{m\in y} W^-_{im}\nu^-_{im}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}, \tag{4}$$

where we defined

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

to ease the notation. P_{iy} (N_{iy}) denotes the weighted share of months per year that record a rise (decline) in expected sales. 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 about annual sales growth, x^e_{iy}, 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^e_{iy}. We turn to exactly this next.

We know that for each firm i, realized sales growth in year y is the sum of expected sales growth in that year and a forecast error, $x_{iy} = x^e_{iy} + x^{fe}_{iy}$, where x^{fe}_{iy} denotes the forecast error. Using this expression to replace x^e_{iy} in equation (4) yields after rearranging

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

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. This means, we need to find a proxy variable that is correlated with the unobserved x^{fe}_{iy} to remove its effect from the error. For the rest of this subsection, we construct such a proxy and deal with unobserved firm heterogeneity to obtain an expression of equation (6) that can be estimated.

Forecast Error Proxy. We assumed in ID1 — consistent with Smith and McAleer (1995) — that the quantitative monthly sales forecasts are correlated with their annual counterparts. The same argument applies for the qualitative observed survey responses. Next, we assume that also the qualitative monthly survey forecasts, XS^e_{im}, are correlated with the corresponding quantitative annual forecasts, x^e_{iy}; and analogously that the qualitative monthly survey realizations, XS_{im}, are correlated with the corresponding quantitative annual realizations, x_{iy}.[8] This leads to our second identifying

[8]For the realized values, we show the assumed positive correlation is actually present in our data in Table 3.A of Appendix A.4. A simple intuition supports the correlation between the monthly forecasts and the annual predictions:

assumption (ID2): the forecast error in qualitative monthly survey data is linearly correlated with the forecast error in annual quantitative data. This means that we can use the survey data to find a proxy variable that is correlated with the unobserved forecast error, x^{fe}_{iy}, as the following auxiliary regression holds

$$x^{fe}_{iy} = x_{iy} - x^{e}_{iy} \quad = \quad \delta_0\big(XS_{im} - XS^{e}_{im}\big) + \nu^{fe}_{iy}, \qquad \text{[ID2]} \tag{7}$$

where $\delta_0 > 0$ is the correlation coefficient, and ν^{fe}_{iy} the approximation error. Multiplying both sides of this equation with the monthly weights W_{im} and summing up all the monthly observations within a year gives[9]

$$x^{fe}_{iy} = \delta_0 XS^{fe}_{iy} + \nu^{fe}_{iy}, \qquad \text{where} \qquad XS^{fe}_{iy} \triangleq \sum_{m\in y} W_{im}[XS_{im}] - \sum_{m\in y} W_{im}[XS^{e}_{im}]. \tag{8}$$

We can now use the observed weighted monthly quantitative survey forecasts as a proxy for the unobserved (omitted) variable x^{fe}_{iy} in equation (6) and obtain the following equation for which all variables are observable

$$x_{iy} \quad = \quad \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \delta_0 XS^{fe}_{iy} + \nu^{fe}_{iy} + \xi_{iy}. \tag{9}$$

This equation can be estimated with Nonlinear Least Squares (NLS). The resulting parameter estimates and the error term can then be used to compute fitted values for the quantitative expected sales growth rates, x^{e}_{iy} via equation (4).

Unobserved Firm Heterogeneity. The last remaining issue with the NLS estimation of equation (9) is to control for unobserved firm heterogeneity. To estimate the coefficients of equation (9), the error term $\nu^{fe}_{iy} + \xi_{iy}$ needs to be mean-independent of all the right-hand side vari-

even though information can be updated every month, a part of it should coincide with the information set used for the annual prediction. As a result, monthly survey forecasts will be correlated with the one for the entire year.

[9] Note that the term $\sum_{m\in y} W_{im}$ in front of x^{fe}_{iy} and ν^{fe}_{iy} is then equal to unity.

ables, i.e. $\mathbb{E}[\nu_{iy}^{fe} + \xi_{iy}|P_{iy}, N_{iy}, XS_{iy}^{fe}] = 0$. A sufficient condition for this equality to hold is that $\mathbb{E}[\nu_{iy}^{fe}|\{XS_{im}^e, XS_{im}\}_{m=1,\dots}] = 0$ and $\mathbb{E}[\xi_{iy}|\{XS_{im}^e, XS_{im}\}_{m=1,\dots}] = 0$.[10] We will focus on these two expressions in turn.

Firstly, we will focus on the error term, ξ_{iy}. From equation (4), we know that this error term is a non-linear function of XS_{im}^e, ν_{im}^+ and ν_{im}^-. Without loss of generality, we can focus on either of the two errors, ν_{im}^+ and ν_{im}^-, and drop the superscripts $+/-$ for the moment to ease the exposition. Our third identifying assumption (ID3a) is to assume that

$$\nu_{im} = \psi_i + \tilde{\nu}_{im}, \qquad \text{and} \qquad \mathbb{E}[\tilde{\nu}_{im}|\{XS_{im}^e, XS_{im}\}_{m=1,\dots}] = 0, \qquad \text{[ID3a]}$$

where $\tilde{\nu}_{im}$ is an idiosyncratic error which is mean-independent of XS_{im} and XS_{im}^e, for all m. The unobserved firm heterogeneity, ψ_i, is in fact an omitted variable hidden in the error term ν_{im} and is endogenous. The reason is that the firm heterogeneity is related to the past, current and future values of XS_{im} and XS_{im}^e, so that $\mathbb{E}[\psi_i + \tilde{\nu}_{im}|\{XS_{im}^e, XS_{im}\}_{m=1,\dots}] = \mathbb{E}[\psi_i|\{XS_{im}^e, XS_{im}\}_{m=1,\dots}] \neq 0$ from assumption IDE3a. To control for the unobserved heterogeneity, we need to approximate $\mathbb{E}[\psi_i|\{XS_{im}^e, XS_{im}\}_{m=1,\dots}]$.[11] However, the unobserved heterogeneity parameter ψ_i appears in equation (3) which has purely expectational terms and hence is independent of the realization variable XS_{im}. Therefore, we only need to approximate $\mathbb{E}[\psi_i|\{XS_{im}^e\}_{m=1,\dots}]$.

A widely used approximation for this purpose is the one suggested in Mundlak (1978).[12] The original Mundlak (1978) specification is linear, but we also include higher order terms due to the non-

[10] In Appendix B.2 we provide a proof of this statement.

[11] The structure of the non-linear equation (4) that we want to estimate does not allow us to derive an estimator for ψ_i analytically, and we cannot use dummy variables either, because the cross-sectional dimension is very large. Another possibility would be to linearise the equation 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 linearised 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.

[12] See e.g. Bartelsman et al. (1994), Semykina and Wooldridge (2010) and Kosova (2010)). The Mundlak (1978) approximation is the standard tool used in non-linear models in panel data and in linear ones it is equivalent to the least squares dummy variable and the standard within estimator.

linearity of our equation. Our fourth identifying assumption (ID4a) is to assume that the conditional expectation of the unobserved firm heterogeneity in the error term ξ_{iy} is

$$\mathbb{E}[\psi_i | \{XS^e_{im}\}_{m=1,\ldots}] \;\; = \;\; \delta' \overline{XS^e_i} + \delta'' (\overline{XS^e_i})^2, \qquad \text{[ID4a]}$$

which results in the following auxiliary regression for ψ_i

$$\psi_i \;\; = \;\; \delta' \overline{XS^e_i} + \delta'' (\overline{XS^e_i})^2 + \omega_i, \tag{10}$$

where ω_i is the approximation error for which $\mathbb{E}[\omega_i | \{XS^e_{im}\}_{m=1,\ldots}] = \mathbb{E}[\omega_i | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}] = 0$, and $\overline{XS^e_i}$ is the simple arithmetic mean of the survey variable XS^e_{im} across time for each firm i.[13] We can now substitute equation (10) for ψ_i in $\nu_{im} = \psi_i + \tilde{\nu}_{im}$, and if we distinguish again between ν^+_{im} and ν^-_{im}, this yields

$$\nu^-_{im} = \delta_1 \overline{XS^e_i} + \delta_2 (\overline{XS^e_i})^2 + \omega^-_i + \tilde{\nu}^-_{im}, \qquad \text{and} \qquad \nu^+_{im} = \delta_3 \overline{XS^e_i} + \delta_4 (\overline{XS^e_i})^2 + \omega^+_i + \tilde{\nu}^+_{im}, \tag{11}$$

where δ_1, δ_2, δ_3 and δ_4 are coefficients.

Next we turn to the requirement on the forecast error proxy residual, $\mathbb{E}[\nu^{fe}_{iy} | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}] = 0$. Analogous to the case above, we can assume that $\nu^{fe}_{iy} = \tilde{\nu}^{fe}_{iy} + \psi^{fe}_i$, where only the unobserved firm heterogeneity ψ^{fe}_i is endogenous and $\tilde{\nu}^{fe}_{iy}$ is the idiosyncratic error which is mean-independent of XS_{im} and XS^e_{im}, for all m. This is our identifying assumption ID3b. So, we need to approximate $\mathbb{E}[\psi^{fe}_i | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}]$. As with ID4a, we use the Mundlak (1978) approximation. We assume as our identifying assumption ID4b that the conditional expectation of the unobserved firm heterogeneity in the forecast error residual ν^{fe}_{iy} is

$$\mathbb{E}[\psi^{fe}_i | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}] \;\; = \;\; \delta_5 \overline{XS^{fe}_i}. \qquad \text{[ID4b]}$$

[13] The fact that the right-hand-side variables of equation (10) are continuous functions of XS^e_{im} and that $\mathbb{E}[\omega_i | \{XS^e_{im}\}_{m=1,\ldots}] = 0$ imply that $\mathbb{E}[\omega_i | \{XS^e_{im}, (XS^e_{im})^2\}_{m=1,\ldots}] = 0$. The fact that the unobserved heterogeneity comes from the expectational equations implies that $\mathbb{E}[\omega_i | \{XS^e_{im}\}_{m=1,\ldots}] = \mathbb{E}[\omega_i | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}]$.

This results in the following auxiliary regression for ψ_i^{fe}

$$\psi_i^{fe} \;=\; \delta_5 \overline{XS_i^{fe}} + \omega_i^{fe} \tag{12}$$

where δ_5 is a coefficient and $\overline{XS_i^{fe}} = \overline{XS_i} - \overline{XS_i^e}$ is an average of the overall firm mis-prediction behavior. ω_i^{fe} is the residual which is mean-independent of XS_{im}, XS_{im}^e.[14] Since the forecast error approximation (equation (8)) is linear, we abstain from including higher order terms in the approximation to economize on variables in the non-linear estimation.

The Final Equation to be Estimated. As we have provided a way to approximate the unobserved firm heterogeneity, we can now derive the final estimable equation. We substitute equations (11) and (12) into the error terms of (9) and obtain

$$\begin{aligned} x_{iy} \;=\; & \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1 \overline{XS_i^e} P_{iy} + \delta_2 (\overline{XS_i^e})^2 P_{iy} + \delta_3 \overline{XS_i^e} N_{iy} + \delta_4 (\overline{XS_i^e})^2 N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} \\ & + \delta_0 XS_{iy}^{fe} + \delta_5 \overline{XS_i^{fe}} + \tilde{\xi}_{iy}, \end{aligned} \tag{13}$$

where

$$\tilde{\xi}_{iy} = \tilde{\nu}_{iy}^{fe} + \omega_i^{fe} + \frac{\sum_{m \in y} W_{im}^{+} (\tilde{\nu}_{im}^{+} + \omega_i^{+}) + \sum_{m \in y} W_{im}^{-} (\tilde{\nu}_{im}^{-} + \omega_i^{-})}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \tag{14}$$

Overall we obtained equation (13) that is estimable because the error term $\tilde{\xi}_{iy}$ is mean-independent of the explanatory variables.[15] This addresses the issue of the unobserved heterogeneity in equation (9), so that we can obtain consistent and unbiased estimates of the coefficients of interest, α, β, γ_1 and γ_2, of the non-linear expectation equation.[16]

Summary of the Quantification Method. Given the identifying assumptions ID1-ID4, we

[14] As previously, mean-independence of ω_i^{fe} from XS_{im} and XS_{im}^e implies also it's mean independence from $\overline{XS_i^{fe}}$.

[15] In Appendix Section B.2 we provide a detailed proof.

[16] The error term, $\tilde{\xi}_{iy}$, in equation (13) 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 both problems — this robust estimator treats errors as clustered within cross-sectional units.

derived two nonlinear equations. Equation (13), relates observed quantitative annual sales to observable variables, and equation (9) relates unobserved quantitative expected sales to observable variables. The key is that both of these relationships depend on the same parameters. We estimate the 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.

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

1. Compute the annualized forecast error proxy from the survey data using equation (8) with weights based on equation (2).

2. Compute the 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).

3. Compute the firm heterogeneity proxies $\overline{XS_i} = \frac{1}{T_i}\sum_m XS_{im}$ and $\overline{XS_i^e} = \frac{1}{T_i}\sum_m XS_{im}^e$, as well as $\overline{XS_i^{fe}} = \overline{XS_i} - \overline{XS_i^e}$ from the monthly qualitative survey data.

4. Estimate equation (13), using NLS including all the fixed effects proxies on the right hand side, as well as the forecast error proxy. Run the estimation separately for the boom ($y \leq 2008$) and bust period ($y > 2008$).

5. Use the NLS estimated coefficients of equation (13) to compute the fitted values for quantified sales growth forecast, $\hat{x}_{iy}^e$, from equation (4). Neither the forecast error proxy nor the fixed effects proxies should be included.

Our parameter estimates of the NLS estimation of equation (13) 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 obviously extends beyond sales growth. It is applicable to any qualitative (survey based) variable on future developments, as long as a qualitative and a quantitative corresponding variable on the subsequent realization is available.

3.2 Quantified Forecasts and Forecast Errors

The subsection above outlines how we can derive annual quantitative forecasts for sales growth, and subsequently forecast errors, using monthly qualitative survey data and annual quantitative data from the financial statements. This subsection first provides an overview of the estimated sales forecasts and then turns to the characteristics of the resulting forecast errors.

Quantified Forecasts. We can use the observable survey data on expected sales developments to benchmark how well our quantified forecasts match the direction of the corresponding survey expectations. 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 survey expectations as well as their annualized counterparts are reported in Appendix B.5. 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 quantified forecasts against the annualized survey forecasts, we split responses in each variable into three categories — positive, zero or negative — and cross-tabulate the three directions. Panel A.1 in Table 2 reports how well our quantified forecasts match the direction of the annualized observable. The main diagonal shows the share of observations that exhibit directionally consistent movements 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 96.89% 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 limited directional guidance of the annualized 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 2 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 24% 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.[17] Overall, these exercises show that our forecasts based on NLS are highly consistent with the direction of the qualitative survey responses.

Table 2: Directional coincidence 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	24.97%	0.00%	0.41%	11.21%	0.00%	0.00%
Zero Forecasts	0.26%	14.71%	0.34%	0.00%	56.96%	0.00%
Positive Forecasts	2.09%	0.00%	57.21%	0.00%	0.00%	31.83%
	Directional Consistency: 96.89%			Directional Consistency: 100.00%		
	Panel B.1: Ordered Logit			**Panel B.2: Ordered Logit**		
	Negative	Zero	Positive	Negative	Zero	Positive
Negative Forecasts	15.26%	0.00%	10.85%	9.30%	0.00%	2.59%
Zero Forecasts	4.30%	0.00%	11.15%	15.68%	0.00%	41.62%
Positive Forecasts	11.27%	0.00%	47.18%	3.03%	0.00%	27.78%
	Directional Consistency: 62.44%			Directional Consistency: 37.08%		
	Panel C.1: Ordered Probit			**Panel C.2: Ordered Probit**		
	Negative	Zero	Positive	Negative	Zero	Positive
Negative Forecasts	15.29%	0.00%	10.82%	9.30%	0.00%	2.59%
Zero Forecasts	4.30%	0.00%	11.15%	15.68%	0.00%	41.62%
Positive Forecasts	11.24%	0.00%	47.20%	3.03%	0.00%	27.78%
	Directional Consistency: 62.49%			Directional Consistency: 37.08%		

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), Ordered Logit (Panel B) and Ordered Probit (Panel C). The restricted sample only considers annualized survey observations for which, in a given year, all underlying monthly observations report forecasts in the same direction.

[17] Results are even fully directionally consistent if we consider annualized observations for which 51% or more of all underlying monthly survey responses of a particular year indicated sales forecasts in the same direction. This comprises 53% of the observations of the full sample used in Panel A.1.

Even though these results speak for the quality of our forecast estimates, there are potentially alternative ways to quantify sales growth forecasts — most notably ordered response models such as logit and probit — which we want to discuss briefly. We outline the details of these alternatives in Appendix B.4. Panels B.1 and C.1 in Table 2 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 62% for both ordered logit and probit. For the restricted sample shown in Panels B.2 and C.2, these shares even drop to 37%, pointing to substantial directional differences between forecasts 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. Our discussion suggests, our sales growth forecasts based on NLS massively outperform alternatives based on ordered response models in terms of alignment with the actual survey data. We next turn to the forecast errors computed using the NLS based sales growth forecasts.

Quantified Forecast Errors. Figure 1 shows the distribution for our forecast errors and we report moments on this distribution in Table 3. The average forecast error in our sample is negative but close to zero (-0.05) and slightly larger than the median (-0.07). This implies the average forecast on sales growth is five percentage points more optimistic than the subsequent average realization. Interestingly, Table 3 shows that the distribution of forecast errors is very stable across the boom and the severe depression in our sample. During both periods their shares are close to the 18% of the full sample which is imposed by construction.[18] Overall, a number of forecast errors made by

[18]Appendix B.5 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

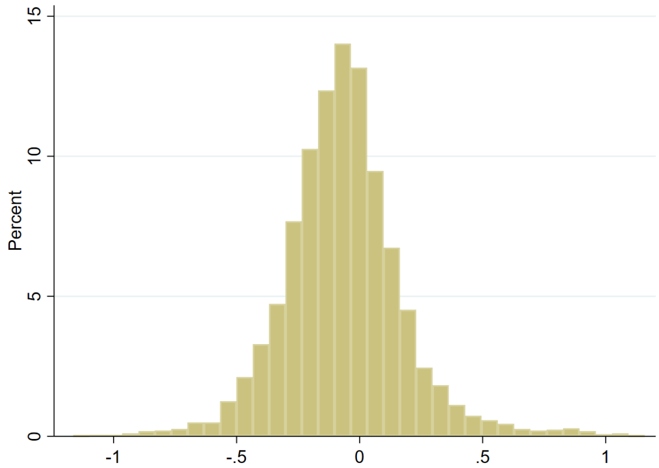

Figure 1: **Distribution of Quantified Sales Growth Forecast Errors.** The 1% of forecast errors at the top of the distribution are omitted to ease visibility.

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 18% of forecast errors to be major, which is in line with the estimates for this threshold obtained in the next section. At the bottom (top) 18 percentiles, firms expected sales growth to be 25.6 (12.0) percentage points higher (lower) than subsequently realized. Hence, also a large number of the remaining 64% of forecast errors in the center of the distribution, which we call minor, are still economically significant.

Table 3: Descriptive Statistics for Quantified Sales Growth Forecast Errors.

	Mean	Median	Stand. Dev.	Share of Forecast Errors (in %) Major Negative	Minor	Major Positive
Full Sample	-0.05	-0.07	0.34	18	64	18
Boom	-0.05	-0.07	0.34	17	65	18
Bust	-0.05	-0.07	0.35	19	63	18

Major forecast errors are defined for the purpose of this table as the 18% of forecast errors at the top and bottom of the distribution. The boom (bust) period spans the years 1998-2008 (2009-2015).

How are these major and minor forecast errors distributed over different dimensions of our sample?

bust subperiods. Further, the shares of major positive and negative forecast errors are also very similar across 2-digit sectors — results are available upon request.

Panel A of Table 4 sorts the sample according to the share of major forecast errors in a firm's total number of observations. For a large number of firms (378 of the total of 799 firms) this share is positive up to 60%, 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 (2,663). 126 (196) firms make exclusively major (minor) forecast errors, however these account only for 199 (481) firm year observations in our sample and are hence quite short lived. Panel B of Table 4 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 28% and 40% across the firm size distribution where the larger firms make slightly fewer major forecast errors.

Table 4: 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%	196	0	481		
(0%,30%]	154	239	1274	(0%,30%]	36.52%
(30%,60%]	224	588	1389	(30%,60%]	39.58%
(60%,90%]	85	367	525	(60%,90%]	34.31%
(90%,100%]	126	199	199	(90%,100%]	27.59%

Major forecast errors are defined for the purpose of this table as the upper or lower 18% of forecast error distribution. The percentile of total net assets has been determined using firm's average percentile in the pool distribution.

Table 5: 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 $y-1$	23.67%	53.36%	22.97%	100%
Minor FE in year $y-1$	14.93%	69.40%	15.67%	100%
Positive MaFE in year $y-1$	21.49%	58.72%	19.79%	100%

Major positive (negative) forecast errors are defined for the purpose of this table as the upper (lower) 18% of forecast error distribution.

The evidence in Table 4 shows that, independent of their size, most firms make major as well as minor forecast errors. Table 5 provides the average year-on-year transition matrix between minor, positive and negative major forecast errors for the pooled data. It suggests that firms don't tend to make many consecutive major forecast errors, but that major and minor forecast errors are likely to alternate. Following either a negative or a positive major forecast error in year $y-1$, the probability of making another major forecast error of either type in year y is always lower than the probability of making a minor forecast error. Further, 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 observations of individual firms. 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 very 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 analysis in the next section on the full sample.

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 forms a forecast using an estimated autocorrelation coefficient, ρ^e, such that $z_t^e = \rho^e z_{t-1}$, and the forecast error is $z_t - z_t^e = (\rho - \rho^e) z_{t-1}$. If $\rho - \rho^e = 0$, then firms correctly extrapolate without any bias and forecast errors are not predicable from past realizations and are purely random. If $\rho - \rho^e \neq 0$, this is a violation of the *efficiency property* (Pesaran (1987)) of the full information rational expectations (FIRE) hypothesis. For $(\rho - \rho^e) < (>)0$, the firm estimates the autocorrelation of the variable with a positive (negative) bias.

We estimate the extrapolation bias, $\varphi \triangleq \rho - \rho^e$, using the following equation

$$x_{iy}^{fe} = \varphi x_{i,y-1} + \Psi_y + \Psi_i + \eta_{iy}, \tag{15}$$

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

$$x_{iy}^{fe} = \varphi_1 x_{i,y-1} * (1 - FEL_{i,y-1}^q) + \varphi_2 x_{i,y-1} * FEL_{i,y-1}^q + \varphi_3 FEL_{i,y-1}^q + \Psi_y + \Psi_i + \eta_{iy}, \tag{16}$$

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 φ_1, whereas following a major forecast error, firm's bias is φ_2. Given the estimated cut-off $q\%$, if $\varphi_1 = 0$ and $\varphi_2 \neq 0$, then forecast errors are predictable only following major forecast errors. φ_3 indicates whether the occurrence of a major forecast error has any effect on the forecast error in the following period.

We estimate equation (16) using a slightly adapted version of 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 (16) for all the values of $q\%$ in the pre-determined interval $q\% = 5\%, 6\%, 7\%...30\%$ to obtain the value of the objective function of the estimator.[19] 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)).

[19] 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 50% of all observations.

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), for all values of $q\%$ we find the one that minimizes the objective function which then determines the final estimates for φ_1, φ_2 and φ_3, as well. For our baseline estimation of equation (16) 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)).[20] 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)).[21] Finally, we also estimate the linear equation (15) with the Arellano and Bover (1995) Forward Orthogonal Transformations GMM.[22]

In Table 6, we document the results from the estimation of equations (15) and (16). Column (1) reports results from the simple linear equation and documents a highly significant negative estimate for the coefficient on $x_{i,y-1}$.[23] This result is in line with findings in the literature which documents that

[20]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.

[21]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.

[22]We collapse the instruments and use Windmeijer (2005) corrected standard errors. We use only four lags (system is just identified) because the Hansen criterion indicated weak instruments for more lags. We document this in Appendix C.1.

[23]In all estimations in Section 4 the two-step estimated covariance matrix is singular and therefore the generalized inverse is used to construct robust standard errors. This does not affect the consistency of our estimated coefficients, but only indicates that the estimated standard errors are not the ideally efficient (Roodman (2009)). Although this could signal problems of over-identification or of instrument proliferation, we have have avoided this by limiting the lag length and collapsing the instruments. We further discuss the matter in Appendix C.1. There, we also show that

firm's forecast errors are predictable by past realizations (see e.g. Gennaioli et al. (2016), Massenot and Pettinicchi (2018) and Bordalo et al. (2018)). This violation of the efficiency property of the FIRE hypothesis implies firms use a biased estimate of the autocorrelation coefficient to forecast sales growth. Column (2) reports estimation results for equation (16). The threshold for major forecast errors is estimated to include those observations at the top and bottom $q = 18\%$ of the distribution. These forecast errors are substantial and economically significant. A forecast error at 18% (82%) of the distribution implies that sales growth was expected to be 25.6 (12.0) percentage points higher (lower) than the subsequent realization.[24] Importantly, the coefficient of the lagged realization is now only statistically significant when it is interacted with $FEL^q_{i,y-1}$. Our results in column (2) of Table 6 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 different lag length and to the one step estimator. We also document why the original Seo and Shin (2016) FD GMM would be unsuitable

all of our estimates are extremely similar to using the one step estimator.

[24]We attempted to estimate the threshold equation with distinct thresholds for upper and lower cut-off %. However, the estimation is unreliable for a number of reasons (results can be made available upon request). Firstly, none of the estimated coefficients were significant, indicating low efficiency. Secondly, none of the cut-off values nor the coefficient estimates were 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 variables. In the extreme case of uncorrelated instruments the asymptotic normality and variance of the estimated coefficients do not hold, i.e. inference is not valid. Stock et al. (2002) indicate that sensitivity of estimates to different lag lengths can be a symptom of weak identification. We stress that our baseline estimates with a symmetric threshold do not deviate much from ones based on specifications with different lag lengths.

with our data.

Table 6: Predictability of firms' forecast errors of sales growth.

	(1)	(2)
Estimation	FOT	FOT
Stand. Errors	Windmeijer corrected	Windmeijer corrected
Lags as Instruments	2-5	2-6
Estimated Threshold q	N.A.	18%
	Dependent Variable: Sales Forecast Error, x^{fe}_{iy}	
$x_{i,y-1}$	-0.150***	–
$x_{i,y-1} * (1 - FEL^q_{i,y-1})$	–	-0.0803
$x_{i,y-1} * FEL^q_{i,y-1}$	–	-0.177**
$FEL^q_{i,y-1}$	–	0.0192
Observations	2,805	2,069
# of Firms	590	432
Over-identified	No	Yes
Hansen p-value	N.A.	0.393
m2 test p-value	0.534	0.570

Column (1) shows estimates of equation (15) 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 (16). Instruments are collapsed in both specifications. In column (1), the instruments are with lags dated from $y-2$ to $y-5$ (more lags yielded Hansen p value < 0.1). In column (2) 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. y dummies are included in all estimations, but are omitted to maintain a simple representation. 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 $q = 18\%$ of its empirical pool distribution. ***, ** and * indicates statistical significance at the 1%, 5% and 10% 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 that again any findings crucially depend on the size of forecast errors.

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

$$s^{fe}_{iy} = \rho x^{fe}_{i,y-1} + \Psi_y + \Psi_i + \eta_{iy}, \tag{17}$$

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

$$x_{iy}^{fe} = \rho_1 x_{i,y-1}^{fe} * (1 - FEL_{i,y-1}^{q}) + \rho_2 x_{i,y-1}^{fe} * FEL_{i,y-1}^{q} + \rho_3 FEL_{i,y-1}^{q} + \Psi_y + \Psi_i + \eta_{iy}, \tag{18}$$

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 ρ_1, while following a major forecast error, forecast errors are autocorrelated with coefficient ρ_2. If only ρ_2 is statistically significant for the estimated threshold $q\%$, then forecast errors show persistence only following a major forecast error.

We estimate equations (17) and (18) using the exact same estimators and specifications as for the corresponding equations (15) and (16) on forecast error persistence. Table 7 shows the estimation results for former two equations. Column (1) reports that based on the simple linear equation, forecast errors are negatively autocorrelated.[25] 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 equation (18) 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 $x_{i,y-1}^{fe} * FEL_{i,y-1}^{q}$ is highly significant and suggests a negative autocorrelation of forecast errors. In contrast, the estimate on the coefficient of $x_{i,y-1}^{fe} * (1 - FEL_{i,y-1}^{q})$ is not significantly different from zero so that minor forecast errors are not autocorrelated. It is reassuring that

[25] In the literature (see e.g. Gennaioli et al. (2016) and Tanaka et al. (2019)) unobserved firm heterogeneity is usually dealt with using the within estimator. However, the within estimator is negatively biased when the time dimension of the panel is finite (Nickell (1981)). Instead, we use the unbiased, consistent and efficient Arellano and Bover (1995) Forward Orthogonal Transformations (FOT) GMM.

the cut-off $q = 17\%$ for the threshold is remarkably close to the one, independently estimated, for the equation on forecast error predictability (18%). For the non-linear threshold model both 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 to the use of different lag length and the one step estimator. We also document why the original Seo and Shin (2016) FD GMM would be unsuitable for our dataset.

Overall, we have documented in this section that with respect to autocorrelation and predictability, only for major forecast errors firms violate the FIRE hypothesis. For smaller (absolute) 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 forecasts. 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.

Table 7: Autocorrelation of firms' forecast errors on sales growth.

	(1)	(2)
Estimation	FOT	FOT
Stand. Errors	Windmeijer corrected	Windmeijer corrected
Lags as Instruments	2-5	2-6
Estimated Threshold q	N.A.	17%
	Dependent Variable: Sales Growth Forecast Error, x^{fe}_{iy}	
$x^{fe}_{i,y-1}$	-0.190***	–
$x^{fe}_{i,y-1} * (1 - FEL^q_{i,y-1})$	–	-0.0118
$x^{fe}_{i,y-1} * FEL^q_{i,y-1}$	–	-0.201**
$FEL^q_{i,y-1}$	–	0.0106
Observations	2,069	2,069
# of Firms	432	432
Over-identified	No	Yes
Hansen p-value	N.A.	0.479
m2 test p-value	0.688	0.725

Column (1) shows estimates from equation (17) 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 (18). In both specifications instruments are collapsed. In column (1), the instruments are with lags dated from $y-2$ to $y-5$ (Hansen p with more lags indicated invalid instruments). In column (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. y dummies are included in all estimations, but are omitted to maintain a simple representation. 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 17% of its empirical pool distribution. ***, ** and * indicates statistical significance at the 1%, 5% and 10% level, respectively.

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 $\mathbb{E}\epsilon_y = 0$, $\mathbb{E}\epsilon_y^2 = \sigma^2_\epsilon$ 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,

$$x_y = \rho_0 + \rho x_{y-1} + u_y, \tag{19}$$

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

At time y the firm wants to obtain a one period ahead forecast, $\tilde{x}_{y+1}$, that minimizes the expected squared forecast error, but its information set only includes the noisy signal s_y and not the true value x_y.[26] Then the optimal forecast, x^e_{y+1}, is[27]

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

The first order condition yields $x^e_{y+1} = \mathbb{E}\big[x_{y+1}|s_y\big]$ and using the fact that x_{y+1} follows the AR(1) process (19), we obtain

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

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

$$x^e_{y+1} = \rho\mathbb{E}\big[x_y|s_y\big] = \rho\lambda_0 + \lambda\rho s_y, \qquad \text{where} \qquad \lambda = \frac{Cov(x_y, s_y)}{V(s_y)}, \quad \text{and} \quad \lambda_0 = (1-\lambda)\mu = 0. \tag{20}$$

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 ϵ_y, it follows that

$$\lambda = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\epsilon^2}. \tag{21}$$

Equation (21) shows that in presence of noise, $\sigma_\epsilon^2 > 0$, λ is strictly between 0 and 1. This has implications on the optimal forecast (20), which is, when applying the definition for the signal,

$$x^e_{y+1} = \lambda\rho x_y + \lambda\rho\epsilon_y. \tag{22}$$

[26]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.

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

Equation (22) 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 firm's optimal forecast underestimates the persistence, ρ, of sales growth, since $0 < \lambda < 1$. Under perfect information (in the absence of noise $\sigma_\epsilon^2 = 0$), $\lambda = 1$ and equation (22) 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 λ; and technically, in the framework above this is equivalent to varying the noise variance σ_ϵ^2. If the firm pays attention to all information the noise variance equals zero and $\lambda = 1$. For a positive noise variance, attention to information is limited and $0 < \lambda < 1$. 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, λ.

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 wants to make an optimal forecast given the utility

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

where the parameter λ determines the firm's degree of attention to the signal about current sales growth.[28] The full information optimal forecast ($\lambda = 1$) would be $x_{y+1}^e = \rho s_y$ which is in line with

[28]As implicitly assumed 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

the underlying AR(1) process (19) 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 λ. If we substitute the optimal forecast, $x^e_{y+1}(\lambda)$, into the utility function, we obtain the indirect utility function

$$U(\lambda) = W\left(x^e_{y+1}(\lambda), \lambda s_y\right), \tag{23}$$

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

$$C(\lambda, c_y) = c_y K(\lambda), \tag{24}$$

where $K(\lambda)$ is a continuous increasing and convex function in λ. Note that this function depends on the cost shock c_y, which is assumed to be independently and identically distributed across time and is bounded between zero and a positive upper bound.[29] The firm observes the cost shock at the beginning of the period prior to any choice on the level of attention.

Given the above assumptions, the firm first chooses an optimal level of attention λ^*, and condi-

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 makes our model much more tractable. It is similar to an assumption in Mackowiak and Wiederholt (2009) who assume that past realizations of the state variable are never observed.

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

tional on this choice, it obtains in a second step the optimal forecast for sales growth.[30] 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 (23) and the cost function (24). 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, λ^* by solving the following intratemporal problem

$$\lambda_y^* \triangleq \arg\max_{\lambda}\Big[-\frac{1}{2}\sigma_s^2(1-\lambda)^2 - c_y K(\lambda)\Big], \tag{25}$$

where σ_s^2 denotes the variance of the signal. It becomes apparent now that, given the time varying cost c_y, also the optimal level of attention fluctuates over time. The first order condition is then[31]

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

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 $K(\cdot)$.[32] However, to briefly provide intuition about how the optimal level of inattention depends on the information cost and the variance of the signal, we specify $K(\lambda) = \lambda^a$ where $a \geq 1$. Then the first order condition has, for the cases $a = 1$

[30]The reason why we can equivalently have a two-step approach is that in the first step the decision is independent of sales growth, x_y.

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

[32]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).

and $a = 2$, the following simple analytical solutions

$$\begin{aligned} \lambda_y^* &= \frac{\sigma_s^2 - c_y}{\sigma_s^2}, \qquad \text{for} \quad a = 1. \\ \lambda_y^* &= \frac{\sigma_s^2}{\sigma_s^2 + 2c_y}, \qquad \text{for} \quad a = 2. \end{aligned} \tag{26}$$

These two parameterizations exemplify that the optimal level of attention is positively related to the variance of the signal on sales growth, σ_s, and negatively related to the cost shock, c_y. In other words the firm acquires more information if the variance of the signal is high, and it reduces the level of attention in light of an increase in the cost of information acquisition. In general, for $c_y = 0$ so there is no cost for information, $\lambda_y^* = 1$. 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 (25), the firm's optimal forecast is given by

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

so that the optimal forecast is

$$x_{y+1}^e = \lambda_y^* \rho s_y. \tag{27}$$

As in the simple signal extraction problem above, the firm underestimates 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. In this section, 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 (19) and the optimal forecast (27), the ex-post forecast error in the framework with rational inattention is given by

$$x_{y+1}^{fe} \triangleq x_{y+1} - x_{y+1}^{e} = (1 - \lambda_y^*)\rho x_y - \lambda_y^* \rho \epsilon_y + u_{y+1}, \tag{28}$$

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 will result in larger absolute forecast errors and a violation of the full information rational expectations hypothesis.

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

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

which is purely random. Note that the noise, ϵ_y, is zero for $\lambda^* = 1$ as implied by equation (21). A positive cost, $c_y > 0$, will reduce λ_y^* to positive values strictly lower than unity. In this case the absolute forecast error is

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

Since $\left|(1-\lambda^*)\rho x_y - \lambda^* \epsilon_y\right|$ will typically be larger than zero, this absolute forecast error for $0 < \lambda^* < 1$ will be larger than the one for the case $\lambda^* = 1$.[33] In presence of positive cost, firms subsequently underestimate the persistence of sales growth by $100 \cdot \lambda_y^*\%$. The forecast error's dependence on underestimated 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

[33] 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.

differences between major and minor forecast errors. In this section, we document that the estimated coefficient on past sales growth in equation (16) — which corresponds to $(1-\lambda_y^*)\rho$ in model equation (28) — is significantly different from zero for large absolute forecast errors. Hence, for these major forecast errors it implies $0 < \lambda_y^* < 1$.[34] For minor forecast errors though, we find the coefficient estimate is not statistically significant anymore which implies that λ_y^* is (close to) unity.

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 (28) and the discussion of Result 1. As explained above, following an increase in the cost c_y from zero to a positive value, the attention parameter λ_y^* reduces from unity to positive values strictly lower than one. For $\lambda_y^* = 1$, the forecast error as given in equation (28) 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 (28) is negative, which only is the case if $\rho < 0$. In Appendix B.6 we provide empirical evidence from our dataset that the autocorrelation of sales growth is indeed negative. Taken together, Results 1 and 2 are also consistent with the empirical findings based on equation (16) 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.

[34] It will become clear in the discussion of Result 2 why the coefficient estimate for large forecast errors is negative in equation (16).

Substituting $x_y^e + x_y^{fe}$ for x_y in equation (28) 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 (18) 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 showed 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, λ_y^*, that depends on the cost for information governed by c_y. The literature on rational inattention often remains agnostic about the specific drivers of the cost for information in such models. We use our dataset to 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 don't 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.[35] These could for example be changes in the specific market environment, or 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. 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

[35]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 Acharya and Wee (2020) or Macaulay (2019).

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*, pages 1–17.

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

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

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

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

Bachmann, R. and Zorn, P. (2013). What Drives Aggregate Investment? Evidence from German Survey Data. Technical report.

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

Botsis, A., Görtz, C., and Sakellaris, P. (2020). 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.

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. Technical report, 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. (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.

Stock, J. H., Wright, J. H., and Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*, 20(4):518–529.

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*, pages 105–120.

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

Supplementary Appendix

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.[36] 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

[36] 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.[37,38] 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

[37]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.

[38]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.[39] 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,

[39]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.[40] 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.

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

$$D3_{im} = \beta_0 + \boldsymbol{\beta}\Big[INV_{im}, SS^e_{im}, L^e_{im}, U_{im}\Big]' + \psi_i + \psi_y + \eta_{im}, \tag{29}$$

where the vector $\boldsymbol{\beta} = [\beta_1, \beta_2, \beta_3, \beta_4]$, ψ_i and ψ_y control for firm and year fixed effects respectively,

[40] 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.

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.

and η_{im} is the idiosyncratic error. The variables $D3_{im}$, INV_{im}, SS^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.[41]

[41] 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'*. SS^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;

We estimate equation (29) twice: first, by eliminating ψ_i using standard fixed effects tools and second, by substituting NACE sector dummies for ψ_i. In Panel A of Table 2.A we report the results from estimating equation (29). 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 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:

$$QS_{im} = \beta_0 + \beta_1[U_{i,m-1} - U_{i,m-3}] + \psi_i + \psi_y + \eta_{im}, \tag{30}$$

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, ψ_i and ψ_y control for firm and year fixed effects respectively, and η_{im} is the idiosyncratic error.[42] As previously, we estimate equation (30) in two ways: firstly, we eliminate ψ_i using standard fixed effects tools and secondly, we substitute NACE sector dummies for ψ_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.

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 firm i in the income statements, x_{iy}, should be positively correlated with the survey question A.2

+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.

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

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***
SS^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.

concerning the evolution of current sales, XS_{im}.

We examine this by estimating the following linear equations

$$XS_{im} = \beta_0 + \beta_1 x_{iy} + \psi_i + \psi_y + \eta_{im}, \tag{31}$$

where ψ_i and ψ_y control for firm and year fixed effects respectively and η_{im} is the idiosyncratic error. As previously, we estimate equation (31) 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.

Overall, based on the results in Tables 2.A and 3.A, we find that survey responses are consistent,

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 ***.

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 substitute equation (3) into (1)

$$x_{iy}^e = \sum_{m \in y} W_{im}^+ \left[\alpha + \gamma_1 x_{iy}^e + \nu_{im}^+ \right] + \sum_{m \in y} W_{im}^- \left[-\beta + \gamma_2 x_{iy}^e + \nu_{im}^- \right].$$

Then, using the definition for W^+_{im} and W^-_{im}, we get

$$\begin{aligned} x^e_{iy} &= \big[\alpha + \gamma_1 x^e_{iy}\big] \sum_{m\in y} W_{im} \mathbb{1}_{[XS^e_{im}=1]} + \sum_{m\in y} W_{im} \mathbb{1}_{[XS^e_{im}=1]} \nu^+_{im} \\ &+ \big[-\beta + \gamma_2 x^e_{iy}\big] \sum_{m\in y} W_{im} \mathbb{1}_{[XS^e_{im}=-1]} + \sum_{m\in y} W_{im} \mathbb{1}_{[XS^e_{im}=-1]} \nu^-_{im}. \end{aligned} \tag{32}$$

To simplify the notation we define

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

where P_{iy} (N_{iy}) denotes the weighted share of months within a year indicating that indicate a rise (fall) in expected sales. Rearranging to solve equation (32) for x^e_{iy} yields

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

which is equation (4) in Section 3.1.

B.2 Proof Related to the Estimation of Equation (9)

To estimate the coefficients of equation (9), the error term $\nu^{fe}_{iy} + \xi_{iy}$ needs to be mean-independent of all the right-hand side variables, i.e. $\mathbb{E}[\nu^{fe}_{iy} + \xi_{iy} | P_{iy}, N_{iy}, XS^{fe}_{iy}] = 0$ — this is 'Assumption NLS.1' in Wooldridge (2010). The underlying mathematical form of this (and any) conditional expectation is $\mathbb{E}[\nu^{fe}_{iy} + \xi_{iy} | \sigma(P_{iy}, N_{iy}, XS^{fe}_{iy})]$, where $\sigma(P_{iy}, N_{iy}, XS^{fe}_{iy})$ is the minimal sigma-algebra generated by P_{iy}, N_{iy} and XS^{fe}_{iy}. Given that all P_{iy}, N_{iy} and XS^{fe}_{iy} are continuous (and hence Borel) functions of XS^e_{im} and XS_{im}, then from the Doob-Dynkin Lemma (see Proposition 3 in Rao and Swift (2006)) we know that $\sigma(P_{iy}, N_{iy}, XS^{fe}_{iy}) \subset \sigma(XS^e_{im}, XS_{im})$. As a result, from the standard properties of the conditional expectations we have that $\mathbb{E}[\mathbb{E}[\nu^{fe}_{iy} + \xi_{iy} | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}] | P_{iy}, N_{iy}, XS^{fe}_{iy}] = \mathbb{E}[\nu^{fe}_{iy} + \xi_{iy} | P_{iy}, N_{iy}, XS^{fe}_{iy}]$. Therefore, if $\mathbb{E}[\nu^{fe}_{iy} | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}] = 0$ and $\mathbb{E}[\xi_{iy} | \{XS^e_{im}, XS_{im}\}_{m=1,\ldots}] = 0$, then $\mathbb{E}[\nu^{fe}_{iy} + \xi_{iy} | P_{iy}, N_{iy}, XS^{fe}_{iy}] = 0$.

We provided a way to approximate the unobserved firm heterogeneity, and we derived the final estimable equation (13). For equation (13), by the same principles as above, it suffices to prove that $\mathbb{E}[\tilde{\xi}_{iy}|\{XS^e_{im}XS_{im}\}_{m=1,...}] = 0$. Then, $\tilde{\xi}_{iy}$ is also mean-independent of all the right hand side variables of equation (13) which are all continuous functions of XS^e_{im} and XS_{im}. This means that Wooldridge (2010)'s 'Assumption NLS.1' is satisfied. Indeed, from equation (14)

$$
\begin{aligned}
\mathbb{E}\big[\tilde{\xi}_{iy}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] &= \mathbb{E}\big[\tilde{\nu}^{fe}_{iy}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] + \mathbb{E}\big[\omega^{fe}_{i}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] \\
&+ \frac{1}{1-\gamma_1 P_{iy} - \gamma_2 N_{iy}} \sum_{m\in y} W^{+}_{im} \mathbb{E}\big[(\tilde{\nu}^{+}_{im} + \omega^{+}_{i})|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] \\
&+ \frac{1}{1-\gamma_1 P_{iy} - \gamma_2 N_{iy}} \sum_{m\in y} W^{-}_{im} \mathbb{E}\big[(\tilde{\nu}^{-}_{im} + \omega^{-}_{i})|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] \\
&= 0,
\end{aligned}
$$

where the terms W^{+}_{im}, W^{-}_{im}, P_{iy} and N_{iy} 'go outside' the conditional expectation as they are continuous functions of XS^e_{im} and XS_{im} and therefore $\sigma(XS^e_{im}, XS_{im})$-measurable. This follows from the Doob-Dynkin Lemma and the standard properties of the conditional expectations. From ID3a we have that $\mathbb{E}\big[\tilde{\nu}^{+}_{im}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] = \mathbb{E}\big[\tilde{\nu}^{-}_{im}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] = 0$; and from ID3b that $\mathbb{E}\big[\tilde{\nu}^{fe}_{iy}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] = 0$. From ID4a we have that $\mathbb{E}\big[\omega^{+}_{i}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] = 0$ and $\mathbb{E}\big[\omega^{-}_{i}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] = 0$; and from ID4b that $\mathbb{E}\big[\omega^{fe}_{i}|\{XS^e_{im}, XS_{im}\}_{m=1,...}\big] = 0$. This completes the proof.

B.3 Nonlinear Least Squares Estimation

Table 4.B reports the results of the NLS estimation of equation (13). Column (1) shows estimation results for the boom period up to 2008 and column (2) for the following recession. The signs of the coefficients are as expected by economic intuition, although the denominator coefficients γ_1 and γ_2 are only statistically significant for the bust period post 2008.

Table 4.B: NLS Estimation of Equation (13)

	(1)	(2)
Coefficients	Dependent Variable: x_{iy}	
α	0.321***	0.133***
β	0.316***	0.153***
γ_1	-0.0338	0.464**
γ_2	-1.397	0.467**
δ_0	0.168***	0.149***
Firm-Year Observations	2,471	1,397
R^2	0.073	0.084
Period	$y \leq 2008$	$y > 2008$

Fixed effects of equation (13) 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.

B.4 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 $XS^{e,*}_{im}$ which defines the outcome of the observed survey response, XS^{e}_{im}, as follows

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

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

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

with $a_1, a_2 \in \mathbb{R}$ being the unobserved threshold parameters. Now assume that $XS^{e,*}_{im}$ is linearly determined by a vector of explanatory variables, $XS^{e,*}_{im} = \delta X^{XS}_{im} + \psi_i + e_{im}$, with ψ_i being the unobserved firm heterogeneity and e_{im} the idiosyncratic error term. The assumed distribution of e_{im} determines whether the model will be probit (standard normal) or logit (logistic). The explanatory variables X^{XS}_{im} can be from both the survey and the financial statements. We can eliminate the unobserved heterogeneity ψ_i using the Mundlak (1978) approximation, that is the cross-time firm-specific averages of all the panel dependent variables $\psi_i \approx \frac{1}{T_i}\sum_m^{T_i} X^{XS}_{im}$, 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, $\hat{XS}^{e,*}_{im}$. These will be the quantified values of the survey variable, $\hat{XS}^{e}_{im}$ That is $\hat{XS}^{e}_{im} = \hat{XS}^{e,*}_{im} = \hat{\delta} X^{XS}_{im}$. The estimated $\hat{XS}^{e}_{im}$, are the quantified value of the firm's monthly response conditional on X^{XS}_{im}.[43] Finally, we can derive their annualized quantified values using the weighted average $\hat{x}^{e}_{iy} \triangleq \sum_{m \in y} W_{im}[\hat{XS}^{e}_{im}]$, using the weights given in equation (2).

Table 5.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^{XS}_{im}, are (i) $XS^{e}_{m} = (N_m)^{-1} \sum_i XS^{e}_{im}$, 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.[44]

B.5 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 responses, 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

[43] 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 $\frac{1}{T_i} \sum_{m}^{T_i} X^{XS}_{im}$ 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 $\frac{1}{T_i} \sum_{m}^{T_i} X^{XS}_{im}$ 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.

[44] 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 5.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}	
XS^e_m	1.631***	1.565***	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. XS^e_{im} is 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.

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.

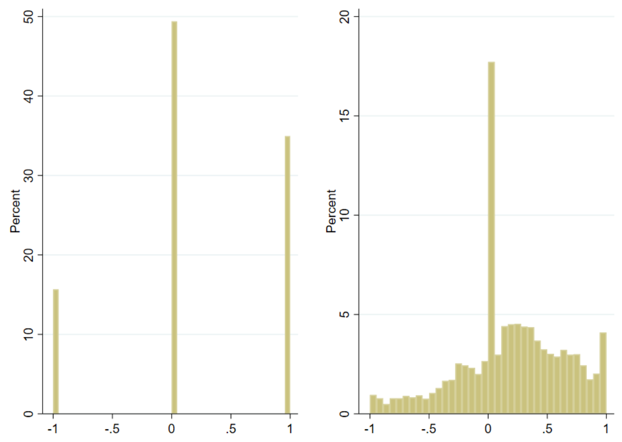

Figure 1.B: **Distribution of Sales Forecasts based on Qualitative Survey Data.** The figure on the left shows the distribution of firm-month sales forecasts based on survey question D.2. The figure on the right shows the distribution of the survey based firm-year sales forecasts when the monthly survey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

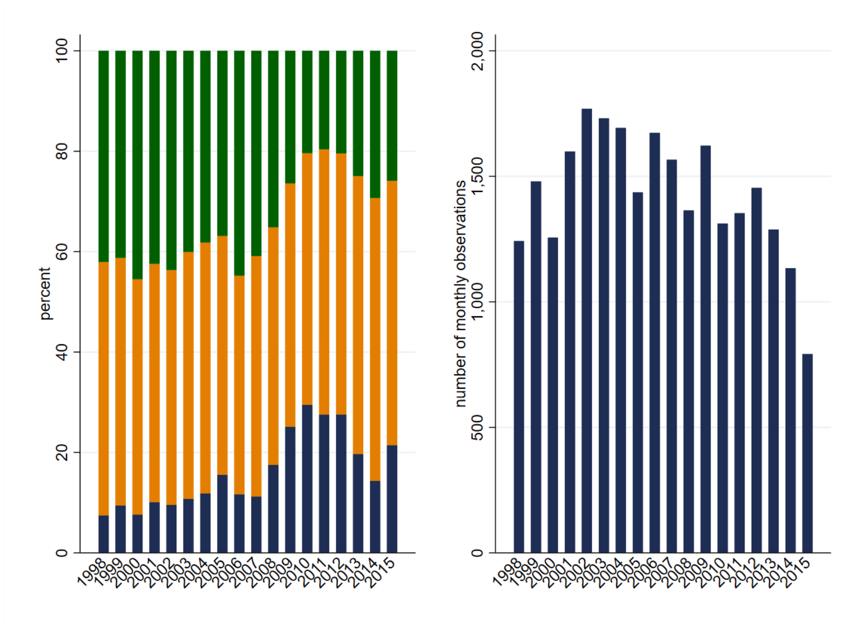

Figure 2.B: **Qualitative Survey Responses on Expected Sales Growth over Time (Survey Question D.2).** The figure on the left shows the responses indicating an increase/unchanged/decrease in green/orange/blue as share of total monthly observations per year. The figure on the right shows the total number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

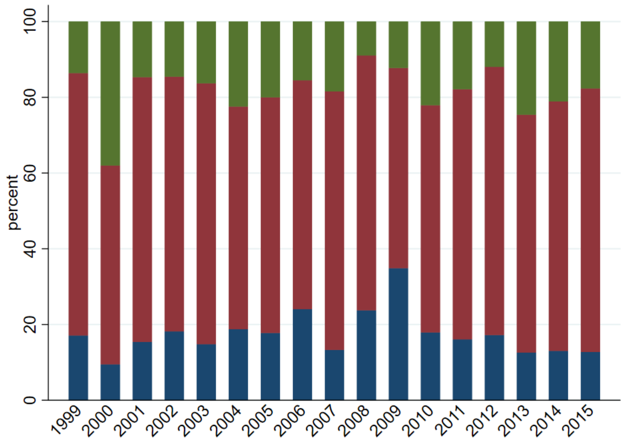

Figure 3.B: **Distribution of Quantified Sales Growth Forecast Errors across years.** Blue (green) indicates the share of major negative (major positive) forecast errors and red stands for the share of minor forecast errors. Major forecast errors are defined for the purpose of this figure as the 18% of forecast errors at the top and bottom of the distribution.

B.6 Autocorrelation of Sales Growth

In Table 6.B 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 used 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 collapsed the instruments and we used the Windmeijer (2005) corrected standard errors (Roodman (2009), Caselli and Tesei (2016)). Table 6.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)).

Table 6.B: Autocorrelation of firms' realized sales growth

	(1)	(2)	(3)	(4)
Estimation	LSDV		FOT	
Stand. Errors	Robust		Windmeijer corrected	
Lags as Instruments	N.A.	2-11	2-6	2-4
	Dependent Variable: Sales Growth Forecast Error, 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 values 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.

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) and to using the less efficient one step estimator. Our results also show that the estimates of the original Seo and Shin (2016) FD GMM are sensitive to the lags of instruments and close to the biased LSDV estimates. This further justifies our choice for FOT.

Before getting into the details of our robustness checks we want to discuss the singularity of the two-step weighting matrix. Generally, this could signal problems of excessive over-identification coming from instrument proliferation. This could lead to biased estimates and Hansen p-values. We demonstrate throughout our exposition that with our choice of instruments, we have taken all the necessary steps and have avoided excessive over-identification. Our instruments do not use the entire lag length and are collapsed in all estimations. In all of our estimations we report the number of lags we use as instruments. We also report the Hansen p-values which are all below 0.55 — these are hence not upwards biased resulting in values towards unity which would be indicative of excess over-identification due to a large number of instruments. Finally, our results are not sensitive to the lags used as instruments which further demonstrates that over-identification is not the reason why the covariance matrix is singular.

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 7.C summarizes results of alternative estimations of the predictability without the threshold, equation (15). Column (1) is estimated with the Arellano and Bover (1995) FOT GMM, but with more lags than our baseline estimate. The results in column (1) have Hansen p-value < 0.05 indicating weak instruments, which justifies our choice of fewer lags in the baseline estimation of the equation without the thresholds. 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 8.C provides robustness on the estimation of equation (16) which includes the threshold. Columns (1), (6) and (7) are estimated using the biased LSDV for pre-estimated threshold cut-off values. In column (1), we document a coefficient estimate of -0.271 for $x_{i,y-1}FEL^q_{i,y-1}$, compared to the -0.177 in our baseline estimation. This indicates that we corrected the dynamic panel bias in the baseline specification. Column (2) shows that using the one step estimator for the standard errors does not affect our estimates: the estimated threshold is the same as our baseline, 18%, and the estimated coefficient is -0.185 vs. -0.177 in our baseline estimation. Column (3) shows that our baseline specification with fewer lags as instruments delivers estimated coefficients (-0.18 vs.

Table 7.C: Predictability of firms' forecast errors of sales growth – Robustness Checks for the Specification without Threshold.

	(1)	(2)
Estimation	FOT	LSDV
Stand. Errors	Windmeijer (2005) corrected	Robust
Lags as Instruments	2-6	N.A.
	Dependent Variable: Sales Forecast Error, x_{iy}^{fe}	
$x_{i,y-1}$	-0.148***	-0.238***
Constant	–	0.121***
Observations	2,805	3,559
# of Firms	590	754
Over-identified	Yes	N/A
Hansen p-value	0.014	N.A.
m2 test p-value	0.527	N.A.

Table shows alternative estimations of equation (15) 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 lags dated from $y-2$ to $y-6$ collapsed. The two-step covariance matrix is singular, so the generalized inverse is used instead. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. y dummies are included in all estimations, but are omitted to maintain a simple representation. x_{iy}^{fe} is the forecast error of sales growth for year y; $x_{i,y-1}$ is the lagged realized sales growth. ***, ** and * indicates statistical significance at the 1%, 5% and 10% level, respectively.

-0.177 in our baseline estimation) and a threshold cut-off (16% vs. 18% in baseline) close to the corresponding figures based on our baseline setup. Our estimates in Section 4.1 are robust to using fewer lags and the less efficient one step estimator. In columns (4) and (5), with the FD GMM, the estimated threshold cut-off is very sensitive to the lag length of instruments: 18% with lags from $y-2$ to $y-6$, and only 6% with lags $y-2$ and $y-3$. In addition, the coefficient estimates (-0.23) are close to the biased LSDV ones (-0.27). These two facts justify our choice of FOT for our baseline results.

Overall, evidence in this section corroborates our baseline result and choice of estimation methodology. The extrapolation coefficient becomes non-zero following a major forecast error which 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 results sensitive to the lag length of the instruments, which justifies our choice of Arellano and Bover (1995) FOT estimator as a baseline.

Table 8.C: Predictability of firms' sales growth forecast errors – Robustness Checks for the Threshold Specifications.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Estimation	LSDV	FOT	FOT	FD	FD	LSDV	LSDV
Stand. Errors	Robust	1-Step, Robust	Windmeijer (2005) corrected			Robust	
Lags as Instruments^P	N.A.	2-6	2-3	2-6	2-3	N.A.	N.A.
Estimated Threshold q	^{P}18%	18%	16%	18%	6%	^{P}18%	^{P}6%
	Dependent Variable: Gross Sales Forecast Error, x_{iy}^{fe}						
$x_{i,y-1} * (1 - FEL_{i,y-1}^q)$	-0.112*	-0.0817	-0.0860	-0.0951	-0.0886	-0.112*	-0.186***
$x_{i,y-1} * FEL_{i,y-1}^q$	-0.271***	-0.185***	-0.180***	-0.231**	-0.234***	-0.271***	-0.272***
$FEL_{i,y-1}^q$	0.0171	0.0251	0.0304	0.0254	-0.0550	0.0171	-0.000298
Constant	0.100***	–	–	–	–	0.100***	0.105***
Observations	2,643	2,069	2,069	1,915	1,915	2,643	2,643
# of Firms	574	432	432	423	423	574	574
Over-identified	N.A.	Yes	No	Yes	No	N.A.	N.A.
Hansen p-value	N.A.	0.393	N.A.	0.392	N.A.	N.A.	N.A.
m2 tes pt-value	N.A.	0.566	0.546	0.831	0.663	N.A.	N.A.

Instruments in all specifications are collapsed; ^{P} indicates pre-estimated threshold cut-off value. The table shows alternative estimations of equation (16). Columns (1), (6) and (7) 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 6; for column (6) the threshold is estimated in column (4); for (7) the threshold is estimated in column (5). Columns (2) and (3) are the adapted Dynamic Panel Threshold estimator using the Arellano and Bover (1995) FOT GMM: (2) is with the one-step estimator; (3) with fewer lags as instruments than the baseline specification. Columns (4) and (5) are with original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$ and to $y-3$. In (3), (4) and (5) the two-step covariance matrix is singular, so the generalized inverse is used instead. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. y dummies are included in all estimations, but are omitted to maintain a simple representation. x_{iy}^{fe} 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 * indicates statistical significance at the 1%, 5% and 10% level, respectively.

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 9.C, summarizes results of alternative estimations of the predictability equation (17) without the threshold. Column (1) is estimated with the Arellano and Bover (1995) FOT GMM and five lags in instruments. The results in column (1) have Hansen p-value < 0.05 indicating weak instruments, which justifies our choice of fewer lags in the baseline estimation of the equation without the thresholds. Column (2) shows estimates 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 9.C: Autocorrelation of firms' forecast errors of sales growth – Robustness the Specification without Threshold.

	(1)	(2)
Estimation	FOT	LSDV
Stand. Errors	Windmeijer (2005) corrected	Robust
Lags as Instruments	2-6	N.A.
	Dependent Variable: Sales Forecast Error, x^{fe}_{iy}	
$x^{fe}_{i,y-1}$	-0.191***	-0.262***
Constant	–	0.0795**
Observations	2,069	2,643
# of Firms	432	574
Over-identified	Yes	N.A.
Hansen p-value	0.0489	N.A.
m2 test p-value	0.712	–

Table shows alternative estimations of equation (17) 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 dated from $y-2$ to $y-6$ (Hansen criterion indicates invalid instruments). The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. y dummies are included in all estimations, but are omitted to maintain a simple representation. x^{fe}_{iy} is the forecast error of sales growth for year y. ***, ** and * indicates statistical significance at the 1%, 5% and 10% level, respectively.

Table 10.C provides robustness on the estimation of equation (18) which includes the threshold. Columns (1), (6) and (7) are estimated using the biased LSDV for pre-estimated threshold cut-off values. In (1) we observe the estimated coefficient to be at -0.273 compared to the -0.201 in our baseline estimation. This indicates that we corrected the dynamic panel bias. From column (2), we observe that the one-step estimates are very close to our baseline figures: the threshold is at 18% (17% in baseline) and the coefficient is -0.195 (-0.201 in baseline). From column (3), we see that our baseline estimator with fewer lags as instruments also delivers estimated coefficients, -0.202 (-0.201 in our baseline estimation), and a threshold cut-off, 16% (17% in baseline), that are close to the ones in our baseline results. Therefore, our estimates are very robust to using fewer lags or the one-step estimator. In columns (4) and (5), with the FD GMM, estimated threshold cut-off is very sensitive to the lag length of instruments: it is 24% with lags from $y-2$ to $y-6$ and 31% with lags $y-2$ and $y-3$. In addition, the coefficient estimates are much lower -0.23 and closer to the

biased LSDV ones, −0.27. These two facts justify our choice of FOT for our baseline results. Note also that the sample size drops substantially when using the FD GMM due to the first-differencing.

Overall, this section shows our baseline result that the forecast errors are autocorrelated following a major forecast error is robust to the lag length and to using the one-step estimator. Moreover, we showed that with our data the original Seo and Shin (2016) estimation with the Arellano and Bond (1991) delivers biased estimates sensitive to the lag length of the instruments, fully justifying our choice of Arellano and Bover (1995) FOT estimator in our baseline results.

Table 10.C: Autocorrelation of firms' forecast errors on sales growth – Robustness for the Threshold Estimation.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Estimation	LSDV	FOT	FOT	FD	FD	LSDV	LSDV
Stand. Errors	Robust	1-Step, Robust	Windmeijer (2005) corrected			Robust	
Lags as Instruments	N.A.	2-6	2-3	2-6	2-3	N.A.	N.A.
Estimated Threshold q	^{P}17%	18%	16%	24%	31%	^{P}24%	^{P}31%
	Dependent Variable: Sales Forecast Error, x_{iy}^{fe}						
$x_{i,y-1}^{fe} * (1 - FEL_{i,y-1}^{q})$	-0.0952	-0.0315	-0.0469	0.0583	-0.0254	-0.0783	-0.216
$x_{i,y-1}^{fe} * FEL_{i,y-1}^{q}$	-0.273***	-0.195***	-0.202***	-0.225**	-0.228***	-0.266***	-0.263***
$FEL_{i,y-1}^{q}$	0.00147	0.0139	0.0165	-0.00745	-0.0138	-0.0141	-0.000537
Constant	0.0894***	–	–	–	–	0.0915***	0.0812**
Observations	2,643	2,069	2,069	1,915	1,915	2,643	2,643
# of Firms	574	432	432	423	423	574	574
Over-identified	N.A.	Yes	No	Yes	No	N.A.	N.A.
Hansen p-value	N.A.	0.470	N.A.	0.547	N.A.	N.A.	N.A.
m2 test p-value	N.A.	0.669	0.710	0.911	0.952	N.A.	N.A.

Instruments in all specifications are collapsed; ^{P} indicates pre-estimated threshold cut-off value. The table shows alternative estimations of equation (18). Columns (1), (6) and (7) 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 7; for column (6) the threshold is estimated in column (4); for (7) the threshold is estimated in column (5). Columns (2) and (3) are the adapted Dynamic Panel Threshold estimator using the Arellano and Bover (1995) FOT GMM: (2) is with the one-step estimator; (3) is with fewer lags as instruments than the baseline specification. Columns (4) and (5) are with original Seo and Shin (2016) with the Arellano and Bond (1991) First-Difference GMM (FD), with lags dated from $y-2$ to $y-6$ and to $y-3$. In (3), (4) and (5) the two-step covariance matrix is singular, so the generalized inverse is used instead. The Arellano-Bond p-value (m2 test) shows no serial correlation of order two in the errors. y dummies are included in all estimations, but are omitted to maintain a simple representation. x_{iy}^{fe} 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 * indicates statistical significance at the 1%, 5% and 10% level, respectively.

D Model Derivations

D.1 Derivation of Equation (25)

To derive equation (25) 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, $\lambda = 1$,[45]

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

where $o(\lambda^3) = 0$, because the utility is quadratic, so higher order derivatives with respect to λ are zero. $U(\lambda)$ is given by equation (23), 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 $U_1 \triangleq \partial U/\partial A$, $U_2 \triangleq \partial U/\partial B$, $U_{11} \triangleq \partial^2 U/\partial A^2 = -1$, $U_{22} \triangleq \partial^2 U/\partial B^2 = -1$ and $U_{12} \triangleq \partial^2 U/\partial A \partial B = 1$.

Recall that $x^e_{y+1}(\lambda) \triangleq x^e_{y+1}\big(\lambda s_y\big) = \arg\max_{x_{y+1}} U(x_{y+1}, \lambda s_y)$. The first order condition implies $U_1\big(x^e_{y+1}(\lambda), \lambda s_y\big) = 0$. Therefore, 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:

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

[45] 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 λ. However, with the Taylor expansion around $\lambda = 1$, this term reduces to $x^e_{y+1}(1)$ which is the known rational expectations solution.

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

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

because $U_1 = 0$ at the optimum (recall that 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 (25) 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.

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. (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.