

Quantifying qualitative survey data with panel data

Alexandros Botsis, Christoph Görtz, Plutarchos Sakellaris

Angaben zur Veröffentlichung / Publication details:

Botsis, Alexandros, Christoph Görtz, and Plutarchos Sakellaris. 2024. "Quantifying qualitative survey data with panel data." *Journal of Economic Dynamics and Control* 167: 104929. <https://doi.org/10.1016/j.jedc.2024.104929>.

Nutzungsbedingungen / Terms of use:

CC BY 4.0

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

CC-BY 4.0: Creative Commons: Namensnennung

Weitere Informationen finden Sie unter: / For more information see:

<https://creativecommons.org/licenses/by/4.0/deed.de>



Quantifying qualitative survey data with panel data [☆]Alexandros Botsis ^a, Christoph Görtz ^{b,c,*}, Plutarchos Sakellaris ^d^a Economic Statistics Centre of Excellence, United Kingdom^b University of Augsburg, Department of Economics, Germany^c University of Birmingham Centre on Household Assets and Savings Management (CHASM), United Kingdom^d Athens University of Economics and Business, Department of Economics, Greece

ARTICLE INFO

JEL classification:

C53
C83
D22
D84
E32

Keywords:

Expectations
Firm data
Forecast errors
Survey data

ABSTRACT

We develop a novel methodology to quantify forecasts based on qualitative survey data. The methodology is generally applicable when quantitative information is available on the realization of the forecasted variable, for example from firm balance sheets. The method can be applied to a wide range of panel datasets, including qualitative surveys on firm-level forecasts or household expectations. As an application, we employ a panel of Greek manufacturing firms and quantify firms' forecast errors of own sales growth. In this context, we conduct a variety of exercises to demonstrate the methodology's validity and accuracy.

1. Introduction

Given the dynamic nature of economic decisions, expectations play a major role in economic behavior. Economic models that describe economic agents' 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, since survey-based measures for expectations are typically categorical, some important questions cannot be answered. For instance, whether firms make substantial errors in their forecasts and what are their statistical properties. Our paper provides a remedy to this obstacle, as we develop a novel methodology that converts categorical survey data on expectations to continuous quantities. Given the broad agreement in the profession on the advantages and usefulness of quantitative forecasts, there is a recent move towards surveys designed to include

[☆] We thank the editor James Bullard, an anonymous associate editor and two reviewers for their guidance and suggestions. We also thank Fabio Antoniou, Danilo Cascaldi-Garcia, Wanyu Chung, Yiannis Dendramis, Martin Ellison, Christos Genakos, Bruce Hansen, Menelaos Karanasos, Yiannis Karavias, Tryphon Kollintzas, Alistair Macaulay, Bartosz Mackowiak, Kaushik Mitra, Thanasis Stengos, Mirko Wiederholt and Klaus Wohlrabe as well as seminar participants at the IAAE 2021 conference and at the University of Birmingham for extremely useful comments and suggestions. We particularly thank Foteini Thomaidou from IOBE for detailed insights about the survey data. All remaining errors are our own. This work was produced using statistical data from ONS, and we are also thankful for their support in data provision. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates. This paper was previously circulated as part of the work titled "Quantifying Qualitative Survey Data: New Insights on the (Ir)Rationality of Firms' Forecasts".

* Corresponding author.

E-mail addresses: botsis.alex@gmail.com (A. Botsis), christoph.goertz@uni-a.de (C. Görtz), plutarch@aub.gr (P. Sakellaris).

quantitative features. As the move to quantitative surveys is recent, they do not provide long time series or historical data and might be restrictive in terms of the forecasting horizons they cover. Furthermore, many useful surveys remain qualitative. Therefore, our methodology will remain relevant and important in the future.

We propose a novel methodology that uses qualitative survey data on the direction of firm sales growth forecasts and quantitative data on realized sales growth from financial statements, to derive a quantitative estimate for firms' sales growth forecasts. We apply this method to a dataset that matches confidential information on firms' monthly qualitative forecasts on their own sales growth together with annual quantitative balance-sheet information on sales growth. The dataset covers Greek manufacturing firms for the period of 1998 to 2015.

In order 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. We use higher-frequency (monthly) qualitative survey data on expected sales growth together with lower-frequency (annual) quantitative data on realized sales growth to estimate quantified expected annual sales growth. Retaining the panel dimension comes with new challenges, such as dealing with unobserved heterogeneity and an omitted variable problem. This requires identifying assumptions that allow us to derive two nonlinear equations. The first one relates unobserved quantitative expected sales growth to observable variables, and the second one relates observed realized quantitative annual sales growth to observable variables. The key is that both of these relationships depend on the same parameters. Then, we estimate the common parameters from the second equation using Nonlinear Least Squares (NLS), and use these estimated parameters in the first equation to derive fitted values for quantitative expectations on sales growth.

In particular, our quantification model uses the fact that the firms' qualitative survey expectations are posted each month with a forecasting horizon of three months. During a given year, firms base their qualitative survey expectations on unobserved quantitative monthly expectations. We first assume that the quantitative expectations are linearly correlated with the annual quantitative forecast the firms formed for the whole of this year during the year before. Second, we assume that the firms' annual forecast can be decomposed to the quantitative monthly forecasts. These two assumptions permit us to algebraically derive a nonlinear equation that maps the quantitative annual forecast to the observed qualitative survey responses. Finally, using the fact that the realized growth rate is the sum of the expectations plus the forecast errors we can modify this first nonlinear equation to map the realized growth rates to the qualitative survey expectations. We can estimate the parameters of this second nonlinear equation with NLS.

This methodology can be applied to a wide range of applications and datasets and is not limited to quantifying firm forecasts. The only requirement is that the researcher can combine two types of data: (i) categorical survey data on expectations, with high time-series frequency within each unit; (ii) quantitative realizations of the corresponding variables with lower time-series frequency within each unit. Effectively, our quantification model aggregates the higher frequency qualitative responses into the lower frequency quantitative variable. For example, in our paper, we aggregate the categorical survey expectations from the firm-month frequency to the firm-year frequency that the quantitative realizations have. We additionally demonstrate that our quantification methodology is applicable to datasets with high- and low-frequency panels that have a short time-series dimension.

We provide evidence of external validity and accuracy for our methodology in four ways. First, we show that our quantified estimates on sales growth expectations are fully consistent in terms of sign with the corresponding qualitative survey-based expectations. Second, in an analysis of firm sales growth, we construct a small dataset of UK manufacturing firms that contains monthly qualitative survey expectations and the annual realizations from balance sheets, which allows us to use our methodology to derive estimates for annual forecast errors. Importantly, for each firm, the dataset also includes annual quantitative survey expectations, which we employ as a benchmark. Comparing our estimated annual forecast errors with the directly observable benchmark forecast errors confirms the accuracy of our quantification methodology. Such an exercise can only be conducted using a dataset that includes quantitative forecasts, made by the same forecaster at a high as well as a lower frequency. In practice, this is challenging to do due to the rare availability of such data on quantitative firm-level expectations. In fact, this dearth of data highlights the need for and value of our quantification methodology, which allows researchers to utilize the large number of qualitative surveys to quantify expectations. Third, we perform a Monte Carlo exercise that provides a benchmark based on simulated data. We find forecast errors based on our methodology are highly accurate when compared with forecast errors based on the underlying artificial 'true' data. Fourth, we run extensive robustness checks and provide additional evidence which add validity to the assumptions that underpin our model.

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.¹ Theil (1952) and Anderson (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. Furthermore, it is not limited to the quantification of firms' production, but can be applied to any variable in principle, given the data requirements outlined above. Importantly, our method can be used on the entire sample rather than only on a potentially small subset of firms.

¹ While this holds in the context of the quantification of individual qualitative forecasts, there is a large literature on the quantification approaches for the cross-sectional disagreement measured via qualitative survey responses, see e.g. Mokinski et al. (2015) and the references therein.

Important early work on the use and pitfalls of survey data to analyze how firms form expectations includes de Leeuw and McKelvey (1981) and Nerlove (1983). Our work is part of a now fast growing literature that uses information from surveys to understand firms' decision making. Born et al. (2023) 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. (2020) use novel Japanese data to study how firm characteristics affect their GDP forecasts. To the best of our knowledge, these two datasets are the only ones constructed so far to contain categorical firm survey data with corresponding quantitative data, e.g. from balance sheets or national accounts. We contribute to the survey literature by providing a novel dataset that combines responses to a rich firm-level survey with the corresponding balance sheet information for Greece. Our empirical results point to the importance of further work on merging existing quantitative datasets with qualitative survey data.² Applying our quantification methodology would then allow for a deeper understanding of how firms or households form expectations and their economic impact. There are many other contributions in the literature that use survey data to help our understanding of firm-level and aggregate variables. Enders et al. (2019) for example use German data from the IFO Business Survey to study how monetary policy announcements affect firms' expectations. Bachmann and Zorn (2020) use the IFO Investment 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.

The rest of the paper is organized as follows. Section 2 discusses the data. Section 3 lays out our methodology to quantify firms' forecasts and describes the characteristics of the estimated forecasts and the resulting forecast errors. Section 4.1 applies our quantification methodology to derive quantitative forecasts of Greek firms' own sales growth. Section 4.2 provides evidence of external validity, validity and accuracy of our methodology and some robustness checks. Section 5 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 firm-level balance sheets and income statements. We obtain this data from ICAP S.A., a private consultancy firm, which collects and digitizes this information from official publicly available records. The financial statements are compiled by certified auditors (chartered accountants) and are used, among other things, for reporting to tax authorities and investors, by commercial banks for credit decisions, and by the central bank for credit rating information. They are available to us at an annual frequency from 1998 to 2015, which determines the time span of our sample. 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 firm-level sales from the financial statements, which is deflated using the implicit gross value added deflator from Eurostat.³

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

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 focus on the sales growth expectations in the manufacturing sector which includes 38% of survey observations and 36% of observations in the financial statements data. The relevant (translated) question in the survey is

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

The qualitative survey responses are coded in the data as +1/0/-1 indicating increase/remain unchanged/decrease, respectively. In the following, we label the variable that includes the responses of firm i in month m to question D.2 as XS_{im}^e . The qualitative survey variable on expected sales developments, XS_{im}^e , has a quantitative counterpart with realized 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 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 the Online Appendices A.2 and A.3. Our cleaned and merged dataset includes 799 firms with 25,764 monthly responses from the survey on the above two questions and 4,104 annual balance sheet observations on sales. Table 1 provides an overview of the firms in our sample. Our sample includes very small firms but also large firms with close to 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

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

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.

Table 1
Sample Characteristics.

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

the Online Appendix A.4 we provide evidence that our sample is representative for the manufacturing sector and establish in several exercises the high quality of the survey responses. In the Online Appendix B.5 we show the distribution of survey expectations on sales growth and document their evolution over time.

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. In this section we develop a novel quantification methodology to derive a quantitative estimate for firms' sales growth forecasts.

3.1. Quantifying expected sales growth

Consider the expected annual growth rate of sales for firm i in year y , $x_{iy}^e \triangleq \mathbb{E}[x_{iy} | \mathcal{F}_{i,y-1}]$, that is based on an information set \mathcal{F} at the end of year $y-1$. Additionally, we define firm i 's expectation about average sales growth in the next three months as $x_{im}^e \triangleq \mathbb{E}[x_{i,\{m,m+1,m+2\}} | \mathcal{F}_{i,m-1 \in y}]$, where $x_{i,\{m,m+1,m+2\}}$ is the average growth rate of sales for the following three-month period. Note that this expectation is formed based on an information set at the end of month $m-1$. This quantitative monthly sales forecast is consistent, in terms of the information set and forecasting horizon, with the qualitative survey forecast XS_{im}^e .

One can describe a firm's expected annual sales growth with its monthly components as

$$x_{iy}^e = \mathbb{E}_{i,y-1} \sum_{m \in y} x_{im}^e. \quad (1)$$

Intuitively, equation (1) states that the forecast that a firm makes in $y-1$ for the whole of year y can be decomposed to its forecast in $y-1$ about its subsequent monthly frequency forecasts made during y . For further details, see also Online Appendix B.1. To simplify our exposition, we ignore for now any seasonality in the monthly growth rates, but we address this at a later step when we discuss our estimation strategy.

While we do not observe quantitative expectations of sales growth in equation (1) — x_{iy}^e and x_{im}^e — our dataset includes qualitative survey responses on the expected change in sales, XS_{im}^e . The aim of this section is to derive a quantitative estimate for annual expected sales growth, x_{iy}^e , using the observed qualitative survey responses and the realized annual sales growth rates from the firm's financial statements.

As a first step towards this, we follow Pesaran (1987) and assume that for each firm its 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 quantitative monthly expectation variable is positive, $x_{im}^{e,+}$, or negative, $x_{im}^{e,-}$. This is the first identifying assumption (ID1) we make to quantify firms' forecast errors. It can be formalized as

$$x_{im}^{e,+} = \alpha + \gamma_1 x_{iy}^e + v_{im}^+, \quad \text{and} \quad x_{im}^{e,-} = -\beta + \gamma_2 x_{iy}^e + v_{im}^-, \quad [\text{ID1}] \quad (2)$$

where v_{im}^+ and v_{im}^- are the error terms.⁵ Any potential monthly serial autocorrelation and correlation across firms (month-specific fixed effects) in these error terms are not of concern, because we show later that the aggregation at the firm-year frequency eliminates them. 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 refrain from accounting for this state dependence in the notation for now to ease the exposition.

⁵ For our dataset this assumption directly links to the monthly survey forecasts in which the information updates each month. However, for other datasets this can be relaxed to

$$\mathbb{E}_{i,y-1} x_{im}^{e,+} = \alpha + \gamma_1 x_{iy}^e + \mathbb{E}_{i,y-1} v_{im}^+, \quad \text{and} \quad \mathbb{E}_{i,y-1} x_{im}^{e,-} = -\beta + \gamma_2 x_{iy}^e + \mathbb{E}_{i,y-1} v_{im}^-. \quad (3')$$

We thank an anonymous referee for pointing this out.

Equations (2) are not formulated to conduct any inference about how firms make their monthly forecasts, but merely to reflect that for each firm the annual expected growth rate should be correlated with the corresponding monthly components. In fact, this linear correlation in equations (2) can be used to eliminate the unobserved variable x_{im}^e from equation (1). If we rewrite x_{im}^e as $x_{im}^e = x_{im}^{e,+} + x_{im}^{e,-}$ in the sum operator of equation (1) and also combine it with (2) we obtain (detailed derivations are shown in Appendix A)

$$x_{iy}^e = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \quad \text{with} \quad \xi_{iy} = \frac{\psi_i}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}, \quad (3)$$

where ψ_i is the unobserved firm heterogeneity (fixed effect), and we define

$$P_{iy} \triangleq \sum_{m \in y} \mathbb{1}_{[XS_{im}^e=1]}, \quad \text{and} \quad N_{iy} \triangleq \sum_{m \in y} \mathbb{1}_{[XS_{im}^e=-1]}, \quad (4)$$

to ease the notation. P_{iy} (N_{iy}) denotes the number of months per year that record a rise (decline) in expected sales of firm i . These qualitative variables are directly available from the survey data so that we can observe P_{iy} and N_{iy} .

To derive equation (3), we have additionally assumed that

$$\mathbb{E}_{i,y-1} P_{iy} = P_{iy} \quad \text{and} \quad \mathbb{E}_{i,y-1} N_{iy} = N_{iy}, \quad (5)$$

where both sides of these two equations refer to firm forecasts. The assumption embedded in equation (5) is that, during year y , firm i makes as many forecasts of positive and negative sales growth when responding to the survey as it expected to do at the end of year $y - 1$. Note that the assignment to individual months of positive and negative forecasts is unconstrained by (5) as long as the proportion is constant. Intuitively, equation (5) states that firms do not drastically update their information during year y when making monthly forecasts. This may happen because updating information on a monthly basis is costly. In Online Appendix B.3, we provide some analysis supporting that this assumption is realistic for our dataset.⁶

To estimate equation (3) we need to take some additional steps since we do not observe the quantitative expectations of annual sales growth, x_{iy}^e , in the data. In fact, deriving quantitative sales growth expectations was our goal in the first place. Instead, if we had estimates for the parameters and knowledge of the error term — and given that we observe P_{iy} and N_{iy} — we could use equation (3) to derive fitted values for x_{iy}^e . The next steps of the derivation are undertaken to facilitate exactly this.

We know that for each firm i , realized sales growth in year y is the sum of expected sales growth in that year and a forecast error, $x_{iy} = x_{iy}^e + x_{iy}^{fe}$. Using this expression to replace x_{iy}^e in equation (3) yields after rearranging

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

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

Omitted Variable Problem. To ease the notational burden in this section, we use equation (3) to define the conditional expectation of the quantitative sales growth expectations as

$$\tilde{x}_{iy}^e \triangleq \mathbb{E}[x_{iy}^e | P_{iy}, N_{iy}] = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}, \quad (7)$$

which can be thought as the ‘econometrician’s estimate’ of the firm’s expectation.

To obtain consistent estimates of the parameters in equation (6), we need the composite error term $x_{iy}^{fe} + \xi_{iy}$ to be mean independent of the nonlinear function \tilde{x}_{iy}^e (see Davidson and MacKinnon (2004)). We proceed now to show this. Note that the forecast error, x_{iy}^{fe} , is mean independent from the forecast x_{iy}^e .⁷ Since $\mathbb{E}[x_{iy}^{fe} | x_{iy}^e] = 0$ holds, it also implies that the firm’s forecast error is mean independent from the econometrician’s estimate of that forecast, $\mathbb{E}[x_{iy}^{fe} | \tilde{x}_{iy}^e] = 0$. We provide a proof of this statement in Online Appendix B.2 using the law of iterated expectations and the fact that in our model the econometrician’s estimate of the firm’s forecast is not more informed than the firm’s as it is based on the firm’s own monthly forecasts. Intuitively, firms’ expected sales growth, x_{iy}^e , whether rational or, not cannot ex-ante forecast their forecast error, otherwise firms would have incorporated this information in their expectation to reduce the forecast error. The same must hold then also for any estimates, \tilde{x}_{iy}^e , of firms’ sales growth expectations.

⁶ The assumption in equation (5) might be violated during major events (e.g. the 2020-2021 pandemic) when a large shift of expectations can occur within that turbulent year. We recommend dropping these years from the estimation sample. In our dataset, the rapid expansion of the noughties in Greece (until 2008) was followed by a bust and a prolonged contraction (output declined more than 25% by the end of 2014). For this reason we split our sample in two periods, 1999-2008 and 2009-2015, when we estimate the parameters α , β , γ_1 and γ_2 .

⁷ Indeed, $\mathbb{E}[x_{iy}^{fe} | x_{iy}^e] = \mathbb{E}[x_{iy} - x_{iy}^e | x_{iy}^e] = x_{iy}^e - x_{iy}^e = 0$. This does not imply rational expectations, because mean independence from the firm’s forecast does not imply mean independence from the information set that was used by the firm for that forecast.

Having established the forecast error's mean independence of \tilde{x}_{iy}^e , and in order to obtain consistent estimates of the parameters in equation (6), it remains to be shown that $\mathbb{E}[\xi_{iy}|\tilde{x}_{iy}^e] = 0$. A sufficient condition for mean independence of the error term, $\mathbb{E}[\xi_{iy}|\tilde{x}_{iy}^e] = 0$, to hold is that $\mathbb{E}[\xi_{iy}|\{XS_{im}^e\}_{m \in y}] = 0$. In Online Appendix B.2 we provide a formal proof of this statement. This leaves us with the task to control for the unobserved firm heterogeneity that is likely to make ξ_{iy} correlated with $\{XS_{im}^e\}_{m \in y}$. We turn to this next.

Unobserved Firm Heterogeneity. From equation (3), we know that the numerator of the error term ξ_{iy} is the unobserved firm heterogeneity ψ_i . We need to account for the effect of the unobserved heterogeneity which is in fact an omitted variable and is endogenous. The reason is that firm heterogeneity is related to the entire history of XS_{im}^e , so that $\mathbb{E}[\psi_i|\{XS_{im}^e\}_{m=1,\dots,T_i}] \neq 0$. Note that the notation $\{XS_{im}^e\}_{m=1,\dots,T_i}$ denotes the entire history of months m for variable XS_{im}^e , where T_i is firm i 's total number of monthly observations.⁸

To control for unobserved heterogeneity, we need to approximate $\mathbb{E}[\psi_i|\{XS_{im}^e\}_{m=1,\dots,T_i}]$. The structure of the nonlinear equation (3) 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. A widely used approximation for this purpose is the one suggested in Mundlak (1978).⁹ The original Mundlak (1978) specification is linear, but in the following we additionally include a second-order term due to the nonlinearity of equation (3), and we show later in Section 4.2.4 that this quadratic approximation is adequate. Therefore, our second identifying assumption is that the conditional expectation of the unobserved firm heterogeneity in the error term ξ_{iy} is

$$\mathbb{E}[\psi_i|\{XS_{im}^e\}_{m=1,\dots,T_i}] = \delta_1 \overline{XS_i^e} + \delta_2 (\overline{XS_i^e})^2, \quad [\text{ID2}]$$

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

$$\psi_i = \delta_1 \overline{XS_i^e} + \delta_2 (\overline{XS_i^e})^2 + \omega_i, \quad (8)$$

where ω_i is the part of the firm-specific heterogeneity that is mean independent from the survey expectations, that is $\mathbb{E}[\omega_i|\{XS_{im}^e\}_{m=1,\dots,T_i}] = 0$; and $\overline{XS_i^e} = \frac{1}{T_i} \sum_{m=1}^{T_i} XS_{im}^e$ is the simple arithmetic mean of the survey variable XS_{im}^e across time for each firm i . Intuitively, ID2 and equation (8) control for the firms' forecasting behavior and their overall firm-specific optimism or pessimism when they respond to the survey. We can now substitute equation (8) for ψ_i in the numerator of ξ_{iy} , obtaining

$$\xi_{iy} = \frac{\delta_1 \overline{XS_i^e} + \delta_2 (\overline{XS_i^e})^2 + \omega_i}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \quad (9)$$

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

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

where

$$\tilde{\xi}_{iy} = x_{iy}^{fe} + \frac{\omega_i}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}. \quad (11)$$

Overall, equation (10) is estimable because the error term $\tilde{\xi}_{iy}$ is mean independent of the explanatory variables. We provide a formal proof of this statement in Online Appendix B.2. This addresses the issue of the unobserved heterogeneity in equation (6), so that we can obtain consistent estimates of the coefficients of interest, α , β , γ_1 and γ_2 .¹⁰

Controlling for seasonality. The exposition above describes our quantification method, yet we can refine it by controlling for the seasonality in the monthly forecasts of the firm. We define the weights

$$\mathcal{W}_{im}^+ = W_{im} \mathbb{1}_{[XS_{i,m}^e = +1]}, \quad \text{and} \quad \mathcal{W}_{im}^- = W_{im} \mathbb{1}_{[XS_{i,m}^e = -1]}. \quad (12)$$

They consist of two components. The first component in each weight, W_{im} , accounts for the fact that some months have a higher level of firm sales than others and therefore represent a larger share of the final annual outcome, which is the seasonality. It is defined as

⁸ This notation is distinct from $\{XS_{im}^e\}_{m \in y}$, used above, which refers to all months m in year y .

⁹ See e.g. Bartelsman et al. (1994), Semykina and Wooldridge (2010), Kosova (2010) and Triguero and Córcoles (2013)). The Mundlak (1978) approximation is the standard tool used in nonlinear models in panel data. In linear models, it is equivalent to the least squares dummy variable and the standard within estimator.

¹⁰ The error term, $\tilde{\xi}_{iy}$, in equation (10) is likely to be heteroscedastic and autocorrelated within each firm. When we estimate such an equation, we will use the heteroscedasticity robust estimator for the standard errors, which addresses both problems — this robust estimator treats errors as clustered within cross-sectional units.

$$W_{im} \triangleq \frac{w_{im}}{\sum_{m \in y} w_{im}}, \quad (13)$$

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 due to seasonality, and this month is more important than others for the annual outcome. Our theoretical decomposition allows for individual weights for each firm i , but 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.¹¹ The second component of the weights \mathcal{W}_{im}^+ and \mathcal{W}_{im}^- is the indicator variables that we used in deriving equation (3) and correspond to expectations of positive or negative sales growth.

The introduction of the seasonal weights does not affect the algebraic manipulations that lead to equation (3) nor does it affect our omitted variable problem and the approximation of the firm-specific heterogeneity. One simply has to use these weights to compute the variables P_{iy} and N_{iy} for equation (10) as follows

$$P_{iy} = \sum_{m \in y} \mathcal{W}_{im}^+, \quad \text{and} \quad N_{iy} = \sum_{m \in y} \mathcal{W}_{im}^-. \quad (14)$$

Summary of the Quantification Method. We have derived two nonlinear equations. Equation (10) relates observed quantitative annual sales to observable variables and the identifying assumption ID2 ensures that the coefficient estimates are consistent. Equation (3) relates unobserved quantitative expected sales growth to observable variables. The key is that both of these relationships depend on the same parameters. We estimate the parameters from equation (10) using Nonlinear Least Squares (NLS), and use these estimated parameters in equation (3) to derive fitted values for quantitative expectations on sales growth.

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

1. Compute the weighted shares of months per year that record a rise (decline) in expected sales P_{iy} (N_{iy}) from survey data, using equation (14), with the weights defined in equations (12) and (13).
2. Compute the firm heterogeneity proxies, \overline{XS}_i^e and $(\overline{XS}_i^e)^2$, based on the arithmetic mean (across time for each firm i) of the qualitative survey variable XS_{im}^e .
3. Estimate equation (10) using NLS. Run the estimation separately for the boom ($y \leq 2008$) and bust period ($y > 2008$).¹²
4. Use the NLS estimated coefficients of equation (10) to compute the fitted values, \hat{x}_{iy}^e , for quantified sales growth forecasts from equation (3).

Our parameter estimates of the NLS estimation of equation (10) are documented in Section 4.1.1 below. The difference between the sales growth rate available from the financial statements, x_{iy} , minus 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} .

Our methodology to quantify forecasts and forecast errors is generally applicable to variables other than sales growth. It is applicable to any qualitative (survey based) variable on future developments, as long as a quantitative corresponding variable on realization is available at a lower frequency. Even in datasets where the time dimension of the panels is short, our methodology remains applicable. First, if the dimension of the panel with the high-frequency observations is ‘shorter’ than the monthly frequency in our data, then there would be some loss of accuracy in the estimated parameters of the NLS equation. We show that the loss of accuracy is small using a Monte Carlo simulation in Section 4.2.2. Second, where the low-frequency panel is short, there are no consequences for accuracy nor for consistency. Indeed, to achieve consistency we have shown that the omitted variable problem can be ignored and the unobserved firm heterogeneity can be proxied with the Mundlak (1978) fixed effects proxy. The omitted variable is not affected by the time-series length of the panels, and equation (8) would still provide a valid proxy of the unobserved firm heterogeneity if the time-series length per panel at the low-frequency data was short.

4. Application of the quantification methodology

In this section, we apply our quantification methodology using the data set introduced in Section 2. We estimate Greek firms’ forecast errors of their own sales growth in Section 4.1.1 and provide the reader with an overview about the panel-data estimates in Section 4.1.2. Section 4.2 uses the forecast error estimates for Greek firms, as well as artificially generated data, and data of UK firms for various tests on the accuracy and validity of our quantification methodology.

¹¹ We use 2-digit seasonally unadjusted 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.

¹² In Online Appendix B.4 we outline an alternative procedure that allows the parameters to be state dependent and tests for these state dependencies.

Table 2
NLS Estimation of Equation (10).

Coefficients	(1)	(2)
	Dependent Variable: x_{iy}	
α	0.190**	0.104**
β	0.151*	0.238***
δ_1	-0.0255	-0.127***
δ_2	-0.00215	-0.0530
γ_1	-0.366	-0.446
γ_2	-0.179	0.0712
Firm-Year Observations	2,471	1,397
R^2	0.043	0.057
Period	$y \leq 2008$	$y > 2008$

We use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance. We also use seasonality weights based on equation (14). Column (1) shows results for the boom period up to 2008 and column (2) for the following recession.

4.1. An application on Greek firms' forecasts of own sales growth

4.1.1. Baseline nonlinear least squares estimation results

Table 2 reports the results of the NLS estimation of equation (10). Column (1) shows estimation results for the boom period up to 2008 and column (2) for the following recession. As a reminder, α and $-\beta$ are the constant terms in the positive and negative continuous monthly forecasts of ID1. We observe that the constant of the positive monthly forecasts is larger during the boom than in the bust which is consistent with our economic intuition. Moreover, the constant of the negative monthly forecast is lower during the bust than in the boom, which is also consistent with our economic intuition.

4.1.2. Descriptive statistics on the quantified forecast errors

Fig. 1 shows the distribution of forecast errors. We report moments on this distribution in Table 3. The average forecast error in our sample is zero and slightly larger than the median (-0.03). This implies that the median forecast on sales growth is three percentage points more optimistic than the subsequent realization. Overall, a number of forecast errors made by firms are small (in absolute value), as these are centred close to zero, but still a significant number of forecast errors made are quite substantial given the high standard deviation. A non-negligible number of firms make forecast errors that imply 50% higher or lower sales than expected.

In Table 3, we also observe that the mean value of the quantified forecast errors appears to be procyclical (significance at 1%). Quantified forecasts are on average 2% more optimistic than realizations during the bust period and 1% more pessimistic during the boom, which accords with our economic intuition. The standard deviation is not countercyclical as the literature has found (see for example Bachmann et al. (2013)). We note, however, that the observable qualitative survey-based forecast errors do not display countercyclical standard deviation either, which suggests that it is not the quantification that has eliminated this property.¹³ We also note that the case of Greece is particularly different from other advanced economies, because it experienced a prolonged period of expansion (boom) that was followed by a prolonged and particularly severe contraction (bust).

4.2. External validity, validity of our assumptions and robustness checks

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

Additionally, we run a number of robustness checks to demonstrate that our quantified forecast errors remain robust to relaxing some of our assumptions. First, we use alternative weights to compute the variables P_{iy} and N_{iy} for equation (10). Second, we use a cubic approximation of the Mundlak (1978) fixed effects that proxy for the unobserved firm-specific heterogeneity. Third, we relax our assumption that the parameters α , β , γ_1 and γ_2 are common for all firms and allow them to vary with i .

¹³ We compute the qualitative survey-based forecast errors following Bachmann et al. (2013). The IOBE survey includes a qualitative question on current sales growth which is framed as follows:

Question A.2: During the previous 3 months, your total sales have increased/remaind unchanged/decreased.

We label the qualitative question A.2 as XS_{im}^f , and we define the qualitative forecast error as $XS_{im}^{fe} = XS_{i,m+3} - XS_{im}^e$.




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

Table 3
Descriptive Statistics for Quantified Sales Growth Forecast Errors.

	Mean	Median	Stand. dev.
Full Sample	0.00	-0.03	0.34
Boom	0.01	-0.02	0.34
Bust	-0.02	-0.05	0.35

The boom (bust) period spans the years 1998-2008 (2009-2015). The mean quantified forecast error differs between the boom and the bust with statistical significance at 1%. We used a random effects regression of the forecast errors on the indicator variable of boom with clustered (robust) standard errors.

4.2.1. Directional consistency of estimated forecasts with the survey data

We can use the observed survey data on the direction of expected sales growth to benchmark how well our quantified forecasts match the direction of expected sales growth. To facilitate the comparison of the monthly survey data with our annual forecast estimates, we annualize the survey responses by computing a weighted yearly average $\sum_{m \in y} \mathcal{W}_{im} [X S_{im}^e]$, where the weights are based on equation (14). The distributions of the raw monthly and annualized survey expectations are reported in the Online 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 estimates of quantified forecasts against the annualized survey-based qualitative forecasts, we split responses in each of these two variables into three categories — positive, zero or negative — and cross-tabulate the three directions. Table 4 reports how well our quantified forecasts match the direction of the annualized observable ones. The main diagonal shows the share of observations that are directionally consistent across the two variables when classified as either positive, zero or negative. Overall, the direction of our quantified forecasts is highly consistent with the ones of the annualized survey responses — their direction coincides for 93.98% of all 3,868 firm-year observations (the sum of the main diagonal).

The small share of observations for which the directions do not coincide can be explained by the absence of information on scale in the qualitative survey data. In practice, even if the majority of all monthly forecasts in one year point in the same direction, a single large monthly forecast in the opposite direction could dominate the annual response. This however cannot be captured by annualizing purely qualitative monthly forecasts. For this reason, we also report in Table 4 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 B shows results for this restricted sample which comprises 26% (999 firm-year observations) of the observations of the full

Table 4

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

	Panel A: Entire Sample Quantified forecasts			Panel B: Restricted Sample Quantified forecasts		
	Negative	Zero	Positive	Negative	Zero	Positive
Negative Forecasts	23.94%	0.00%	1.45%	11.21%	0.00%	0.00%
Zero Forecasts	0.26%	14.71%	0.34%	0.00%	56.96%	0.00%
Positive Forecasts	3.98%	0.00%	55.33%	0.00%	0.00%	31.83%
	Directional Consistency: 93.98%			Directional Consistency: 100.00%		

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. The restricted sample only considers annualized survey observations for which, in a given year, all underlying monthly observations report forecasts in the same direction. Panel A with the 'entire sample' comprises 3,868 firm-year observations. Panel B with the 'restricted sample' comprises 999 firm-year observations.

sample used in Panel A. It is evident that now the direction of all quantified forecasts is consistent with the ones of the annualized survey responses.¹⁴

Overall, our exercise shows that forecasts based on our quantification methodology are fully consistent with the direction of sales growth implied by the qualitative survey responses. We next turn to a Monte Carlo exercise that uses simulated data to infer how precisely our estimates match the magnitude of underlying true forecast errors.

4.2.2. Monte Carlo experiment

It is important to understand how well forecast errors based on our methodology match, in terms of magnitude, true quantitative forecast errors. In practice, this is challenging to do due to the unavailability of data on quantitative firm-level expectations. This dearth of data was, indeed, the key motivation for developing the quantification methodology proposed in this paper. The vast majority of surveys contain qualitative questions about firms' future developments. If quantitative survey-based expectations are available at all, then they either focus on aggregate rather than firm-specific variables or have a limited sample size. To overcome this obstacle, we perform a Monte Carlo exercise that provides a benchmark based on simulated data. In particular, we simulate data on firm (continuous) annual sales growth realizations, as well as corresponding qualitative and quantitative expectations. We then use the data on realized sales growth and qualitative expectations as inputs to the quantification methodology of Section 3.1 and generate estimates for quantified sales growth expectations. Subsequently, we evaluate the accuracy of the estimated forecast errors in comparison to those based on the underlying artificial 'true' data.

We generate 1,000 sets of random artificial data, each one of which mimics the structure of the true dataset in terms of number of firms and its unbalanced nature of firm-year-month observations. Details about the data generation are provided in the Online Appendix B.6. This Appendix documents that the underlying processes and their calibration to generate the artificial data are carefully guided by the characteristics and statistics of the observable financial statements and the survey data. We further highlight in this appendix that the simulated datasets match closely moments and statistics in the empirical data that have not been targeted during the calibration.

Panels A and B of Table 5 show the distribution of the true forecast errors and of the estimated ones, all based on the artificial datasets. The mean and median of their distribution are very close — for both moments the difference is only about one percentage point of sales growth. This is very small, particularly when recalling from Fig. 1 and Table 3 that the absolute median forecast error in our data is three percentage points and the empirical distribution has non-negligible mass at forecast error values as large as 50 percentage points of sales growth. The close correspondence between the estimated and the true forecast error can also be illustrated in a scatter plot. Fig. 2 contains the scatter plot for one artificial dataset (randomly chosen among the 1,000 draws). The forecast error pairs conform to the 45 degree line quite closely.

In Panel C of Table 5 we show the distribution of the estimated quantitative forecast error and the true quantitative forecast error when using *only the quarterly survey responses*. We do this exercise to assess the degree of the loss of accuracy in the case that the dimension of the panel with the qualitative survey-based observations is 'shorter' than the monthly frequency that we have in our data. Observe that the two distributions of the quantified forecast errors in Panels B and C are very close so any measurement error resulting from using quarterly instead of monthly observations is not substantial. Also the standard deviation of the moments across the 1,000 trials of the Monte Carlo is slightly higher in Panel C than B. This is to be expected because the reduction in the survey-based observations will result in losses in efficiency.

¹⁴ Results are fully directionally consistent even if we consider annualized observations for which at least 67% of underlying monthly survey responses of a firm for a particular year indicated sales forecasts in the same direction. This comprises 39% (1,492 firm-year observations) of the observations of the full sample used in Panel A.

Table 5

Distribution of the estimated quantitative forecast error and the true quantitative forecast error (both based on artificial data).

	5%	10%	25%	Median	Mean	75%	90%	95%
Panel A: True forecast errors								
Average	-0.654	-0.510	-0.268	0.000	0.002	0.269	0.511	0.654
St. dev.	0.018	0.014	0.012	0.011	0.009	0.012	0.014	0.018
Panel B: Quantification using all monthly survey responses								
Average	-0.646	-0.501	-0.259	0.011	0.013	0.281	0.524	0.669
St. dev.	0.014	0.011	0.010	0.009	0.008	0.010	0.012	0.014
Panel C: Quantification using only quarterly survey responses								
Average	-0.645	-0.499	-0.254	0.018	0.020	0.291	0.536	0.682
St. dev.	0.019	0.015	0.013	0.013	0.011	0.014	0.017	0.021

We report the average and standard deviation (St. dev.) across 1,000 random samples of artificial data of the descriptive statistics.




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

4.2.3. Validating forecast error accuracy and ID1 in a sample of UK firms

As explained above, the vast majority of firm-level surveys contain only qualitative questions. If surveys have quantitative features, these are typically limited. While this highlights the importance of developing methodologies to quantify qualitative survey responses, it makes it difficult to validate our methodology against survey-based quantitative forecast errors. In principle, we can do so if a dataset contains firm-level information on (i) monthly qualitative survey based forecasts for the three-month period ahead, (ii) quantitative annual survey forecasts, and (iii) annual realizations of the underlying variable. We have managed to obtain this information for a very limited sample of firms in the UK manufacturing sector. To the best of our knowledge, this is the only dataset that contains all three types of data required to inspect the accuracy of our methodology. In particular, we consider quantitative annual forecasts on firm's own turnover growth from the Management and Expectations Survey which was conducted by the Office for National Statistics (ONS) in 2017. During the same year, the Confederation of British Industry (CBI) independently collected qualitative monthly survey forecasts on firm's output growth.¹⁵ To obtain the annual realizations on turnover growth we match the survey data with the Financial Statements from Bureau Van Dijk's FAME dataset.¹⁶ Since the ONS and CBI surveys are conducted independently, the resulting matched sample is very small. It consists of 173 firm-month observations for qualitative survey forecasts on output growth, and 47 observations for annual quantitative forecasts on turnover growth and the corresponding realizations.

First, we implement our quantification methodology as follows. In the interest of statistical power, we fit the nonlinear equation (10) to the realized turnover growth from FAME for a sample of 2,502 firm-year observations, for the period from 2000 until 2016.

¹⁵ Details about the ONS survey can be found in Awano et al. (2018) and Bloom et al. (2021). The CBI's survey on forecasts has a similar structure as the one from the IOBE for Greece, as both are used to construct EU-wide index of business climate by the Directorate-General for Economic and Financial Affairs (see DGECPIN (2017)).

¹⁶ We thank Nick Bloom, Paul Mizen, Rebecca Riley and Michael Mahony for sharing the survey data and linking tables.

Table 6

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

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

Table 7

Distribution of baseline quantified forecast errors and quantified forecasts based on alternative weighting schemes.

	Min	5%	10%	25%	Median	Mean	75%	90%	95%	Max
Baseline	-0.993	-0.370	-0.282	-0.148	-0.028	-0.001	0.093	0.243	0.407	5.226
Constant weighting	-0.992	-0.370	-0.281	-0.147	-0.027	-0.001	0.093	0.242	0.407	5.226
Decreasing weighting	-0.967	-0.363	-0.276	-0.141	-0.025	0.004	0.100	0.248	0.419	5.226

We then compute the forecast errors according to the methodology outlined in Section 3.1 for the 47 firms for which quantitative annual survey forecasts are available. We compare these forecast error estimates with the quantitative forecast errors from the ONS survey. The distribution of the differences between the estimated and survey-based forecast errors is summarized in Table 6.

The overall distribution for differences in forecast errors shown in Table 6 is rather tight. Both the mean and median of this distribution are very close to zero suggesting unbiased estimation. We note that the standard deviation of the observable forecast errors is 0.31 with mean value 0, which means that 95 percent of the forecast errors likely lie between -0.6 and 0.6. Therefore, even a discrepancy between the estimated and the observed forecast error of 0.135 should not be considered substantial. This is striking also because the monthly survey question is concerned with output growth and the annual survey question with turnover growth, which are closely related, but may not be perceived by respondents as exactly equal.¹⁷ Overall, the results in Table 6 (UK firms data) and those in Table 5 (artificial data) suggest that our quantification methodology is reliable with a reasonably low measurement error. Table 5 further demonstrates that even in the tails of the distribution of the forecast errors our methodology remains reliable.

Second, we give evidence that supports our identifying assumption ID1 that the monthly expectations during a given year are linearly correlated with the corresponding annual expectations about the same year. We find that the monthly qualitative survey forecasts are correlated with their annual quantitative counterparts. We also find that this linear correlation does not vary within the year.¹⁸

4.2.4. Robustness checks

Alternative weighting schemes This section shows results based on two alternative weighting schemes used in equation (14). In particular, while our baseline weighting controls for seasonalities within the year, we consider as an alternative that all observations are weighted equally per year as well as with decreasing monthly weights.

The decreasing monthly weights are motivated by the fact that an expected increase in sales in the first months of the year might have a larger effect on the overall forecast of sales growth for the entire year. For example, consider a case (a) in which a firm expects an increase in sales in the first three months of the year and then monthly sales are expected to stabilize at a higher level for the rest of the year. Consider also an alternative example, case (b) where the same firm would expect monthly sales to remain constant during the first nine months of the year and expect an increase in the last three months. The true expected sales growth for the entire year would be higher in case (a) than in case (b), but the quantified expectations from our methodology would be equal for the two cases as the variables P_{iy} and N_{iy} would also be equal between case (a) and case (b).

To control for this, in our robustness exercise we apply weights in equation (13) that are decreasing with the months. That is, January has a weight of $w_{im} = 12$; February has $w_{im} = 11$; March has $w_{im} = 10$; ...; December has $w_{im} = 1$.

Table 7 shows the distribution of the quantified forecasts using (i) our baseline weighting scheme, (ii) constant monthly weights, and (iii) decreasing monthly weights. We observe that the differences in the three distributions are minimal. That is, the quantified forecast errors are robust to using alternative weighting schemes.

Cubic approximation of firm fixed effects To control for unobserved heterogeneity, we employ a quadratic approximation. In this section, we show that a finer approximation of the firm fixed effect is not required. In Table 8, we re-estimate our baseline equation (10) including the cubic term $(XS_i^e)^3$ to proxy for the firm fixed effect. That is

$$x_{iy} = \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1 \overline{XS_i^e} + \delta_2 (\overline{XS_i^e})^2 + \delta_3 (\overline{XS_i^e})^3}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \tilde{\xi}_{iy}, \quad (15)$$

where

¹⁷ In Online Appendix B.7, we show the close correspondence between the imputed and the true forecast errors in the MES and CBI data illustrated in a scatter plot.

¹⁸ We have more details about this exercise in Online Appendix B.8.

Table 8
Robustness check: Cubic approximation of firm-level
fixed effects.

Coefficients	Dependent Variable: x_{ij}	
α	0.196*	0.104**
β	0.151*	0.241***
δ_1	-0.0334	-0.147***
δ_2	-0.0104	-0.0380
δ_3	0.0225	0.0521
γ_1	-0.406	-0.370
γ_2	-0.215	0.0505
Firm-Year Observations	2,471	1,397
R^2	0.043	0.057
Period	$y \leq 2008$	$y > 2008$

Table shows estimates of equation (15). We use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance. Columns (1) and (3) show results for the boom period up to 2008 and columns (2) and (4) for the following recession.

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

The estimates clearly demonstrate that the coefficient of the cubic term is not statistically significant in any of the periods. The estimates are not sensitive to the inclusion of the cubic term and are close to those reported in Table 2.

Correlated random coefficients We turn now to relaxing our assumption that the parameters α , β , γ_1 and γ_2 are common for all firms and allow them to vary with i .

Our baseline estimation controls for firm-varying α and β , because we have assumed that the firm heterogeneity is part of the error terms of equations (2), and we treat it as firm fixed effects.¹⁹ In Appendix A we show how the firm fixed effects appear on the nonlinear equation (3), while ID2 and equation (8) control for the endogeneity they introduce.

Let us now examine the possibility that the parameters γ_1 and γ_2 vary with i in equations (3). Suppose that the true coefficients in the population are $\gamma_{1i} = \gamma_1 + \tilde{\gamma}_{1i}$ and $\gamma_{2i} = \gamma_2 + \tilde{\gamma}_{2i}$, where $\tilde{\gamma}_{1i}$ and $\tilde{\gamma}_{2i}$ are centred around 0 and distinguish the firm-specific component from the common one. We can rewrite the final equation that we will estimate as

$$x_{iy} = \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1 \overline{X S_i^e} + \delta_2 (\overline{X S_i^e})^2}{1 - \gamma_{1i} P_{iy} - \gamma_{2i} N_{iy}} + \xi_{iy}$$

$$= \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1 \overline{X S_i^e} + \delta_2 (\overline{X S_i^e})^2}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy} - \tilde{\gamma}_{1i} P_{iy} - \tilde{\gamma}_{2i} N_{iy}} + \tilde{\xi}_{iy},$$

where

$$\tilde{\xi}_{iy} = x_{iy}^{fe} + \frac{\omega_i}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy} - \tilde{\gamma}_1 P_{iy} - \tilde{\gamma}_2 N_{iy}}.$$

Owing to the sample restrictions in our data and the nonlinear form of our model we are unable to model the term $-\tilde{\gamma}_{1i}P_{iy} - \tilde{\gamma}_{2i}N_{iy}$ for each firm by estimating for instance the nonlinear equation for each firm independently.

As a robustness check of the consistency of our estimates, we approximate the firm-specific components of γ_1 and γ_2 using two firm-specific variables: (i) the age of the firm in its first appearance in the sample, age_i , (ii) and the size of the firm in its first appearance in the sample measured as the decile (values 1-10) of the value of the firm's real total net assets, K_i . We report that firms stay in the sample for on average 5 years and during this period we do not observe large swings in their net assets growth (each firm grows on average by 3.8% during its presence in the survey sample), so their decile size in the first appearance in the sample remains a reliable proxy for the size of the firm during its overall sample presence. Essentially, we assume that $\gamma_{1i} = \gamma_1 + \tilde{\gamma}_1 Z_i$ and $\gamma_{2i} = \gamma_2 + \tilde{\gamma}_2 Z_i$, where $Z_i = age_i, K_i$. That is, we assume a specific form for $\tilde{\gamma}_{1i}$ and $\tilde{\gamma}_{2i}$ in which age or size captures the firm specific effect. We used age because there is evidence in the literature that it affects forecast accuracy (see Tanaka et al. (2020)) and size following our economic intuition that larger firms might have more resources to make more rational forecasts. Then, we estimate in Table 9 the following equation

¹⁹ Indeed, we could rewrite the equations that show the linear correlation between the monthly quantitative expectations and the annual quantitative expectations as follows

$$x_{im}^{e,+} = \alpha + \alpha_l + \gamma_1 x_{iv}^e + \tilde{v}_{im}^+, \quad \text{and} \quad x_{im}^{e,-} = -\beta - \beta_l + \gamma_2 x_{iv}^e + \tilde{v}_{im}^-.$$

Then, we would assume $v_{im}^+ = \tilde{v}_{im}^+ + \alpha_i$ and $v_{im}^- = \tilde{v}_{im}^- + \beta_i$, which gives equation (2).

Table 9NLS Estimation of Equation (16) – Robustness for firm-varying γ_1 and γ_2 .

Coefficients	(1)	(2)	(3)	(4)
	Dependent Variable: x_{iy}			
	$Z_i = age_i$		$Z_i = K_i$	
α	0.176**	0.109**	0.193**	0.105**
β	0.141**	0.232***	0.145*	0.241***
δ_1	-0.0255	-0.128***	-0.0209	-0.130***
δ_2	-0.0321	-0.0579	-0.00824	-0.0477
γ_1	0.555	-0.883	-0.198	-0.717
γ_2	0.363	0.0438	0.449	0.155
$\tilde{\gamma}_1$	-0.0330*	0.0161	-0.0376	0.0521
$\tilde{\gamma}_2$	-0.0201	0.00196	-0.130	-0.0194
Firm-Year Observations	2,461	1,395	2,461	1,395
R^2	0.047	0.057	0.043	0.056
Period	$y \leq 2008$	$y > 2008$	$y \leq 2008$	$y > 2008$

The table shows estimates of equation (16). Columns (1) and (2): $Z_i = age_i$, i.e. age of the firm in its first appearance in the sample. Columns (3) and (4): $Z_i = K_i$, i.e. the size of the firm in its first appearance in the sample (decile of real total net assets). Columns (1) and (3) show results for the boom period up to 2008 and columns (2) and (4) for the following recession. We use robust standard errors and ***, ** and * indicates 1%, 5% and 10% significance.

Table 10Calculating firm-varying γ_1 and γ_2 at the average, $\gamma_l + \tilde{\gamma}_l \overline{Z_i}$ with $l = 1, 2$ and $Z_i = age_i, K_i$.

	(1)	(2)	(3)	(4)	(5)	(6)
	$Z_i = age_i$		$Z_i = K_i$		Baseline, $\tilde{\gamma}_l = 0$	
	Boom	Bust	Boom	Bust	Boom	Bust
$\gamma_1 + \tilde{\gamma}_1 \overline{Z_i}$	-0.19	-0.39	-0.4	-0.467	-0.366	-0.446
$\gamma_2 + \tilde{\gamma}_2 \overline{Z_i}$	-0.14	0.09	-0.2	0.06	-0.179	-0.071

age_i is the firm's age at the first appearance in the sample with an average of 25 years across firms; K_i is the firm's decile of total net assets at the first appearance in the sample with an average of 5 across firms.

$$x_{iy} = \frac{\alpha P_{iy} - \beta N_{iy} + \delta_1 \overline{X S_i^e} + \delta_2 (\overline{X S_i^e})^2}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy} - \tilde{\gamma}_1 \overline{Z_i} P_{iy} - \tilde{\gamma}_2 \overline{Z_i} N_{iy}} + \tilde{\xi}_{iy}, \quad (16)$$

with $Z_i = age_i, K_i$.

In Table 9, we show that the estimates of α and β and their significance during both the boom and the bust periods are very close to the ones we obtained from our baseline model in Table 2, while the coefficients $\tilde{\gamma}_l$ are hardly significant. At first look, the estimates of γ_1 and γ_2 in the new equation are different. However, if we compute the quantities $\gamma_l + \tilde{\gamma}_l \overline{Z_i}$, $l = 1, 2$ which is equivalent to the assumptions of our baseline estimates, then these quantities are very close to the ones we obtained in the baseline estimation — see Table 10.²⁰ The exception is $\gamma_1 + \tilde{\gamma}_1 \overline{age_i}$ and only during the boom period, which has a smaller magnitude than the baseline estimate. Overall, our robustness check suggests that our baseline estimates are unaffected by using age to proxy for firm-specific γ_1 and γ_2 .

5. Conclusion

In this paper, we develop a novel methodology to quantify qualitative survey data on expectations. This methodology is applicable generally when quantitative information is available on the realization of the forecasted variable. We apply this methodology to Greek firm data on sales growth. The survey of firm expectations we use for Greek firms is similar in structure to the ones used by all European Union countries at a monthly frequency. A key component of our methodology to produce quantified expectations estimates of sales growth using the qualitative survey data is to combine firm balance sheet data on realized sales.

Once we have quantitative estimates of firms' forecasts and forecast errors, we can answer important questions about firm expectations formation and economic behavior. Do firms make errors in forecasting their future sales that are predictable and display autocorrelation? If so, what does that reveal about firm behavior and the way they form expectations? In particular, does firms' behavior conform to the Full Information Rational Expectations (FIRE) hypothesis? What are the causes of forecast errors, and how

²⁰ $\overline{Z_i}$ is the average age ($Z_i = age_i$) across the sampled firms and is equal to 25 years; or the average decile of total net assets ($Z_i = K_i$) which is equal to 5.

do these errors affect firm production, investment, and financing decisions? These are important questions to be pursued in future research.

Appendix A. Derivation of Equation (3)

This appendix section shows how equation (3) can be derived using equations (1) and (2).

First, note that we can rewrite (1) as follows

$$x_{iy}^e = \mathbb{E}_{i,y-1} \sum_{m \in y} [x_{im}^{e,+} + x_{im}^{e,-}], \quad (17)$$

and we can naturally ignore the terms where $x_{im}^e = 0$.

Let us now use the indicator variables that take a value of unity if the expected sales growth rate x_{im}^e is either positive, $\mathbb{1}_{[x_{im}^e > 0]}$, or negative, $\mathbb{1}_{[x_{im}^e < 0]}$. Because we observe in the surveys the direction of x_{im}^e , we have that $\mathbb{1}_{[x_{im}^e > 0]} = \mathbb{1}_{[XS_{i,m}^e = +1]}$ and $\mathbb{1}_{[x_{im}^e < 0]} = \mathbb{1}_{[XS_{i,m}^e = -1]}$. Given the nature of the indicator variables, we can rewrite equation (17) as follows

$$x_{iy}^e = \mathbb{E}_{i,y-1} \sum_{m \in y} [\mathbb{1}_{[XS_{i,m}^e = +1]} x_{im}^{e,+} + \mathbb{1}_{[XS_{i,m}^e = -1]} x_{im}^{e,-}]. \quad (18)$$

Second, we substitute equation (2) into (18)

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

Then, we obtain

$$\begin{aligned} x_{iy}^e &= [\alpha + \gamma_1 x_{iy}^e] \mathbb{E}_{i,y-1} \sum_{m \in y} \mathbb{1}_{[XS_{i,m}^e = +1]} + \mathbb{E}_{i,y-1} \sum_{m \in y} \mathbb{1}_{[XS_{i,m}^e = +1]} v_{im}^+ \\ &\quad + [-\beta + \gamma_2 x_{iy}^e] \mathbb{E}_{i,y-1} \sum_{m \in y} \mathbb{1}_{[XS_{i,m}^e = -1]} + \mathbb{E}_{i,y-1} \sum_{m \in y} \mathbb{1}_{[XS_{i,m}^e = -1]} v_{im}^-. \end{aligned} \quad (19)$$

To simplify the notation, we define

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

where P_{iy} (N_{iy}) denotes the share of months within a year that indicate a rise (fall) in *expected* sales.

Next we assume that (equation (5) in main text)

$$\mathbb{E}_{i,y-1} P_{iy} = P_{iy} \quad \text{and} \quad \mathbb{E}_{i,y-1} N_{iy} = N_{iy},$$

where both sides of these equations refer to firm forecasts.

This assumption allows us to rearrange equation (19) to solve for x_{iy}^e :

$$x_{iy}^e = \frac{\alpha P_{iy} - \beta N_{iy}}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}} + \xi_{iy}, \quad \text{with} \quad \xi_{iy} = \frac{\mathbb{E}_{i,y-1} \sum_{m \in y} (\mathbb{1}_{[XS_{i,m}^e = +1]} v_{im}^+ + \mathbb{1}_{[XS_{i,m}^e = -1]} v_{im}^-)}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}}.$$

We can rewrite the term $(\mathbb{1}_{[XS_{i,m}^e = +1]} v_{im}^+ + \mathbb{1}_{[XS_{i,m}^e = -1]} v_{im}^-)$ with a more compact representation that is standard in the literature, by decomposing it into a firm-specific component ψ_i , a time-specific term ψ_m and an idiosyncratic term ψ_{im} . We obtain

$$\mathbb{E}_{i,y-1} \sum_{m \in y} (\mathbb{1}_{[XS_{i,m}^e = +1]} v_{im}^+ + \mathbb{1}_{[XS_{i,m}^e = -1]} v_{im}^-) = \sum_{m \in y} \mathbb{E}(\psi_i + \psi_m + \psi_{im} | \mathcal{F}_{i,y-1}),$$

by the definition of the conditional expectation, with $\mathcal{F}_{i,y-1}$ being the information set of firm i in year $y - 1$. This leaves us with

$$\mathbb{E}_{i,y-1} \sum_{m \in y} (\mathbb{1}_{[XS_{i,m}^e = +1]} v_{im}^+ + \mathbb{1}_{[XS_{i,m}^e = -1]} v_{im}^-) = \psi_i,$$

because the firm's expectation of the shocks ψ_m and ψ_{im} in $y - 1$ for the months of the following year y is 0. Note, that the expectation of ψ_m and ψ_{im} conditional on last year's information is 0, as firms cannot predict shocks. Mathematically, the random shocks ψ_m and ψ_{im} are by definition mean independent of the firm's information set in $y - 1$. Note that $\mathbb{E}\psi_m = \mathbb{E}\psi_{im} = 0$, because $\mathbb{E}v_{im}^+ = \mathbb{E}v_{im}^- = 0$ from the structure of equations (2). Therefore,

$$\xi_{iy} = \frac{\psi_i}{1 - \gamma_1 P_{iy} - \gamma_2 N_{iy}},$$

where the firm fixed effect is a source of endogeneity, but we control for it at a later step. The error term ξ_{iy} also indicates that any potential serial correlation in the monthly errors v_{im}^+ and v_{im}^- that is not the result of the firm-specific unobserved heterogeneity is eliminated and is not of concern.

This completes the derivation of equation (3) in Section 3.1.

Appendix B. Supplementary material

The online appendix related to this article can be found online at <https://doi.org/10.1016/j.jedc.2024.104929>.

References

- Anderson Jr, O., 1952. The business test of the IFO-institute for economic research, Munich, and its theoretical model. *Rev. Inst. Int. Stat.* 20 (1), 1–17.
- Awano, G., Bloom, N., Dolby, T., Mizen, P., Riley, R., Senga, T., Van, J., Reenen, J.V., Wales, P., 2018. A firm-level perspective on micro- and macro-level uncertainty. ESCoE Discussion Paper 10.
- Bachmann, R., Elstner, S., 2015. Firm optimism and pessimism. *Eur. Econ. Rev.* 79, 297–325.
- Bachmann, R., Elstner, S., Sims, E.R., 2013. Uncertainty and economic activity: evidence from business survey data. *Am. Econ. J. Macroecon.* 5 (2), 217–249.
- Bachmann, R., Zorn, P., 2020. What drives aggregate investment? Evidence from German survey data. *J. Econ. Dyn. Control* 115, 103873.
- Bartelsman, E.J., Caballero, R.J., Lyons, R.K., 1994. Customer- and supplier-driven externalities. *Am. Econ. Rev.* 84 (4), 1075–1084.
- Bloom, N., Bunn, P., Chen, S., Mizen, P., Smietanka, P., Thwaites, G., 2019. The impact of brexit on uk firms. National Bureau of Economic Research Working Paper.
- Bloom, N., Kawakubo, T., Meng, C., Mizen, P., Riley, R., Senga, T., Van Reenen, J., 2021. Do well managed firms make better forecasts? Technical report. National Bureau of Economic Research.
- Born, B., Enders, Z., Müller, G.J., Niemann, K., 2023. Firm expectations about production and prices: facts, determinants, and effects. In: *Handbook of Economic Expectations*. Elsevier, pp. 355–383.
- Coibion, O., Gorodnichenko, Y., Kumar, S., 2018. How do firms form their expectations? New survey evidence. *Am. Econ. Rev.* 108 (9), 2671–2713.
- Davidson, R., MacKinnon, J.G., 2004. *Econometric Theory and Methods*, vol. 5. Oxford University Press, New York.
- de Leeuw, F., McKelvey, M.J., 1981. Price expectations of business firms. *Brookings Pap. Econ. Act.* 1981 (1), 299–314.
- DGECFIN, 2017. The joint harmonised EU programme of business and consumer surveys: User guide. European Commission: Directorate-General for Economic and Financial Affairs.
- Enders, Z., Hünnekes, F., Müller, G.J., 2019. Monetary policy announcements and expectations: evidence from German firms. *J. Monet. Econ.* 108, 45–63.
- Kosova, R., 2010. Do foreign firms crowd out domestic firms? Evidence from the Czech Republic. *Rev. Econ. Stat.* 92 (4), 861–881.
- Manski, C.F., 2004. Measuring expectations. *Econometrica* 72 (5), 1329–1376.
- Mokinski, F., Sheng, X.S., Yang, J., 2015. Measuring disagreement in qualitative expectations. *J. Forecast.* 34 (5), 405–426.
- Mundlak, Y., 1978. On the pooling of time series and cross section data. *Econometrica*, 69–85.
- Nerlove, M., 1983. Expectations, plans, and realizations in theory and practice. *Econometrica* 51 (5), 1251–1279.
- Pesaran, M.H., 1987. *The Limits to Rational Expectations*. Blackwell Publishers.
- Pesaran, M.H., Weale, M., 2006. Survey expectations. In: *Handbook of Economic Forecasting*, vol. 1, pp. 715–776.
- Semykina, A., Wooldridge, J.M., 2010. Estimating panel data models in the presence of endogeneity and selection. *J. Econom.* 157 (2), 375–380.
- Smith, J., McAleer, M., 1995. Alternative procedures for converting qualitative response data to quantitative expectations: an application to Australian manufacturing. *J. Appl. Econom.* 10 (2), 165–185.
- Tanaka, M., Bloom, N., David, J.M., Koga, M., 2020. Firm performance and macro forecast accuracy. *J. Monet. Econ.* 114, 26–41.
- Theil, H., 1952. On the time shape of economic microvariables and the Munich business test. *Rev. Inst. Int. Stat.* 20 (2), 105–120.
- Triguero, A., Córcoles, D., 2013. Understanding innovation: an analysis of persistence for Spanish manufacturing firms. *Res. Policy* 42 (2), 340–352.
- Vellekoop, N., Wiederholt, M., 2019. Inflation expectations and choices of households. SAFE Working Paper.