



Firms' expectations and uncertainty during COVID-19: Evidence from Italian firms^{☆,☆☆}

Department of Economics, Ca' Foscari University of Venice, Venice, Italy
Venice Centre for Economic and Risk Analytics (VERA), Italy

0167-2681/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Firms face uncertainty in their decision-making processes as they deal with incomplete information, a consequence of imperfect knowledge well documented in the literature (Duncan, 1972; Baley and Blanco, 2019). This might originate from the absence of causal-link information, limited predictive accuracy, and the inability to assign a prior probability to future outcomes, known as Knightian uncertainty (Knight, 1921; Milliken, 1987). Even when the general long-term direction of certain future events is predictable, uncertainty makes it difficult for economic agents to form expectations about how these events will progress, when they will occur, and what their impact will be.

Economic uncertainty is a documented driver of the business cycle, a predictor of recessions (Jurado et al., 2015; Ercolani and Natoli, 2020), and a source of volatility comovement in financial markets (Engle and Campos-Martins, 2023). It rises sharply after major shocks such as policy and political events (Baker et al., 2016; Kellard et al., 2022), geopolitical conflicts (Caldara and Iacoviello, 2022), climate-related disasters (Tzouvanas et al., 2019), and pandemics (Baker et al., 2020). The economic consequences of such shocks, particularly health shocks, raise the question of how firms respond to environments in which the future is hard to forecast. Bloom (2009) provides the canonical framework, showing that elevated uncertainty depresses hiring and investments as firms postpone irreversible decisions, and that activity rebounds once uncertainty subsides.

The COVID-19 pandemic represents a particular case of induced economic uncertainty. Altig et al. (2020) document that most economic uncertainty measures in the US and UK reached all-time high levels during this period, leading to a sharp drop in industrial production in the range between 12% and 19%.

To frame these economic effects, it is useful to recall the timeline of the outbreak: the SARS-CoV-2 virus emerged in Wuhan, China, in December 2019, spread rapidly across borders, and was officially declared a pandemic on March 11, 2020. Among European countries, Italy was particularly affected early, quickly becoming an epicenter of the crisis. The Italian government declared a state of emergency on January 31, 2020. As the virus continued to spread domestically, a nationwide lockdown was imposed on March 9, 2020, restricting movement, gatherings, and non-essential business activities. These measures were progressively relaxed starting in May 2020, with alternating periods of stricter and looser controls in response to evolving epidemiological dynamics. The state of emergency officially ended on March 31, 2022.¹ Concurrently, economic activity largely rebounded to the levels observed before the pandemic (Commission, 2022).

This paper investigates how the COVID-19 pandemic altered the relationship between Italian firm expectations and uncertainty. We examine four aspects of firms' expectations and ask whether the pandemic produced a departure from what was observed in other periods. Specifically, we examine (i) how current performance shapes firms' growth expectations, (ii) whether these expectations are informative about future outcomes, (iii) how self-reported subjective uncertainty relates to forecast accuracy, and (iv) how cross-sectional dispersion in firm expectations relates to future sectoral volatility. The analysis is also extended to the price domain to test whether the relationships identified for real growth hold in the price dimension.

Although the macroeconomic and firm-level repercussions of the COVID-19 shock have been documented (e.g., Carletti et al., 2020; Buchetti et al., 2022; Guerrieri et al., 2022), little is known about how the pandemic reshaped firms' expectation-formation processes. In this regard, Italy offers a compelling case study for this investigation, as it was among the first nations affected by the new coronavirus and the first Western country to introduce movement-restriction policies (Cerqueti et al., 2022), creating an unprecedented environment for firms.

Furthermore, given its position as the eighth country globally in manufacturing value added and the second in the Eurozone, insights into how Italian firms adapted their expectations under the pandemic's stress offer an interesting perspective on the role of uncertainty in shaping economic outlooks. The severity of COVID-19's economic impact on Italy is evidenced by its 2020 GDP contraction, which ranked among the largest in the Eurozone after Spain and Greece.² This health shock thus provides a unique opportunity to investigate how firms adapt their expectations and decision-making processes in response to extreme uncertainty.

We address these questions using firm-level data from the Bank of Italy's Survey on Industrial and Service Firms (INVIND), an annual panel that includes direct measures of growth expectations, realized turnover, firm characteristics, and perceived uncertainty. Our empirical strategy allows the link between current and expected variables to vary year by year over the full sample, which spans nearly three decades, in order to reveal whether the pandemic period departs from the dynamics observed in earlier years, without imposing an ex post sub-period design.

Results show that, in the real-growth domain, the link between current growth and one-year-ahead growth expectations is negative and stable across the full sample. Over nearly three decades, including the 2008–2009 financial crisis, the 2011–2012 sovereign debt crisis, and the COVID-19 pandemic, we do not detect any significant change in this link.

Italian firms continue to expect a partial reversal of current performance regardless of the macroeconomic regime. Growth forecasts also predict realized growth in every year of the sample, with predictive content that persists through the pandemic. The relationship weakens in 2021 and 2022 but never approaches zero. Firms' expectations retain genuine information about future outcomes even during the most acute phase of the pandemic. The strongest predictive relationship in the sample appears in the 2019 wave, fielded between February and May 2020 as the first national lockdown unfolded. This case reflects nowcasting of the ongoing event rather than genuine pre-shock forecasting.

¹ *Delibera del Consiglio dei Ministri* No. 20A00737 (Council of Ministers Resolution, 31 January 2020) declared the national state of emergency; *Decreto del Presidente del Consiglio dei Ministri* (DPCM) No. 20A02179, 9 March 2020, imposed a nationwide lockdown restricting movement and non-essential business activities; DPCM No. 20G00051, 26 April 2020, set out the phased reopening that began in May 2020; DPCM No. 22G00034, 24 March 2022, formally terminated the state of emergency.

² Data retrieved from the World Bank for Manufacturing value added (mnemonics NV.IND.MANF.CD) and GDP growth (mnemonic NY.GDP.MKTP.KD.ZG).

By contrast, the link between firms' self-reported uncertainty and the size of their forecast errors is not constant over time. It is strong in years of macroeconomic stress, close to zero in calmer years, and weak during the most acute phase of the pandemic, when firms faced a common exogenous shock and disagreed less about how uncertain they were.

The picture changes when the analysis is extended to the price domain. The response of expected prices to realized prices is positive on average but unstable over time. It is weak in the 1990s and breaks down during the 2008–2009 financial crisis and again during the post-pandemic inflation surge.

The link between subjective uncertainty and price forecast errors, by contrast, is positive throughout the available window and shows much less time variation than the real-growth case, with only one of six year-specific interactions individually significant. The asymmetry between stable growth-expectation formation and time-varying price-expectation formation runs through the analysis.

At the aggregate level, the cross-section of firms' expectations carries information about the cross-section of subsequent outcomes. Higher disagreement among firms in one year is followed by higher dispersion in their realized growth and forecast errors the next. The relationship holds in both the real-growth and price domains and at several levels of aggregation, and it does not stem from the autoregressive persistence of dispersion or from firms extrapolating recent volatility into their forecasts.

At the firm level, exploiting the dedicated COVID-19 module of the INVIND 2020 wave, we identify workforce shortages as the dominant channel through which the pandemic affected expectations. Firms that interrupted production due to workforce unavailability altered both how they formed growth expectations and the predictive content of their forecasts, more strongly than firms affected by regulatory closures (which influenced expectation formation alone), and unlike firms exposed through demand, supply, or financial channels.

The remainder of this paper is structured as follows. Section 2 positions our contribution within the existing literature. Section 3 introduces the INVIND survey and the variables used in the analysis. Section 4 presents our hypotheses and empirical strategy. Section 5 reports and discusses the results. Finally, Section 6 concludes.

2. Related literature

Prior research highlights the role of firm expectations and uncertainty in driving strategic economic decisions such as investment, hiring, and production (Bloom et al., 2007; Bachmann et al., 2013; Enders et al., 2022). Heterogeneity in expectations is a well-known feature that potentially arises from alternative information sources and interpretation abilities (Coibion et al., 2018). Studies also indicate that firms' past performance and local information influence their expectations about their future and the aggregate economy (Massenot and Pettinicchi, 2018; Altig et al., 2022; Enders et al., 2022; Andrade et al., 2022; Dovern et al., 2023), suggesting in some cases, a degree of "rational inattention" to aggregate signals unless they are particularly salient (Lucas, Jr., 1977; Maćkowiak and Wiederholt, 2009). Furthermore, a line of research shows that higher firm-level uncertainty is typically associated with lower predictive accuracy (Altig et al., 2022; Bachmann et al., 2013; Bloom et al., 2020), and that uncertainty consistently intensifies during recessions at both macroeconomic and firm levels, acting as both a trigger for negative shocks and an amplifier of economic slowdowns (Bloom, 2014).

Building on this literature, this study contributes to research on firm expectations and uncertainty, the real effects of uncertainty shocks, and the specific impact of the COVID-19 pandemic on firms. By providing micro-level evidence from Italy, we document how a severe health shock might affect firm-level forecasting behavior and the aggregate signals contained in expectations.

As firm-level uncertainty is not directly observable, researchers commonly rely on survey data to obtain proxies for it. One approach uses subjective probability distributions over future outcomes to construct measures of expectations (the first moment) and uncertainty (the second moment). Alternatively, surveys may ask respondents for a point forecast alongside the minimum and maximum plausible realizations, with the point estimate capturing expectations and the reported range providing a proxy for uncertainty (e.g., Coibion et al., 2018; Altig et al., 2022).

Examples of firm-level surveys include the Survey of Business Uncertainty (SBU) in the US (Altig et al., 2022), the Decision Maker Panel (DMP) in the UK (Bunn et al., 2024), the IFO Business Survey in Germany (Lehmann, 2023), and the Japanese Annual Survey of Corporate Behavior (ASCB) in Japan (Lewin et al., 1995). Studies using this type of survey show that firm expectations can be more predictive of their outcomes than historical data (Bloom et al., 2020), and that subjective uncertainty increased during the COVID-19 pandemic (Altig et al., 2020; Chen et al., 2021).

Recent work has further refined the understanding of expectation formation using these surveys. For instance, Tanaka et al. (2020) combined firm accounting data with GDP growth forecasts from the ASCB, showing that forecast accuracy is associated, among others, with firm size and age. Using the German IFO survey, Buchheim et al. (2022) analyzed sentiment during the pandemic and showed that variation in expected shutdown duration primarily reflects sentiment.

Other surveys have been used to study uncertainty. Altig et al. (2022), using the SBU, show that firm-level expectations predict actual growth rates and that their perceived uncertainty is related to forecast errors and revisions. Following the SBU, Chen et al. (2021) utilized the Research Institute of Economy, Trade and Industry (RIETI)'s Business Plans and Expectations Survey (BPES) and showed that during the COVID-19 pandemic, there was an increase in subjective uncertainty for Japanese firms, especially for firms with direct trade and investment ties to China.

The current paper leverages the Bank of Italy's INVIND survey, a long-running panel that enables the study of expectations over time. Prior research using INVIND has established that demand uncertainty negatively affects investment (Guiso and Parigi, 1999); however, this effect may be mitigated by factors such as market competition and labor flexibility (Bontempi et al., 2010). Fiori and Scoccianti (2023) developed a novel measure from INVIND data to show that firm-level uncertainty is persistent and primarily affects

outcomes through its downside risk. Bottone (2025) demonstrates the superiority of point forecasts over the mean of probability forecasts for predicting future sales and that alternative indicators of uncertainty provide different rankings among firms.

While the existing literature establishes the link between survey-based uncertainty and economic outcomes, our contribution is distinct. We use the COVID-19 pandemic to identify structural changes in the expectation formation process itself. We investigate whether the established relationships between past performance, firm characteristics, and future expectations remained stable across the pandemic shock, and analyze how the dispersion of micro-level expectations translated into aggregate volatility. This sheds light on how forecasting behavior responds to high uncertainty.

3. Data

The Bank of Italy's INVIND survey collects a wide range of firm-level data, including prices, sales, employment, investment, financial conditions, and turnover trends. Established in 1984, the panel initially focused on manufacturing firms with 50 or more employees, its coverage subsequently expanding to all of industry in the strict sense (1999), private non-financial services (2002), and construction (2006). The smallest size class, firms with 20–49 employees, was incorporated in 2001 for industry and progressively for the other sectors.³ Despite these expansions, the survey maintains a specific focus, excluding sectors such as finance, public administration, education, and healthcare. To ensure a representative sample of the Italian business landscape, the INVIND survey employs a stratified random sampling methodology, drawing on diverse data sources, including INPS archives and Infocamere.

The survey uses a panel design, with firms usually interviewed annually as long as they remain within the target population. Exclusions may occur due to cessation of operations, mergers, or changes in size or activity classification. Non-participating firms are replaced with similar entities to maintain sample representativeness. Interviews, conducted usually between February and May by the Bank's branches, focus on the prior year's activities as the reference year t .⁴ This timing is particularly interesting, as it means expectations for the upcoming year are captured within the first semester of that year.

The questionnaire comprises two sections: a fixed component and a variable component. The fixed component gathers general information about the firm, its structure, and annual data on key metrics, including investment, employment, turnover, financial results, capacity utilization, and debt. The variable component, conversely, covers different themes each year, adapting to evolving economic landscapes and research priorities. Additionally, the INVIND 2020 wave (fielded in spring 2021) included a dedicated COVID-19 module that recorded firm-level pandemic exposure.

In our analysis, we consider all firms with at least 50 employees that are observable in the INVIND panel between 1996 and 2024.⁵ We retain all sectors covered by the survey, which span manufacturing (SS1–SS7), energy and utilities (SS8), construction (SS9), transport and communications (SS10), and wholesale and retail trade (SS11). The resulting analysis sample contains 68,003 firm-year observations. A detailed breakdown of the sample by survey year, by sector, and by the cross-tabulation of size class and macro-area is reported in Appendix A of the Online Appendix.

Table 1 presents the variables selected for the empirical analysis, detailing their names from the INVIND dataset and corresponding descriptions. These data are instrumental for our investigation, as described in the next section, and encompass both firm-level characteristics and key financial metrics.

Firm-level characteristics include `ident`, which serves as a unique firm code, and `annoril`, indicating the year of the survey. We also incorporate `v24`, representing average employment in year t as a proxy for firm size (see, for example, Bachmann and Elstner, 2015; Coibion et al., 2018), the logarithm of `annoril` minus the foundation year (variable `v284`) as the firm's age (Anderson and Reeb, 2003), and `settor11`, the firm's business sector.

Regarding financial metrics, the dataset provides `v209` and `v210`, which denote turnover from sales of goods and services in current prices for year $t - 1$ and year t , respectively. Additionally, `v220a` measures the average annual percentage change in billed prices for year t relative to year $t - 1$.

For firm expectations, we use `v540`, which records the expected turnover percentage change for year $t + 1$ net of price effects. The minimum and maximum projected changes in turnover for year $t + 1$ are recorded in `v541` and `v542`, which we use to construct a measure of firm-level subjective uncertainty. These variables are collected in every wave from 1996 to 2024, with one exception in the 2005 wave, which did not include any of the three.⁶ The 2005 firm-year cells for these variables are therefore missing in the archive, with downstream implications for any specification that uses their lagged values.

In addition to the real turnover domain, the INVIND survey contains analogous variables for price expectations. Variable `v440` records the expected average percentage change in billed prices for year $t + 1$ relative to year t , while `minv440` and `maxv440` capture the minimum and maximum of this price-change expectation, respectively. Together, these allow us to construct a price-uncertainty proxy that parallels the turnover-uncertainty measure introduced in Section 4.

To capture firm-level policy exposure during the pandemic, we also incorporate hours of *Cassa Integrazione Guadagni* (CIG, short-time work), recorded as `v29`, which we transform to $\log(1 + \text{CIG hours}_t)$ to reduce right-skewness while retaining zeros. CIG is the main Italian government-subsidized reduced-hours scheme, and its use signals that a firm experienced a temporary demand or supply shock but retained its workforce, which may systematically influence expectation formation.

³ See Banca d'Italia (2024) for the full coverage history.

⁴ See the user documentation for the dataset in Banca d'Italia (2024), section "Rilevazione dei dati e riferimenti normativi".

⁵ The 1996 starting point reflects the first wave in which the survey collected firms' point expectations and expectation-range bounds for revenue growth.

⁶ Variable-coverage information reported throughout this section is drawn from the official INVIND codebook INVIND_IT_03_lista_variabili.xlsx maintained by the Bank of Italy, which documents for each variable the survey waves in which it was collected and the corresponding number of responses. The file is publicly available at https://www.bancaditalia.it/statistiche/basi-dati/rdc/accessomicrodati/INVIND_IT_03_lista_variabili.zip.

Table 1

Selected variables from the INVIND survey used in the empirical analysis, with corresponding descriptions and coverage.

Variable	Description	Domain	Coverage
<i>Panel identifiers and controls</i>			
ident	Firm code		1996–2024
annoril	Year of survey		1996–2024
v24	Average employment (year t)	Size	1996–2024
settor11	Business sector (11 categories)	Sector	1996–2024
v284	Year of establishment of the firm	Age	1996–2024
areag4	Geographic macro-area (NW, NE, C, SI)	Region	1996–2024
peso	Survey weight (inverse probability)	Weight	1996–2024
<i>Realized outcomes (year t)</i>			
v209	Turnover, sales of goods and services, current prices (year $t - 1$)	Nominal turnover	1996–2024
v210	Turnover, sales of goods and services, current prices (year t)	Nominal turnover	1996–2024
v220a	Average annual percentage change in billed prices ($t/t-1$)	Prices	1996–2024
<i>Expectations and uncertainty (formed at t for $t + 1$)</i>			
v540	Expected turnover percentage change, net of price effects ($t+1/t$)	Real turnover	1996–2024 ^a
v541	Minimum projected change in turnover (year $t + 1$)	Real turnover	1996–2024 ^a
v542	Maximum projected change in turnover (year $t + 1$)	Real turnover	1996–2024 ^a
v440	Expected average percentage change in billed prices ($t+1/t$)	Prices	1996–2024
minv440	Minimum expected price change ($t+1/t$)	Prices	2014–2020
maxv440	Maximum expected price change ($t+1/t$)	Prices	2014–2020
<i>Policy exposure and labor-market variables</i>			
v29	Total CIG hours (year t)	Policy	1996–2024
<i>COVID-19 module (INVIND 2020 wave, fielded spring 2021)</i>			
cvd1	Dummy: production stopped due to COVID (0/1)	COVID	2020 wave only
cvd3	Share of production suspended (ordinal, 1–4)	COVID	2020 wave only
cvd4a	Main reason for interruption (categorical, 1–5)	COVID	2020 wave only
cvd5	Perceived probability of permanent closure (ordinal, 1–5)	COVID	2020 wave only
mor20a	Moratorium request (categorical, 1–3)	COVID	2020 wave only

^a The 2005 wave of INVIND did not collect v540, v541, or v542. The corresponding firm-year cells are missing in the archive.

The COVID-19 module deserves a note on timing. Under the INVIND convention, the 2020 wave was administered in spring 2021, when firms reported on activities in reference year $t = 2020$. We extract each firm's COVID exposure record from the 2020 wave and merge it (by firm identifier) into the 2020–2024 subpanel. This design provides a single, predetermined firm-level measure of pandemic disruption that can be interacted with the main regressors to identify heterogeneous effects.

Before turning to the analysis of firm expectations and uncertainty, we assess whether sample attrition undermines the representativeness of the panel and, hence, the potential validity of the results. Attrition bias occurs when the observations that remain in a study over time (the “survivors”) do not represent a random sample of the original data, potentially leading to biased results. Such bias may occur because firms exit the survey for reasons that are correlated with the variables of interest.

In firm-level studies, for example, attrition bias can emerge when a firm's decision to stay in the market depends on its productivity or capital intensity. As Olley and Pakes (1996) note, firms with larger capital stocks may survive even during periods of low productivity because capital acts as a buffer against shocks, whereas highly productive firms tend to persist regardless of their capital. Conversely, highly productive firms may be less likely to exit, regardless of their capital. Observing only surviving firms can therefore yield a spurious negative relationship between capital and productivity, leading to erroneous conclusions.

Appendix B reports an attrition analysis of the INVIND dataset for firms surveyed in at least two consecutive years. Consistent with Di et al. (2016) and Notarnicola (2018), we find that attrition rates are relatively stable across sectors and firm sizes, despite some regional variation, suggesting limited risk of systematic bias.

4. Framework and hypothesis development

A growing body of evidence indicates that firms' growth expectations, together with the uncertainty surrounding those expectations, influence firms' decisions, including investment (Bloom et al., 2007), hiring (Bachmann et al., 2013), and production (Enders et al., 2022). Furthermore, heterogeneity in expectations may arise from differences in information sources and firms' interpretations of the economic environment (Coibion et al., 2018).

Major exogenous shocks offer opportunities to study how firms form expectations and manage the associated uncertainty under stress. The COVID-19 pandemic, unprecedented in its rapid global spread and economic disruption, provides a unique setting to investigate whether, and in which ways, such extreme conditions have altered expectations and uncertainty both at the individual-firm level and in the aggregate. We examine the potential impacts of this event on firm behavior and expectations within the Italian context and exploit the full panel (1996–2024) to measure whether the relationship between expectations and outcomes changed.

Specifically, we interact the variable of interest with a full set of year dummies, allowing the slope coefficient to vary across years. This year-varying approach uses all available observations in a single estimation, nests the constant-slope specification as a testable restriction, and reveals the time profile of any structural change.

To characterize a firm's expectations, we first describe the quantity used as their measure. The log-change in nominal turnover for firm i at time t is

$$s_{it} = 100[\ln S_{it} - \ln S_{it-1}], \quad (1)$$

where S_{it} and S_{it-1} denote nominal turnover from sales of goods and services (INVIND variables v210 and v209, respectively), $i = 1, \dots, N$, $t = 1, \dots, T$. The real turnover growth for firm i at horizon t is defined as

$$g_{it} = s_{it} - \pi_{it}, \quad (2)$$

where $\pi_{it} = 100 \ln(1 + p_{it}/100)$ is the log-transformed price change and p_{it} is the average annual percentage change in billed prices (variable v220a).

The expectation of firm i at time t for the real turnover growth at horizon $t + 1$, $E_t[g_{i,t+1}]$, is recorded in the INVIND variable v540, which reports the expected percentage change in turnover net of price effects. The survey also provides the expected year-on-year change in billed prices, $E_t[p_{i,t+1}]$ (variable v440), which we use as the dependent variable in the price-expectation analysis. Lead-lag alignment within each firm is enforced to use strictly consecutive survey years only: g_{t+1} is assigned to period t only when the firm is observed in both year t and year $t + 1$, to prevent misaligned "next observed" comparisons.⁷

In this respect, we begin outlining our research framework by introducing the first hypothesis.

H1: Firms' realized growth influences their expectations of future growth, and this relationship is stable over time.

We test whether firm-specific realized growth has a significant effect on one-year-ahead expectations, whether the sign and magnitude of this relationship vary across years, and whether the COVID-19 period represents a departure from the historical pattern.

Recent survey studies have revealed systematic patterns in how firms form and revise their expectations. Using the IFO Business Survey of German manufacturing firms, Massenot and Pettinicchi (2018) document that firms extrapolate from their qualitative self-assessment of current business conditions when forming expectations about their future business situation, generating predictable expectation errors over time. Coibion et al. (2018) show that firms devote limited resources to gathering information and adjust their macroeconomic expectations only when salient news arrives. Altig et al. (2022) present evidence from the US SBU indicating that firms' subjective sales growth expectations predict realized growth. Enders et al. (2022) find that German firms with positive (negative) production expectations are more likely to increase (decrease) their current output and prices than are firms that anticipate no change. Andrade et al. (2022) show that French manufacturing firms treat industry-specific shocks as signals of the macroeconomy, persistently adjusting both their aggregate and own-price expectations in ways that vary with firm and industry characteristics, as outlined by the island model in Lucas, Jr. (1972). Dovern et al. (2023) provide further evidence that local information shapes firms' expectations about aggregate growth in line with the model of rational inattention outlined in Maćkowiak and Wiederholt (2009). Collectively, the evidence from previous studies implies that past or surrounding performance often enters the expectation-formation process, motivating us to address our first hypothesis.

To explore this empirically, we estimate the following model of expected real turnover growth:

$$E_t[g_{i,t+1}] = \alpha_i + \beta \cdot g_{it} + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (3)$$

where g_{it} is the firm's realized growth in t , size_{it} and age_{it} control for a firm's scale and maturity,⁸ α_i denotes firm fixed effects, d_t denotes year fixed effects, and v_{it} the error term.

The coefficient on past performance, β , describes the role of momentum in expectation formation. A positive β indicates that higher (or lower) current turnover growth is associated with higher (or lower) expected future growth. This reflects persistence, suggesting that firms expect past performance trends to continue. In contrast, a β close to zero implies that past growth has no measurable influence on expectations, indicating that firms form forecasts independently of their past performance. A negative β , on the other hand, signals antipersistence, indicating that firms believe current growth will reverse in the subsequent period and thus anticipate a turning point in performance.

To assess whether the expectation-formation process has changed over time, we extend Eq. (3) by interacting g_{it} with a complete set of year dummies:

$$E_t[g_{i,t+1}] = \alpha_i + \beta \cdot g_{it} + \sum_{\tau \neq 2018} \delta_\tau \cdot (g_{it} \times D_\tau) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (4)$$

where $D_\tau = \mathbf{1}(t = \tau)$ is a year indicator, d_t denotes year fixed effects, and β is the slope on g_{it} in the reference year 2018. We select 2018 as the reference year because it is the last calendar year whose INVIND wave (fielded February–May 2019) was entirely unaffected by the COVID-19 pandemic. The 2019 wave, by contrast, was fielded during the initial lockdowns (February–May 2020) and therefore incorporates pandemic-period information.

The coefficient δ_τ measures the deviation of the year- τ slope from this baseline, and the total slope in year τ is $\beta_\tau = \beta + \delta_\tau$. If the expectation-formation process is stable, all δ_τ should be individually insignificant. If the COVID-19 period altered this process, the interactions for the corresponding years should be significant.⁹

⁷ For example, if a firm is observed in 2017 and 2019 but not 2018, the 2019 realization is not treated as one-year-ahead of 2017.

⁸ We use $\log(1 + \text{age})$ to avoid the undefined $\log 0$ for newly founded firms.

⁹ As a complementary exercise, all panel specifications are estimated year by year via cross-sectional OLS with sector fixed effects and heteroskedasticity-robust standard errors, using INVIND survey weights. The full set of cross-sectional estimates is reported in Appendix C.

To interpret the sign of β in Eq. (4), we complement the first research hypothesis with a persistence diagnostic of realized growth. We estimate the autoregressive specification

$$g_{it} = \alpha_i + \rho \cdot g_{i,t-1} + \sum_{\tau \neq 2018} \rho_{\tau} \cdot (g_{i,t-1} \times D_{\tau}) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (5)$$

where ρ is the persistence of realized growth in the reference year and $\rho + \rho_{\tau}$ is the total persistence in year τ . If $\rho < 0$, growth is mean-reverting, implying that a good year tends to be followed by a weaker one. In that case, a negative β in Eq. (4) indicates that firms anticipate this reversal rather than extrapolating current growth into the future.

Then, we examine the second hypothesis.

H2: Firms' expectations contain predictive information about future growth.

We investigate whether expected growth at t predicts realized growth at $t + 1$, beyond what is captured by current growth, and whether the predictive content varies over time, including the COVID-19 period.

Building on the view that firms draw on both aggregate, industry-level trends and firm-specific data, the literature provides evidence that firm managers embed forward-looking information into their growth forecasts (see, e.g., Massenot and Pettinicchi, 2018; Avalos et al., 2024). If firms use all available signals, we should observe a systematic relationship between their forecasts and subsequent realized growth. Following Gennaioli et al. (2016), we test this directly by regressing realized outcomes on stated expectations. Since expectations are formed at t and outcomes are realized at $t + 1$, the timing ensures that β captures genuine predictive content. Hence, we estimate the following model:

$$g_{i,t+1} = \alpha_i + \beta \cdot E_t[g_{i,t+1}] + \phi \cdot g_{it} + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_{t+1} + v_{it}, \quad (6)$$

where $g_{i,t+1}$ is realized growth in $t + 1$, $E_t[g_{i,t+1}]$ is the expectation formed at t , and g_{it} controls for the autoregressive component. The fixed effects d_{t+1} are indexed by the outcome year, and α_i absorbs time-invariant firm heterogeneity. A positive β indicates that expectations carry predictive information beyond current growth. The inclusion of g_{it} ensures that β captures incremental forward-looking content rather than mechanical correlation with lagged growth.

As with the first hypothesis, we extend the model with year-varying slopes:

$$g_{i,t+1} = \alpha_i + \beta \cdot E_t[g_{i,t+1}] + \sum_{\tau \neq 2018} \delta_{\tau} \cdot (E_t[g_{i,t+1}] \times D_{\tau}) + \phi \cdot g_{it} + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_{t+1} + v_{it}, \quad (7)$$

where the year-varying slopes on $E_t[g_{i,t+1}]$ are indexed by the formation year t , measuring whether the informational content of expectations changed depending on when they were formed.¹⁰

We then transition to the third hypothesis.

H3: Firms' self-reported uncertainty is informative about their forecast accuracy.

We examine whether the range of expected outcomes reported by firms at t predicts the magnitude of their forecast errors realized at $t + 1$, and whether this relationship varies over the business cycle, including the COVID-19 period.

Empirical evidence suggests that higher subjective uncertainty is linked to lower predictive accuracy (see, for example, Altig et al., 2022; Bachmann et al., 2013; Bloom et al., 2020; Diether et al., 2002; Epstein and Wang, 2004). We quantify firm-level uncertainty using a self-reported dispersion measure, defined as

$$u_{i,t+1|t} = E_t[g_{i,t+1}^{\max}] - E_t[g_{i,t+1}^{\min}], \quad (8)$$

where $E_t[g_{i,t+1}^{\max}]$ (variable v542) and $E_t[g_{i,t+1}^{\min}]$ (variable v541) are the maximum and minimum real sales growth rates anticipated by the firm given in t , respectively.

To measure firms' forecast accuracy, we define the absolute forecast error as

$$\text{AFE}_{it} = |g_{it} - E_{t-1}[g_{it}]|, \quad (9)$$

where g_{it} is the firm's realized real sales growth in period t and $E_{t-1}[g_{it}]$ is the prediction formed at $t - 1$.

We model the relationship between lagged uncertainty and realized forecast errors as

$$\text{AFE}_{it} = \alpha_i + \lambda \cdot u_{it|t-1} + \sum_{\tau \neq 2018} \delta_{\tau} \cdot (u_{it|t-1} \times D_{\tau}) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (10)$$

where $u_{it|t-1}$ is the firm's self-reported uncertainty about growth in t as forecasted at $t - 1$, and λ is the slope in the reference year. A positive λ indicates that firms reporting wider ex ante forecast ranges also experience larger ex post absolute errors. If the relationship is state-dependent, we expect δ_{τ} to be positive and significant during episodes of economic turmoil.

The same year-varying design extends to the *price domain*, where the INVIND survey collects analogous expectation and uncertainty variables. Andrade et al. (2022) show that French manufacturing firms use industry-specific conditions to update their own-price expectations, consistent with the island model of Lucas, Jr. (1972), while Enders et al. (2022) show that German firms'

¹⁰ Each firm-year observation has a unique formation year t and a unique outcome year $t + 1$. Using formation-year fixed effects in place of outcome-year fixed effects produces identical fitted values and slope coefficients.

production expectations are closely linked to their current price-setting decisions. These findings suggest that price-expectation formation may follow a different pattern from quantity-expectation formation, motivating a separate investigation.

We examine price-expectation formation by estimating

$$E_t[\pi_{i,t+1}] = \alpha_i + \beta \cdot \pi_{it} + \sum_{\tau \neq 2018} \delta_\tau \cdot (\pi_{it} \times D_\tau) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (11)$$

where π_{it} is the realized price change (variable v220a) and $E_t[\pi_{i,t+1}]$ is the expected price change (variable v440). This tests whether the relationship between realized and expected price changes is stable over time or varies significantly across periods.

We then assess whether price uncertainty is informative about price-forecast accuracy. We construct price uncertainty as

$$u_{i,t+1|t}^\pi = E_t[\pi_{i,t+1}^{\max}] - E_t[\pi_{i,t+1}^{\min}], \quad (12)$$

where $E_t[\pi_{i,t+1}^{\max}]$ (variable maxv440) and $E_t[\pi_{i,t+1}^{\min}]$ (variable minv440) are the maximum and minimum firm-level price-change expectations. As discussed in Section 3, these bounds are available only for the 2014–2020 waves, restricting the estimation sample for this specification to 2015–2021. We estimate

$$\text{AFE}_{it}^\pi = \alpha_i + \lambda \cdot u_{it|t-1}^\pi + \sum_{\tau \neq 2018} \delta_\tau \cdot (u_{it|t-1}^\pi \times D_\tau) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (13)$$

where $\text{AFE}_{it}^\pi = |\pi_{it} - E_{t-1}[\pi_{it}]|$ is the absolute price forecast error.

Lastly, we examine the fourth hypothesis.

H4: Cross-sectional dispersion in firms' expectations predicts future economic volatility.

We investigate if larger cross-firm dispersion in expected growth predicts increased group-level volatility in realized growth and forecast errors, and test the direction of predictability through reverse regressions.

Several empirical studies show that cross-sectional dispersion in firm expectations contains valuable information about future macroeconomic conditions. Bloom (2014) observes that economic uncertainty rises during recessions across all data levels, from aggregate macro series to industry, firm, and plant observations. Bachmann et al. (2013) used German business survey data to demonstrate that upward shocks to sector-level business uncertainty cause prolonged downturns in sectoral economic activity. Altig et al. (2022) construct an uncertainty index by taking size-weighted averages of firm-level subjective forecast distributions and show that the index correlates positively with the financial Volatility Index (VIX).

Motivated by these findings, we shift our focus from firm-specific uncertainty to its aggregate manifestation, the dispersion of expectations across a group. Establishing such a link implies that the collective uncertainty perceived by firms can serve as a signal for macroeconomic fluctuations. As a group-level indicator of realized volatility, we aggregate firm-specific growth rates by computing the interquartile range (IQR) of turnover growth within each group j in period t .

Let g_{ijt} denote real turnover growth of firm i that belongs to group j (defined at four aggregation levels: sector×macro-area, sector, macro-area, and overall) at date t , and let $Q_\alpha(x)$ be the α th-percentile values of the cross-firm distribution for the variable x . The aggregated dispersion measure is then

$$\text{IQR}(g_{jt}) = Q_{0.75}(g_{jt}) - Q_{0.25}(g_{jt}),$$

which captures the spread of the central 50 percent of firm-level growth observations within each group.

Following prior literature, we interpret a wider IQR of cross-sectional firm outcomes as an indicator of greater volatility (e.g., Berger and Vavra, 2013). We first examine whether greater dispersion in firms' expectations is followed by greater dispersion in realized outcomes, by regressing the IQR of realized growth on the IQR of one-period-ahead expectations:

$$\text{IQR}(g_{jt}) = \alpha_j + \beta \cdot \text{IQR}(E_{t-1}[g_{jt}]) + d_t + \varepsilon_{jt}, \quad (14)$$

where d_t captures common time effects. A positive β indicates that wider disagreement among firms about future growth precedes higher realized volatility within the group.

A natural benchmark is the persistence of realized volatility itself. If realized dispersion is highly autocorrelated, the predictive content of expectation dispersion in Eq. (14) may simply mirror this persistence. We document the persistence directly,

$$\text{IQR}(g_{jt}) = \alpha_j + \rho \cdot \text{IQR}(g_{j,t-1}) + d_t + \varepsilon_{jt}, \quad (15)$$

and then estimate a joint specification that includes both lagged regressors,

$$\text{IQR}(g_{jt}) = \alpha_j + \beta \cdot \text{IQR}(E_{t-1}[g_{jt}]) + \rho \cdot \text{IQR}(g_{j,t-1}) + d_t + \varepsilon_{jt}. \quad (16)$$

A significant β in Eq. (16) indicates that expectation dispersion carries forward-looking information beyond what is accounted for by the persistence of realized volatility.

If expectation dispersion predicts realized volatility, it may also predict the dispersion of forecast errors, since firms that disagree more about the future are likely to be, on average, less accurate. We test this with a parallel specification in which the IQR of absolute forecast errors is the dependent variable,

$$\text{IQR}(\text{AFE}_{jt}) = \alpha_j + \beta \cdot \text{IQR}(E_{t-1}[g_{jt}]) + d_t + \varepsilon_{jt}. \quad (17)$$

The specifications above establish whether expectations predict outcomes, but do not rule out the possibility that realized volatility drives future expectation dispersion. If firms simply extrapolate from recent turbulence, we would observe a significant relationship in the reverse direction. We test this by estimating

$$\text{IQR}(E_t[g_{j,t+1}]) = \alpha_j + \beta \cdot \text{IQR}(g_{j,t-1}) + d_t + \varepsilon_{jt}. \quad (18)$$

The regressor is lagged by one period relative to the dependent variable, mirroring the lag structure of Eq. (14). If β is significant in Eq. (14) but not in Eq. (18), the predictive direction runs from expectations to outcomes, establishing that disagreement among firms contains forward-looking information rather than reflecting past experience. We estimate all five equations at each of the four aggregation levels, separately for real growth and price changes. We apply the same specifications to inflation expectations, replacing realized growth g_{jt} with the realized price change π_{jt} , expectations $E_t[g_{j,t+1}]$ with $E_t[\pi_{j,t+1}]$, and forecast errors AFE_{jt} with AFE_{jt}^π .

To retain common time effects and measure the impact of the pandemic, we recover an average COVID-19 effect from the estimated year fixed effects. Consistent with the main specifications, we use 2018 as the reference year ($\mu_{2018} = 0$) and define the average pandemic-period shift as

$$\bar{\Delta}_{\text{COV}} = \frac{1}{3} \sum_{t=2020}^{2022} d_t, \quad (19)$$

testing $H_0: \bar{\Delta}_{\text{COV}} = 0$ via a Wald t -test on the linear combination. For completeness, we also report a joint Wald test of $H_0: \mu_{2020} = \mu_{2021} = \mu_{2022} = 0$.

COVID-19 heterogeneity. The year-varying specifications identify whether the expectation–realized outcome relationships changed during the pandemic period. A distinctive feature of the COVID-19 pandemic, relative to earlier macroeconomic crises, is that the disruption operated through identifiable firm-level channels: regulatory closures halted production in non-essential sectors, demand collapsed for some goods while surging for others, workforce availability was reduced by lockdowns, quarantine, and health concerns, and supply chains were disrupted by movement restrictions and logistics bottlenecks (see, e.g., Carletti et al., 2020; Guerrieri et al., 2022).

Unlike the 2008–2009 financial crisis or the 2011–2012 sovereign debt crisis, where the transmission to individual firms was largely mediated by aggregate financial and demand conditions, the pandemic generated firm-specific variation in the nature and intensity of the disruption (Hassan et al., 2023; Lamorgese et al., 2024). The INVIND 2020 wave (fielded in spring 2021) includes a dedicated COVID-19 module that records, for each firm, whether production was interrupted (cvd1_i) and the main reason for the interruption (r_i). We exploit these firm-level exposure measures on the 2020–2024 subpanel, with 2020 as the reference year, as predetermined and time-invariant firm attributes:

$$E_t[g_{i,t+1}] = \alpha_i + \beta \cdot g_{it} + \sum_{\tau \neq 2020} \delta_\tau \cdot (g_{it} \times D_\tau) + \delta^{\text{cvd}} \cdot (g_{it} \times \text{cvd1}_i) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (20)$$

where δ^{cvd} measures the differential effect of past growth on expectations for firms that interrupted production relative to those that did not. A significant δ^{cvd} indicates that the expectation-formation process differs between disrupted and non-disrupted firms, conditional on firm and year fixed effects.

We estimate analogous specifications for the second and third hypotheses, on the 2020–2024 subpanel with 2020 as the reference year. For Eq. (7), the dependent variable is $g_{i,t+1}$, the regressor interacted with the exposure variable is $E_t[g_{i,t+1}]$, and g_{it} enters as an additional control. For Eq. (10), the dependent variable is AFE_{it} and the interacted regressor is $u_{it|t-1}$.

The interruption dummy captures whether a firm was disrupted but does not distinguish among the reasons. To identify the specific channel, we replace the binary interaction with a categorical one:

$$E_t[g_{i,t+1}] = \alpha_i + \beta \cdot g_{it} + \sum_{\tau \neq 2020} \delta_\tau \cdot (g_{it} \times D_\tau) + \sum_{k=1}^5 \theta_k \cdot (g_{it} \times \mathbf{1}[r_i = k]) + \gamma_1 \cdot \text{size}_{it} + \gamma_2 \cdot \text{age}_{it} + d_t + v_{it}, \quad (21)$$

where $\mathbf{1}[r_i = k]$ is an indicator for firms reporting reason k as the main cause of their production interruption (1=regulatory closure, 2=insufficient demand, 3=workforce unavailability, 4=supply-chain disruption, 5=other), and the reference category $r_i = 0$ consists of firms that did not interrupt production. Each θ_k measures the differential expectation-formation slope for firms disrupted through channel k relative to non-disrupted firms.

This specification is estimated for each of the three hypotheses, allowing us to identify which disruption channel accounts for the heterogeneity observed in the full-sample year-varying estimates. We also estimate analogous interactions using the share of production suspended, the perceived probability of closure, and the moratorium indicator.

5. Results

This section presents the empirical results organized as follows. Section 5.1 addresses the first research hypothesis (expectation formation) for real growth, includes a persistence diagnostic of realized growth, and extends the analysis to price expectations. Section 5.2 presents the results for the second hypothesis on the predictive content of expectations. Section 5.3 reports the findings

Table 2

Expectation formation (H1). Dependent variable: expected real growth $E_t[g_{t+1}]$. Panel estimation with firm and year Fixed Effects (FE), INVIND survey weights.

	(1)	(2)	(3)
g_t		−0.181*** (0.057)	−0.106 (0.199)
$g_t \times D_\tau$			Yes
Size _{<i>t</i>}	−0.001** (0.001)	−0.001** (0.001)	−0.001** (0.001)
log(1 + age)	−0.702 (0.467)	−1.165** (0.479)	−1.282** (0.523)
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Clustering	Firm	Firm	Firm
Observations	38,670	38,670	38,670
R^2	0.253	0.276	0.325
Within R^2	0.000	0.031	0.096

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Column (1): controls only. Column (2): constant slope (Eq. (3)). Column (3): year-varying slope (Eq. (4), reference year=2018); 27-year interactions included, none significant at 10%. Full specification in Appendix Table D.13.

for the third hypothesis (forecast accuracy and the calibration of subjective uncertainty) for both real growth and prices.¹¹ Section 5.4 discusses the cross-sectional dispersion regressions of the fourth hypothesis. Section 5.5 examines the COVID-19 heterogeneity through firm-level exposure channels. For each hypothesis, we report a constant-slope specification, a year-varying specification with interactions d_τ (reference year 2018), and slope-path figures displaying the total year-specific coefficients with 95% confidence intervals.

5.1. Expectation formation (H1)

Table 2 presents the empirical results addressing the first hypothesis (H1), which investigates the influence of firms' performance on their expectations of future growth and explores the impact of the COVID-19 shock on this relationship. The table reports estimates of Eq. (3), where expected growth is regressed on experienced growth, along with additional controls, including firm size and age. All specifications include firm and year fixed effects to control for unobserved heterogeneity, with standard errors clustered at the firm level. Firm-level clustering accounts for serial correlation in firm-specific shocks to expectations and forecast errors across years of the panel. The year fixed effects absorb common aggregate shocks across firms, and firm-level clustering accounts for the residual serial dependence at the firm level.

We estimate this relationship on the full sample (1996–2024) using three specifications: controls only (column 1), a constant slope on g_{it} (column 2), and a year-varying slope via interactions $g_{it} \times D_\tau$ with 2018 as the reference year (column 3, Eq. (4)). The year-varying specification allows the expectation-formation slope to vary each year and nests the constant-slope model as a testable restriction.

The constant-slope specification in column (2) returns $\hat{\beta} = -0.181$, negative and significant at the 1% level. A negative coefficient is unusual in the firm-level expectations literature, where Gennaioli et al. (2016) document positive extrapolation in CFO forecasts of own firm earnings. Italian firms in INVIND forecast turnover growth instead, and a natural interpretation is that they incorporate the historical regularity that growth shocks tend to be partially reversed in the following year, rather than extrapolating current growth. Whether realized growth exhibits this regularity, and at what magnitude, is the question addressed by the persistence diagnostic below. Firm age enters with a negative sign, consistent with evidence that younger firms anticipate more rapid expansion (Jovanovic, 1982).

Column (3) extends the constant-slope specification by allowing year-specific deviations δ_τ relative to the 2018 baseline. None of the δ_τ is individually significant at conventional levels, implying that the expectation-formation slope is statistically indistinguishable from its constant-slope value in every year of the 1996–2024 period.¹²

We recover the year-specific slope as the sum of the constant-slope estimate and the year-specific deviation, $\hat{\beta}_\tau = \hat{\beta} + \hat{\delta}_\tau$. Fig. 1 plots these estimates with 95% confidence intervals over 1996–2024. The slope fluctuates within a narrow band around the constant-slope estimate of −0.181, with the vast majority of year-specific estimates negative and no year statistically distinguishable from this estimate. Neither the 2008–2009 financial crisis, the 2011–2012 sovereign debt crisis, nor the COVID-19 pandemic (2020–2022) produces a coefficient that differs significantly from the constant-slope estimate. The 2019 wave, fielded at the onset of the first lockdown (February–May 2020), shows a point estimate close to zero but is not statistically distinguishable from the constant-slope estimate, indicating that the pandemic affected realized outcomes but not the expectation-formation slope itself.

¹¹ The baseline coefficients of the three real-growth analyses above are robust to the inclusion of CIG hours as an additional firm-level control. The missing-rate audit of the candidate controls, and the corresponding sample cascade are documented in Appendix E of the Online Appendix.

¹² The full set of year-specific interactions is reported in Table D.13 of the Online Appendix.




Fig. 1. Expectation formation (H1). Total slope $\hat{\beta}_\tau = \hat{\beta} + \hat{\delta}_\tau$ of expected growth $E_t[g_{t+1}]$ on realized growth g_t , by year (Eq. (4)), with 95% confidence intervals.

Notes: Dashed red line: reference year 2018, selected as the last calendar year whose INVIND wave (fielded February–May 2019) was entirely unaffected by the COVID-19 pandemic. Year 2005 is absent because the INVIND survey did not collect the expected revenue-growth variable in the 2005 wave. The 2011 estimate ($\hat{\beta}_{2011} = -1.53$, not significant) is omitted for readability; see Appendix Table D.13. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The persistence diagnostic complements the expectation-formation result with an autoregressive specification of realized growth. If realized growth is mean-reverting at the firm level, a firm with a high current g_t would on average face a weaker year next, and the negative $\hat{\beta}$ tracks that regularity rather than indicating bias. The slope on the lag in Eq. (5) measures the magnitude of this regularity directly.

Column (1) of Table 3 reports the constant slope on the lag, $\hat{\rho} = -0.134$, negative and significant at the 1% level. After absorbing permanent firm-level differences and common aggregate shocks through the fixed effects, the remaining firm-specific component of realized growth is mean-reverting, where an above-average year is typically followed by a correction. The negative coefficient in Table 3, together with the negative slope on g_t in Table 2, indicates that Italian firms' forecasts are consistent with the mean-reverting structure of their own realized growth, anticipating it rather than extrapolating current growth into the future.

Column (2) extends the specification to year-specific persistence coefficients $\hat{\rho}_\tau = \hat{\rho} + \hat{\rho}^{(\tau)}$, where $\hat{\rho}^{(\tau)}$ is the year-specific deviation from the constant baseline. The persistence of realized growth changes over time, in contrast to the stability of the expectation-formation slope reported in Table 2. Mean reversion is stronger in post-crisis recoveries, notably after the 2009 financial crisis ($\hat{\rho}_{2010} = -0.40$) and after the 2020 contraction ($\hat{\rho}_{2021} = -0.33$). In expansion years, growth shows positive persistence, as in the recovery from the sovereign-debt crisis ($\hat{\rho}_{2017} = 0.15$) and during the post-pandemic normalization ($\hat{\rho}_{2023} = 0.06$).

The difference between the cyclical persistence of realized growth and the stable expectation-formation slope suggests that Italian firms anticipate the long-run degree of mean reversion in their data, rather than updating their forecasts year by year in response to current cyclical conditions.

For prices, we estimate Eq. (11) on realized price changes π_{it} and expected price changes $E_t[\pi_{i,t+1}]$, with results in Table 4. The constant-slope estimate (column 1) is positive and significant at the 1% level, $\hat{\beta} = 0.227$. Firms expect current price shocks to persist into the following year, consistent with firm-level transmission from realized to expected inflation. The sign is the opposite of the expectation-formation slope for real growth, where firms anticipated mean reversion.

Column (2) reports the year-varying specification where 13 of the 29 year interactions are statistically significant, while Fig. 2 plots the resulting year-specific slopes $\hat{\beta}_\tau$ with 95% confidence intervals.

The year-specific deviations $\hat{\delta}_\tau$ are negative and significant in the late 1990s and early 2000s, in 2008–2009, and in 2022. From 1996 onwards, when Italian inflation was low and stable, realized and expected price changes were weakly related. Between 2010 and 2021, the year-specific slopes range from 0.32 to 0.64; none of the year-specific deviations from the 2018 baseline is statistically significant, indicating a stable expectation-formation relationship across the decade. In 2022, the post-pandemic recovery and the Russia–Ukraine conflict generated an inflation surge whose timing was not anticipated in firms' forecasts, and the slope drops sharply. The relationship recovers in 2023–2024. The expectation-formation slope for prices, therefore, varies substantially across years, in contrast with the stability documented for real growth.¹³

5.2. Predictive content of expectations (H2)

The second hypothesis tests whether firms' expectations carry information about their realized future growth beyond what is predicted by current growth alone. The specification, given by Eq. (6), regresses realized growth $g_{i,t+1}$ on the expectation $E_t[g_{i,t+1}]$

¹³ The full set of coefficients is in Table D.16 of the Online Appendix.

Table 3

Persistence diagnostic. Dependent variable: realized growth g_t . Panel estimation with firm and year Fixed Effects (FE), INVIND survey weights.

	(1)	(2)
g_{t-1}	-0.134*** (0.022)	-0.161** (0.066)
$g_{t-1} \times D_{2010}$		-0.238*** (0.080)
$g_{t-1} \times D_{2017}$		0.308** (0.125)
$g_{t-1} \times D_{2021}$		-0.172** (0.087)
$g_{t-1} \times D_{2023}$		0.225*** (0.081)
$g_{t-1} \times D_t$		Yes
Controls	Yes	Yes
Firm FE/Year FE	Yes	Yes
Clustering	Firm	Firm
Observations	38,466	38,466
R^2	0.251	0.276
Within R^2	0.019	0.052

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Controls: size, and $\log(1 + \text{age})$. Column (1): constant lag coefficient (Eq. (5)). Column (2): year-varying lag coefficient (Eq. (5), ref=2018); 32-year interactions included, 4 significant at 5%, only those shown. The estimation sample (1992–2024) exceeds the Table 2 panel (1996–2024) because the lag regression does not require $E_t[g_{t+1}]$. The lower bound reflects the first wave with non-missing firm founding year, used to construct $\log(1 + \text{age})$.

Table 4

Price expectation formation (H1, price domain). Dependent variable: expected price change $E_t[\pi_{t+1}]$. Panel estimation with firm and year Fixed Effects (FE), INVIND survey weights.

	(1)	(2)
π_t	0.227*** (0.044)	0.543*** (0.162)
$\pi_t \times D_t$		Yes
Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Clustering	Firm	Firm
Observations	48,678	48,678
R^2	0.345	0.417
Within R^2	0.079	0.179

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Controls: size, and $\log(1 + \text{age})$. Column (1): controls only. Column (2): constant slope (Eq. (11) with $\delta_t = 0$ imposed). Column (3): year-varying slope (Eq. (11), reference year = 2018); 29-year interactions included, 13 significant at 10%. The estimation sample is 1995–2024, as the expected price-change variable $E_t[\pi_{t+1}]$ was introduced in the 1995 wave.

formed at time t , controlling for current growth g_{it} so that the coefficient on the expectation isolates its incremental predictive content. As shown in Table 5, both a constant-slope and a year-varying specification are estimated (Eq. (7), with 2018 as the usual reference year).

The constant-slope estimate (column 2) is $\hat{\beta} = 0.585$, positive and significant at the 1% level. Firm-level expectations predict realized future growth beyond what current growth alone implies. The coefficient on g_t drops from -0.170 in column (1) to -0.062 in column (2), with the expectation absorbing most of the apparent mean-reversion in the current growth coefficient.

Column (3) extends the constant-slope specification by allowing year-specific deviations δ_t on $E_t[g_{i,t+1}]$, with 2018 as the reference year ($\hat{\beta} = 0.600$). Three of the 26 year-specific interactions are individually significant at conventional levels, namely $\hat{\delta}_{2008}$ (the onset of the global financial crisis), $\hat{\delta}_{2015}$ (the post-sovereign-debt recovery), and $\hat{\delta}_{2019}$ (the wave fielded as the COVID-19 lockdown unfolded).¹⁴

The 2019 wave returns the largest year-specific slope of the sample, $\hat{\beta}_{2019} = 1.086$. This estimate is not a feature of pre-pandemic forecasting, as the INVIND 2019 wave was fielded between February and May 2020, overlapping with the first Italian lockdown, which began in March 2020. Firms surveyed in this window could observe the unfolding shock and the policy response, and they incorporated this information into their reported expectations for 2020 outcomes. Therefore, the high predictive content reflects

¹⁴ The full set of year-specific interactions is reported in Table D.14 of the Online Appendix.




Fig. 2. Price expectation formation (H1, price domain). Total slope $\hat{\beta}_\tau$ of expected price change $E_t[\pi_{t+1}]$ on realized price change π_t , by year (Eq. (11)), with 95% confidence intervals.

Notes: Dashed red line: reference year 2018. Dashed gray line: zero. The estimation sample is 1995–2024, as the expected price-change variable $E_t[\pi_{t+1}]$ was introduced in the 1995 wave. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 5

Predictive content (H2). Dependent variable: realized growth g_{t+1} . Panel estimation with firm and outcome-year Fixed Effects (FE), INVIND survey weights.

	(1)	(2)	(3)
$E_t[g_{t,t+1}]$		0.585*** (0.045)	0.600*** (0.126)
$E_t[g_{t,t+1}] \times D_{2008}$			0.413*** (0.148)
$E_t[g_{t,t+1}] \times D_{2015}$			0.329* (0.179)
$E_t[g_{t,t+1}] \times D_{2019}$			0.486** (0.189)
$E_t[g_{t,t+1}] \times D_\tau$			Yes
g_t	−0.170*** (0.022)	−0.062** (0.027)	−0.059*** (0.022)
Size _{<i>t</i>}	−0.003*** (0.001)	−0.002** (0.001)	−0.002** (0.001)
log(1 + age)	−1.732** (0.782)	−0.978 (0.656)	−0.939* (0.565)
Firm FE	Yes	Yes	Yes
Year FE (d_{t+1})	Yes	Yes	Yes
Clustering	Firm	Firm	Firm
Observations	28,004	28,004	28,004
R^2	0.273	0.468	0.494
Within R^2	0.029	0.291	0.325

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Controls: size, and log(1 + age). Column (1): g_t and controls only. Column (2): constant slope (Eq. (6)). Column (3): year-varying slope (Eq. (7), ref=2018); 26-year interactions included, 3 significant at 10%, all shown.

nowcasting during the outbreak rather than pre-shock forecasting. In general, all INVIND waves are fielded between February and May of the year $t+1$, but the 2019 wave is exceptional because a major shock occurred during the survey window. Had the pandemic occurred in the second half of 2020, the wave would have been completed before the shock, and the 2019 coefficient would lie within the historical range.

Fig. 3 plots the year-specific slopes $\hat{\beta}_\tau$ with 95% confidence intervals. Every year-specific slope is positive and significantly different from zero. Firms' expectations retain predictive information about future growth beyond the autoregressive component throughout the 1996–2023 window.

Setting the 2019 nowcasting case aside, the slope path shows a moderate decline during 2021–2022, with point estimates falling from 0.600 in 2018 to 0.443 in 2021 and 0.392 in 2022, before recovering to 0.794 in 2023. Although these point estimates are visually lower than the 2018 baseline, the year-specific deviations $\hat{\delta}_{2021} = -0.157$ and $\hat{\delta}_{2022} = -0.208$ are not individually significant. Therefore, the pandemic-period slopes are not statistically distinguishable from the long-run benchmark.




Fig. 3. Predictive content (H2). Total slope $\hat{\beta}_\tau$ of realized growth g_{t+1} on expected growth $E_t[g_{t+1}]$, by formation year (Eq. (7)), with 95% confidence intervals.

Notes: Dashed red line: reference year 2018. Year 2005 is absent because the INVIND survey did not collect the expected revenue-growth variable in the 2005 wave. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 6

Forecast accuracy (H3, real growth). Dependent variable: absolute forecast error AFE_{it} . Panel estimation with firm and year Fixed Effects (FE), INVIND survey weights.

	(1)	(2)
$u_{it t-1}$	0.274*** (0.086)	0.021 (0.025)
$u_{it t-1} \times D_\tau$		Yes
Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Clustering	Firm	Firm
Observations	21,513	21,513
R^2	0.445	0.476
Within R^2	0.073	0.125

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Controls: size, and $\log(1 + \text{age})$. Column (1): constant slope (Eq. (10) with $\delta_\tau = 0$ imposed). Column (2): year-varying slope (Eq. (10), ref=2018); 26 year interactions included, 15 significant at 10%. Year 2006 is absent because the 2005 INVIND wave did not collect the variables used to construct $u_{it|t-1}$. Full set of year interactions in Appendix Table D.15.

5.3. Forecast accuracy (H3)

The third hypothesis tests whether firms' self-reported uncertainty is informative about their forecast accuracy. The specification, given by Eq. (10), regresses the absolute forecast error AFE_{it} on lagged uncertainty $u_{it|t-1}$, the range of expected outcomes that the firm reported at $t - 1$.

As shown in Table 6, the constant-slope estimate (column 1) is $\hat{\lambda} = 0.274$, positive and significant at the 1% level. Firms reporting wider uncertainty ranges record larger forecast errors on average. Column (2) extends the constant-slope specification by allowing year-specific deviations δ_τ on $u_{it|t-1}$, indicating that 15 of the 26 year-specific interactions are individually significant at conventional levels, indicating substantial time variation in the uncertainty-accuracy relationship. The coefficient for 2018 is not statistically significant. In that year, firms reporting wider uncertainty ranges did not record larger forecast errors than those reporting narrower ranges.¹⁵

Fig. 4 plots the year-specific slopes with 95% confidence intervals. The relationship between uncertainty and forecast accuracy appears state-dependent, strengthening during periods of elevated aggregate uncertainty and weakening in calmer years (Altig et al., 2020). In 2018, 2023, and 2024, the total slope is close to zero and not statistically significant; self-reported uncertainty does not predict forecast accuracy in these years.

¹⁵ The full set of year-specific interactions is reported in Table D.15 of the Online Appendix.




Fig. 4. Forecast accuracy (H3, real growth). Total slope $\hat{\lambda}_\tau$ of absolute forecast error AFE_{it} on lagged uncertainty $u_{it|t-1}$, by year (Eq. (10)), with 95% confidence intervals.

Notes: Dashed gray line: zero. Dotted red line: reference year 2018. The uncertainty–accuracy relationship is state-dependent: strong during crisis episodes, near zero in calm years. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

During crisis episodes, the slope rises sharply where significant positive deviations from the 2018 baseline appear in the late 1990s and early 2000s (1999–2001, 2003–2005), during and after the global financial crisis (2007, 2009–2011), through the sovereign-debt crisis (2012–2013), in 2016, and during the post-pandemic inflation surge in 2022. The 2019 interaction is also significant but reflects the same survey-timing effect already discussed, as the wave was fielded between February and May 2020, during the first national lockdown. In these crisis years, firms reporting wider expected ranges produced significantly less accurate forecasts than those reporting narrower ranges, indicating that self-reported uncertainty bounds are informative about that of the macroeconomic environment.

The COVID-19 period is informative, as the 2019 interaction reflects the survey-timing effect previously discussed, where the wave was fielded as the first lockdown unfolded, highlighting nowcasting rather than forecasting. The year interactions for 2020 and 2021 are not statistically significant, meaning that self-reported uncertainty was not informative about forecast accuracy in these two years. The relationship becomes significant again in 2022, when the inflation shock reintroduced firm-specific price and cost uncertainty (Yotzov et al., 2023).

A plausible interpretation is that during 2020–2021, uncertainty was driven primarily by common aggregate shocks (lockdowns, reopenings, sector-wide restrictions) rather than firm-specific factors. In that environment, the cross-sectional variation in self-reported uncertainty bounds loses its discriminating power, since all firms face comparable forecast challenges regardless of the range they report.

Table 7 applies the same analysis to the price domain, using firms' ex ante price uncertainty $u_{it|t-1}^\pi$ to predict subsequent price forecast errors (Eq. (13)).

The constant-slope estimate (column 1) is positive and significant at the 1% level (0.205). Firms reporting wider price uncertainty record larger price forecast errors on average, and the magnitude is comparable to the real-growth coefficient of 0.274. The year-varying specification in column (2) shows that only one of the six year-specific deviations from the 2018 baseline (year 2017) is statistically significant. Unlike the real-growth case, where most interactions are individually significant, the cross-year variation in the price-uncertainty slope is far less systematic.¹⁶

Fig. 5 plots the year-specific slope $\hat{\lambda}_\tau$ with 95% confidence intervals. The highest value is in 2021 (0.408), followed by 2018 (0.263), while the 2020 estimate is not significantly different from zero (0.151). A natural interpretation is that during the pandemic year, common aggregate shocks (lockdowns, demand collapse) reduced the cross-sectional variation in firms' price uncertainty, weakening its ability to predict differences in forecast accuracy. The 2021 increase coincides with the start of the post-pandemic inflation surge, which plausibly heightened firm-specific exposures through differential cost pass-through and pricing power. Under this reading, firms reporting wider price uncertainty at the end of 2020 went on to record larger price forecast errors in 2021, and the price-domain uncertainty–accuracy link responds to the inflation shock one year before the corresponding strengthening for real growth.

¹⁶ The full set of year-specific interactions is reported in Table D.17 of the Online Appendix.

Table 7

Price forecast accuracy (H3, price domain). Dependent variable: absolute price forecast error AFE_{it}^{π} . Panel estimation with firm and year Fixed Effects (FE), INVIND survey weights.

	(1)	(2)
$u_{it t-1}^{\pi}$	0.205*** (0.064)	0.263*** (0.088)
$u_{it t-1}^{\pi} \times D_{\tau}$		Yes
Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Clustering	Firm	Firm
Observations	6210	6210
R^2	0.486	0.494
Within R^2	0.032	0.046

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Controls: size_{*i*} and log(1 + age). Column (1): constant slope (Eq. (13) with $\delta_{\tau} = 0$ imposed). Column (2): year-varying slope (Eq. (13), ref=2018); 6 year interactions included, 1 significant at 10%. The estimation sample is 2015–2021. The minimum and maximum expected-price-change variables used to construct $u_{it|t-1}^{\pi}$ were collected only between the 2014 and 2020 INVIND waves. Full set of year interactions in Appendix Table D.17.




Fig. 5. Price forecast accuracy (H3, price domain). Total slope $\hat{\lambda}_{\tau}$ of absolute price forecast error AFE_{it}^{π} on price uncertainty $u_{it|t-1}^{\pi}$, by year (Eq. (13)), with 95% confidence intervals.

Notes: Dashed gray line: zero. Dotted red line: reference year 2018. The estimation sample is 2015–2021. The minimum and maximum expected-price-change variables used to construct $u_{it|t-1}^{\pi}$ were collected only between the 2014 and 2020 INVIND waves. Full set of year interactions in Appendix Table D.17. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5.4. Cross-sectional dispersion (H4)

The fourth hypothesis examines whether disagreement across firms about future growth predicts the dispersion of subsequent realized growth. Specification 1 is the core test, described in Eq. (14), and regresses the interquartile range of realized growth, $IQR(g_t)$, on the lagged interquartile range of expected growth, $IQR(E_{t-1}[g_t])$.

To control for the possibility that $IQR(g_t)$ is itself persistent, we estimate two auxiliary specifications. Specification 2, Eq. (15), regresses $IQR(g_t)$ on its own lag, providing a benchmark for the persistence of realized dispersion, while specification 3, Eq. (16), includes both lagged regressors jointly to separate the predictive content of expectation dispersion from this persistence.

Specification 4 considers the interquartile range of absolute forecast errors as described in Eq. (17), with $IQR(AFE_t)$ as the dependent variable, testing whether expectation dispersion is associated with higher dispersion in forecast accuracy. Specification 5, presented in Eq. (18), reverses the direction of the core test by regressing $IQR(E_{t-1}[g_t])$ on lagged $IQR(g_t)$, testing whether past dispersion in realized growth shapes the dispersion of subsequent expectations.

We estimate each specification at four aggregation levels to assess whether the predictive content of expectation dispersion is robust to firm grouping. The levels are sector $\times$ macro-area ($11 \times 4 = 44$ groups), sector (11 groups), macro-area (4 groups), and the overall national time series (one observation per year). The choice of aggregation level involves a trade-off: finer aggregations preserve more cross-firm variation but rely on smaller groups, where dispersion measures are noisier. Conversely, coarser aggregations are less noisy but average over underlying differences across sectors and areas.

Table 8

Cross-sectional dispersion (H4, real growth). Each column reports a panel regression of the IQR of the dependent variable on the lagged IQR of the regressor(s), with group and year Fixed Effects (FE, except the overall level, which uses time-series OLS with heteroskedasticity-robust standard errors). Standard errors clustered at the group level in parentheses.

Dep. var.	(1) IQR(g_t)	(2)	(3)	(4) IQR(AFE $_t$)	(5) IQR($E_t[g_{t+1}]$)
<i>Panel A: Sector $\times$ macro-area (44 groups)</i>					
IQR($E_{t-1}[g_t]$)	0.363*** (0.076)		0.328*** (0.072)	0.315*** (0.049)	
IQR(g_{t-1})		0.203*** (0.046)	0.101*** (0.031)		-0.028 (0.068)
Within R^2	0.134	0.046	0.144	0.130	0.001
Observations	813	813	813	813	813
<i>Panel B: Sector (11 groups)</i>					
IQR($E_{t-1}[g_t]$)	1.171*** (0.254)		1.147*** (0.270)	0.921*** (0.222)	
IQR(g_{t-1})		0.286*** (0.052)	0.053 (0.054)		0.076 (0.051)
Within R^2	0.650	0.085	0.653	0.567	0.013
Observations	213	213	213	213	213
<i>Panel C: Macro-area (4 groups)</i>					
IQR($E_{t-1}[g_t]$)	0.305 (0.115)		0.295 (0.124)	0.160 (0.056)	
IQR(g_{t-1})		0.107 (0.047)	0.030 (0.079)		-0.093 (0.072)
Within R^2	0.069	0.011	0.070	0.033	0.011
Observations	96	96	96	96	96
<i>Panel D: Overall (1 group, time-series OLS)</i>					
IQR($E_{t-1}[g_t]$)	1.286*** (0.060)		1.289*** (0.082)	0.616*** (0.060)	
IQR(g_{t-1})		0.411 (0.206)	-0.006 (0.097)		-0.023 (0.137)
R^2	0.872	0.171	0.872	0.657	0.001
Observations	24	24	24	24	24

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Panels A–C: group and year fixed effects; cluster-robust SEs at the group level. Panel D: time-series OLS with intercept; heteroskedasticity-robust SEs (eigenvalue-corrected for positive semi-definite). The number of observations decreases from Panel A to Panel D because the cross-sectional dimension is progressively collapsed: each panel aggregates the firm-level data into fewer groups, leaving group $\times$ year cells as the unit of observation. All IQR measures are computed within each aggregation level shown in the panel headers.

Table 8 reports the results for real growth. The core test in specification 1 finds that expectation dispersion predicts subsequent realized dispersion, with positive and significant coefficients at the sector $\times$ macro-area, sector, and overall aggregation levels (0.363, 1.171, and 1.286, all at the 1% level). At the macro-area level, the coefficient is positive but not significant, suggesting that firms' disagreement is primarily sectoral rather than geographic.

Specification 2 finds significant persistence in realized dispersion at the finer aggregations, while specification 3 shows that the predictive content of expectation dispersion survives controlling for this persistence. The lagged IQR of expectations remains positive, significant, and close to its standalone magnitude (e.g., 0.363 to 0.328 at the sector $\times$ macro-area level), while the lagged IQR of realized growth retains some predictive content at the finest aggregations but is largely absorbed at coarser levels. Expectation dispersion carries forward-looking information about realized volatility that the persistence of realized dispersion does not capture.

Specification 4 repeats the core test with the IQR of absolute forecast errors as the dependent variable, showing that expectation dispersion is again positively associated with realized dispersion at the same aggregation levels. Firms that disagree more about expected growth are also less accurate on average, consistent with both phenomena reflecting underlying heterogeneity in firms' information sets. The reverse-direction test in specification (5) finds no significant relationship at any aggregation level, in contrast with the predictability documented in the forward direction.

Table 9 reports the results for prices, which parallel those for real growth with one key difference in the geographic dimension. For prices, the predictive relationship holds at the macro-area and sector levels but not at the finer sector $\times$ macro-area level, the reverse of the real-growth result, plausibly because inflationary pressures and supply-side conditions in Italy retain a regional component (Vaona and Ascari, 2012).

Specification 1 returns positive coefficients significant at the 5% level or better for the sector (0.875), macro-area (0.567), and overall (1.064) levels. Realized price dispersion is itself persistent (specification 2), but including both regressors jointly does not absorb the coefficient on the lagged IQR of expectations, which remains close to the standalone estimate from specification 1. At the sector $\times$ macro-area level, the persistence term absorbs most of the joint-regression explanatory power, consistent with the limited

Table 9

Cross-sectional dispersion (H4, prices). Each column reports a panel regression of the IQR of the dependent variable on the lagged IQR of the regressor(s), with group and year fixed effects (except the overall level, which uses time-series OLS with heteroskedasticity-robust standard errors). Standard errors clustered at the group level in parentheses.

Dep. var.	(1) IQR(π_t)	(2)	(3)	(4) IQR(AFE $_t^r$)	(5) IQR($E_t[\pi_{t+1}]$)
<i>Panel A: Sector $\times$ macro-area (44 groups)</i>					
IQR($E_{t-1}[\pi_t]$)	0.299 (0.275)		0.249 (0.232)	0.288 (0.305)	
IQR(π_{t-1})		0.295*** (0.050)	0.290*** (0.062)		0.002 (0.012)
Within R^2	0.011	0.089	0.097	0.009	0.000
Observations	1080	1080	1080	1080	1080
<i>Panel B: Sector (11 groups)</i>					
IQR($E_{t-1}[\pi_t]$)	0.875*** (0.084)		0.853*** (0.082)	0.358** (0.137)	
IQR(π_{t-1})		0.187*** (0.037)	0.043 (0.041)		-0.022 (0.040)
Within R^2	0.326	0.046	0.328	0.080	0.001
Observations	276	276	276	276	276
<i>Panel C: Macro-area (4 groups)</i>					
IQR($E_{t-1}[\pi_t]$)	0.567** (0.112)		0.552*** (0.132)	0.077 (0.131)	
IQR(π_{t-1})		0.189*** (0.020)	0.030 (0.125)		0.123 (0.057)
Within R^2	0.196	0.038	0.197	0.007	0.027
Observations	112	112	112	112	112
<i>Panel D: Overall (1 group, time-series OLS)</i>					
IQR($E_{t-1}[\pi_t]$)	1.064*** (0.057)		1.137*** (0.083)	0.661*** (0.075)	
IQR(π_{t-1})		0.557*** (0.171)	-0.097 (0.085)		0.123 (0.165)
R^2	0.872	0.313	0.877	0.625	0.020
Observations	28	28	28	28	28

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Panels A–C: group and year fixed effects; cluster-robust SEs at the group level. Panel D: time-series OLS with intercept; heteroskedasticity-robust SEs (eigenvalue-corrected for positive semi-definite). The number of observations decreases from Panel A to Panel D because the cross-sectional dimension is progressively collapsed: each panel aggregates the firm-level data into fewer groups, leaving group $\times$ year cells as the unit of observation. Sample sizes differ from those of Table 8 because the expected-price variable was introduced one wave earlier than expected real growth and has no internal gap year, whereas the expected real-growth variable was not collected in the 2005 wave. All IQR measures are computed within each aggregation level shown in the panel headers.

variation in price expectations at this fine aggregation. As in the real-growth case, expectation dispersion contains forward-looking information about realized price dispersion that the persistence of realized dispersion does not capture.

The lagged IQR of expectations also predicts the IQR of absolute price forecast errors at the sector (0.358) and overall (0.661) levels, significant at the 5% level or better (specification 4), while the reverse-direction test (specification 5) finds no significant relationship at any aggregation level, matching the real-growth result.¹⁷

Table 10 reports the average COVID-19 effect $\bar{\Delta}_{\text{COV}}$ on the cross-sectional dispersion of realized growth and forecast errors.¹⁸ The pandemic raised cross-firm dispersion at the sector $\times$ macro-area and macro-area levels for real growth (Panel A), but not at the sector level, and at every aggregation level for prices (Panel B), with magnitudes well above zero and uniformly significant across aggregations.

Considered together, the aggregation-level results suggest distinct propagation channels for the two domains. Real-side effects cluster along the geographic dimension and disappear at the sector level, consistent with the regional differentiation of lockdown rules and the unequal incidence of the pandemic across Italian macro-areas (Ascani et al., 2021). Price-side effects are detected uniformly across aggregations, consistent with the broader transmission of supply-chain disruptions and cost-push inflation.

Finally, Fig. 6 plots the year fixed effects $\hat{\mu}_t$ from the sector-level price-domain regression, showing the trajectory of cross-firm dispersion that Table 10 averages over the 2020–2022 pandemic period. Price forecast error dispersion is broadly stable until 2019 and rises substantially during 2020–2022 before moderating in 2023 and 2024. The corresponding figure for real growth is reported

¹⁷ We also estimate a contemporaneous reverse specification, regressing the IQR of expectations on the contemporaneous IQR of realized growth. The coefficient is positive and statistically significant across multiple aggregation levels in both the real-growth and price domains, indicating that firms partly incorporate currently observed dispersion into their expectations.

¹⁸ The shift is computed from the year fixed effects of the regressions in columns 1 and 4 of Tables 8 and 9, real growth and prices, respectively.

Table 10

Average COVID-19 effect on cross-sectional dispersion (H4, COVID period). Each cell reports the average year-effect shift from the year fixed effects of columns (1) and (4) of Tables 8 and 9, with 2018 as the reference year. Standard errors clustered at the group level.

Aggregation level	IQR(g_t)		IQR(AFE $_t$)	
	$\bar{\Delta}_{\text{COV}}$	Joint Wald p	$\bar{\Delta}_{\text{COV}}$	Joint Wald p
<i>Panel A: Real growth</i>				
Sector $\times$ macro-area	4.609***	< 0.001	2.966***	< 0.001
Sector	−0.429	0.274	−1.130	0.247
Macro-area	6.126***	< 0.001	4.242***	< 0.001
<i>Panel B: Prices</i>				
Sector $\times$ macro-area	4.621***	0.007	3.821**	0.044
Sector	2.248***	< 0.001	2.249***	< 0.001
Macro-area	1.953***	< 0.001	2.582***	< 0.001

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Each cell reports $\bar{\Delta}_{\text{COV}} = (\hat{\mu}_{2020} + \hat{\mu}_{2021} + \hat{\mu}_{2022})/3$ from columns (1) and (4) of Tables 8 and 9 (2018 reference year). Stars from a t -test of $H_0: \bar{\Delta}_{\text{COV}} = 0$; the Joint Wald p column tests $H_0: \mu_{2020} = \mu_{2021} = \mu_{2022} = 0$. All IQR measures are computed within each aggregation level shown in the panel headers.




Fig. 6. Year fixed-effect coefficients in the cross-sectional dispersion of price forecast errors, sector level, 1997–2024.

Notes. Blue dots: year fixed effects $\hat{\mu}_t$ from the sector-level regression of $\text{IQR}(\text{AFE}_{jt}^*)$ on the lagged IQR of expectations (Table 9, Panel B, column (4)). Red dot: 2018 reference year ($\hat{\mu}_{2018} = 0$). Vertical bars: 95% confidence intervals from cluster-robust standard errors at the group level. Gray band: COVID-19 period (2020–2022). The real-growth analog (Figure F.1 of the Appendix) does not exhibit a comparable COVID-period pattern. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

in Figure F.1 of the Online Appendix. The year fixed effects are smaller in magnitude and less significant, consistent with the average shift estimated in Table 10.

5.5. COVID-19 exposure and firm-level heterogeneity

The year-varying specifications in the previous sections capture variation across years. We turn to variation across firms within the pandemic period, exploiting differences in how the COVID-19 shock affected individual firms (Lamorgese et al., 2024) to identify which channels of the disruption account for the changes documented above.

Table 11

COVID-19 heterogeneity in expectation formation. Dependent variable: expected real growth $E_t[g_{t+1}]$. Panel estimation on the 2020–2024 subpanel with firm and year fixed effects (2020 as the reference year), INVIND survey weights. Year-varying slopes on g_t included in all specifications.

	(1)	(2)	(3)	(4)	(5)
$g_t \times$ Production stop	−0.313** (0.147)				
$g_t \times$ Reason: regulatory		−0.354** (0.171)			
$g_t \times$ Reason: demand		−0.057 (0.063)			
$g_t \times$ Reason: workforce		−1.497*** (0.461)			
$g_t \times$ Reason: supply chain		0.226 (0.171)			
$g_t \times$ Reason: other		0.081 (0.146)			
$g_t \times$ Stop share			−0.050 (0.036)		
$g_t \times$ Closure probability				0.111* (0.063)	
$g_t \times$ Moratorium request					−0.155 (0.120)
Year-varying slopes on g_t	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Clustering	Firm	Firm	Firm	Firm	Firm
Observations	7228	7198	1259	6584	7168
Adj. R^2	0.249	0.277	0.338	0.227	0.222

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the firm level in parentheses. Controls: size, and $\log(1 + \text{age})$. The reference category for the reason interactions is “no interruption”. All exposure variables are from the INVIND 2020 wave and are time-invariant in the subpanel.

The INVIND 2020 wave (fielded in spring 2021)¹⁹ includes a dedicated COVID-19 module recording whether production activity was disrupted, the firm’s financial response, and its perceived survival prospects. We use five firm-level exposure variables: (i) an indicator for whether production was interrupted ($\text{cvd}1_i$); (ii) the main reason for the interruption (r_i , with categories regulatory closure, insufficient demand, workforce unavailability, supply-chain disruption, and other)²⁰; (iii) the share of production suspended ($\text{cvd}3_i$, ordinal); (iv) the perceived probability of permanent closure ($\text{cvd}5_i$, ordinal from *not at all probable* to *extremely probable*)²¹; and (v) a binary indicator for whether a loan moratorium was requested ($\text{mor}20a_i$).²²

The production-stop interaction in column (1) is negative and significant at the 5% level (−0.313), indicating that firms that interrupted production have a steeper negative slope in current growth than non-disrupted firms. Thus, positive current growth translates into a steeper expected reversal, and negative current growth into a stronger anticipated rebound.

Column (2) breaks this effect down by reason for the interruption. Workforce unavailability dominates with an interaction of −1.497 at the 1% level, followed by regulatory closure, also negative and significant at the 5% level (−0.354); the remaining categories (insufficient demand, supply-chain disruption, and a residual “other” bin) do not register. The further exposure measures yield no significant interaction with the share of production suspended (column 3) or the moratorium request (column 5), and only a weakly positive interaction at the 10% level with the perceived probability of permanent closure (column 4). Direct operational obstacles, workforce shortages, or regulatory closures are the only channels through which the pandemic changed how firms formed expectations on current growth, while intensity of the disruption and reliance on financial support did not register, and subjective survival prospects entered only marginally. The earlier ISECO survey (Lamorgese et al., 2024) identifies demand, supply, and labor shortages as the salient disruption channels for the same population in the same period.

¹⁹ The COVID questions are explicitly phrased as of early 2021 (“*Dall’inizio del 2021 ad oggi...*”). Banca d’Italia indexes each wave under the prior calendar year (the year of the reported final-year accounts), placing these questions under year = 2020 in the dataset. We confirmed this convention with Banca d’Italia, Team Microdati Imprese.

²⁰ The variable r_i is constructed from INVIND variable $\text{cvd}4a$. We extend the original five-category coding (1–5) with a reference category $r_i = 0$ for firms that did not interrupt production. Firms with $\text{cvd}1_i = 1$ but missing $\text{cvd}4a$ are dropped.

²¹ The five $\text{cvd}5$ categories are: 1 = not at all probable [0%–20%], 2 = slightly probable [20%–40%], 3 = fairly probable [40%–60%], 4 = very probable [60%–80%], 5 = extremely probable [80%–100%].

²² The moratorium variable is recoded from the original three-level survey field (1 = requested and granted, 2 = requested and not granted, 3 = not requested) into a binary indicator equal to one if the firm requested a moratorium (categories 1 or 2 pooled) and zero if it did not (category 3). Observations outside this set are treated as missing.

Table 12

COVID-19 heterogeneity in predictive content. Dependent variable: realized growth g_{t+1} . Panel estimation on the 2020–2024 subpanel with firm and outcome-year fixed effects, INVIND survey weights. Year-varying slopes on $E_t[g_{t+1}]$ included in all specifications.

	(1)	(2)	(3)	(4)	(5)
$E_t[g_{t+1}] \times$ Production stop	0.047 (0.099)				
$E_t[g_{t+1}] \times$ Reason: regulatory		0.146 (0.089)			
$E_t[g_{t+1}] \times$ Reason: demand		0.262 (0.245)			
$E_t[g_{t+1}] \times$ Reason: workforce		−0.737*** (0.110)			
$E_t[g_{t+1}] \times$ Reason: supply chain		−0.008 (0.209)			
$E_t[g_{t+1}] \times$ Reason: other		−0.256 (0.247)			
$E_t[g_{t+1}] \times$ Stop share			0.071 (0.094)		
$E_t[g_{t+1}] \times$ Closure probability				0.041 (0.025)	
$E_t[g_{t+1}] \times$ Moratorium request					0.005 (0.073)
Year-varying slopes on $E_t[g_{t+1}]$	Yes	Yes	Yes	Yes	Yes
Controls + g_t	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Outcome-year FE	Yes	Yes	Yes	Yes	Yes
Clustering	Firm	Firm	Firm	Firm	Firm
Observations	4770	4749	828	4377	4744
Adj. R^2	0.376	0.397	0.485	0.393	0.377

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at firm level in parentheses. Controls: size, $\log(1 + \text{age})$, and g_t . See Table 11 notes for variable definitions.

Table 12 reports the heterogeneity results for predictive content about future growth. The production-stop indicator in column (1) has no significant interaction with $E_t[g_{t+1}]$, while column (2) shows that workforce unavailability is the only significant (and negative) channel, indicating that the expectations of firms whose production was interrupted by labor shortages carried significantly less predictive content for realized future growth than those of non-disrupted firms.

None of the remaining exposure measures yields a significant interaction with $E_t[g_{t+1}]$, including the share of production suspended (column 3), the perceived closure probability (column 4), and the moratorium request (column 5). The finding is analogous to the expectation-formation result in Table 11, where workforce unavailability was again the dominant channel. Labor shortages emerge as the only disruption that systematically altered both how firms map current growth into expectations and how informative those expectations were about realized outcomes.

The corresponding analysis for forecast accuracy, reported in Appendix G, finds no significant heterogeneity across the identifiable COVID exposure channels, with the residual “other” reason category as the only positive significant interaction.²³

6. Conclusion

The COVID-19 pandemic of 2020 represented an exogenous shock that generated economic uncertainty and a contraction in activity across sectors and countries. Italy has been among the first European countries to impose nationwide restrictions and among the most severely affected during the first wave, making it an informative setting for studying how firms’ expectations and uncertainty respond to a major macroeconomic shock. We use firm-level panel data from the Bank of Italy’s INVIND survey, which has tracked Italian industrial and service firms since 1996 and collects direct measures of expected and realized growth, expected and realized price changes, and the firm-level uncertainty surrounding these expectations. We investigate four aspects of firms’ expectations and ask whether the pandemic produced a departure from pre-pandemic regularities. We examine (i) how current performance shapes firms’ growth expectations, (ii) whether these expectations are informative about future outcomes, (iii) how self-reported subjective uncertainty relates to forecast accuracy, and (iv) whether cross-sectional dispersion in firm expectations predicts the dispersion of subsequent realized outcomes. We extend the analysis to the price domain and, exploiting the firm-level heterogeneity in pandemic exposure recorded in the dedicated INVIND COVID-19 module, identify the channels through which the shock affected expectation formation.

²³ As a complement, we also examine aggregate-level heterogeneity using a DPCM-based essential-sector indicator and the macro-area-year closure-stringency index of Conteduca and Borin (2022); see Appendix H.

First, results show that the link between current growth and growth expectations is negative and stable. Over nearly three decades, including the 2008–2009 financial crisis, the 2011–2012 sovereign debt crisis, and the COVID-19 pandemic, we do not detect a departure from its long-run value. The pandemic did not alter how Italian firms map current growth into their growth expectations, as firms expect a partial reversal of current performance regardless of the macroeconomic regime. Price-expectation formation behaves differently, with the response of expected prices to realized prices being positive on average but unstable over time, weak in the 1990s, and breaking down during the 2008–2009 financial crisis and again during the post-pandemic inflation surge.

Second, growth forecasts predict realized growth in every year of the sample, and this predictive content persists through the pandemic. Although the relationship weakened moderately in 2021 and 2022, it never approached zero. Firms' expectations retained genuine information about future outcomes even during the pandemic's most acute phase. The strongest predictive relationship in the sample appears in the 2019 wave, which was fielded between February and May 2020 as the first national lockdown unfolded. Firms surveyed in those weeks were able to incorporate the shock into their reports, with this case reflecting nowcasting of the ongoing event. This is a unique circumstance, since had the pandemic broken out a few months later, the wave would have been completed before the shock, and the 2019 result would not stand out from the rest of the analysis.

Third, firms that report wider expected ranges around their growth forecasts also tend to make larger forecast errors, but only in some years. The link is strong during crisis episodes, close to zero in calm years, and weak in 2020–2021, when uncertainty was driven by common aggregate shocks rather than firm-specific factors. The relationship returns in 2022, when the post-pandemic inflation surge reintroduces firm-specific price and cost uncertainty. A possible interpretation is that during the most acute phase of the pandemic, all firms faced comparable forecast challenges, and the cross-sectional variation in self-reported uncertainty was marginal, limiting its ability to discriminate between more and less accurate forecasters. For price expectations, by contrast, the link between subjective uncertainty and forecast errors is positive and stable throughout the available window, without the time variation observed for growth.

Fourth, when firms disagree more about future growth or future prices, the cross-section of subsequent outcomes is also more dispersed. This holds for both the real-growth and the price domain at most aggregation levels and is not driven by the persistence of realized dispersion, since the result survives the inclusion of its own lag as a control. The relationship runs only in one direction, as we do not find any pattern in the reverse regression at any aggregation level, in which lagged realized dispersion is used to predict expectation dispersion. This suggests that disagreement among firms reflects the underlying heterogeneity in their information sets about future conditions, and that expectations therefore carry information about the spread of outcomes still to be observed.

Finally, we exploit firm-level information on pandemic exposure to identify the channels through which the pandemic affected expectations, and find that workforce unavailability is the dominant one. Firms that interrupted production due to workforce shortages anticipate a larger reversal in current performance than firms that did not, and their forecasts are also less informative about future outcomes. By contrast, firms that interrupted production because of regulatory closures share the first effect, but their forecasts remain as informative as those of other firms. Demand contractions, supply-chain disruptions, and the financial channels captured by the moratorium indicator have no significant effect, suggesting that the pandemic shock to firm expectations operated mainly through the labor-supply margin, with channels usually associated with macroeconomic crises playing a smaller role.

These findings document how the COVID-19 pandemic affected the formation, predictive content, accuracy, and aggregation of firms' expectations. The mechanism through which firms map current growth into expected growth remained largely unchanged, while the informativeness of subjective uncertainty about forecast errors and the cross-sectional dispersion of expectations responded to the shock in ways that differ across the real-growth and price domains. At the aggregate level, cross-sectional dispersion among firms predicts future volatility in a unidirectional way that does not run from outcomes to expectations, and at the firm level, the disruption was concentrated on the labor-supply margin, with workforce unavailability emerging as the dominant channel. These results inform our understanding of how firms cope with episodes of high Knightian uncertainty. They also point to the workforce channel as a transmission mechanism for aggregate shocks to firm-level expectations, distinct from the demand and financial channels emphasized in earlier crisis literature.

Future research could extend the analysis to other surveys and countries and examine how firm expectations respond to other shocks, including geopolitical tensions and trade-policy disruptions that have unfolded in the post-pandemic period.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jebo.2026.107716>.

Data availability

The firm-level data used in this study are not in the authors' possession. They remain under the exclusive control of the Bank of Italy and can be accessed only through REX, the Bank's remote execution environment. Researchers upload code to REX, where it is executed on the secure server hosting the INVIND database; only the output that passes the Bank's disclosure review is returned. We therefore submitted our scripts to REX and received the vetted results without ever downloading or handling the raw firm-level data locally.

References

- Altig, D., Baker, S., Barrero, J.M., Bloom, N., Bunn, P., Chen, S., Davis, S.J., Leather, J., Meyer, B., Mihaylov, E., et al., 2020. Economic uncertainty before and during the COVID-19 pandemic. *J. Public Econ.* 191, 104274.
- Altig, D., Barrero, J.M., Bloom, N., Davis, S.J., Meyer, B., Parker, N., 2022. Surveying business uncertainty. *J. Econometrics* 231 (1), 282–303.
- Anderson, R.C., Reeb, D.M., 2003. Founding-family ownership and firm performance: evidence from the S&P 500. *J. Financ.* 58 (3), 1301–1328.
- Andrade, P., Coibion, O., Gautier, E., Gorodnichenko, Y., 2022. No firm is an island? How industry conditions shape firms' expectations. *J. Monet. Econ.* 125, 40–56.
- Ascani, A., Faggian, A., Montresor, S., 2021. The geography of COVID-19 and the structure of local economies: The case of Italy. *J. Reg. Sci.* 61 (2), 407–441.
- Avalos, E., Barrero, J.M., Davies, E., Iacovone, L., Torres, J., 2024. Business Uncertainty in Developing and Emerging Economies. Tech. Rep., CESifo Working Paper, No. 11368.
- 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.
- Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring economic policy uncertainty. *Q. J. Econ.* 131 (4), 1593–1636.
- Baker, S.R., Bloom, N., Davis, S.J., Terry, S.J., 2020. Covid-Induced Economic Uncertainty. Tech. Rep., National Bureau of Economic Research.
- Baley, I., Blanco, A., 2019. Firm uncertainty cycles and the propagation of nominal shocks. *Am. Econ. J.: Macroecon.* 11 (1), 276–337.
- Banca d'Italia, 2024. Indagine sulle imprese industriali e dei servizi. <https://www.bancaditalia.it/statistiche/basi-dati/rdc/accessomicrodati/>, Divisione Indagini Campionarie del Servizio Analisi statistiche, documentazione per l'utilizzo dell'archivio.
- Berger, D., Vavra, J.S., 2013. Volatility and Pass-Through. Tech. Rep., National Bureau of Economic Research Cambridge, MA, USA.
- Bloom, N., 2009. The impact of uncertainty shocks. *Econometrica* 77 (3), 623–685.
- Bloom, N., 2014. Fluctuations in uncertainty. *J. Econ. Perspect.* 28 (2), 153–176.
- Bloom, N., Bond, S., Van Reenen, J., 2007. Uncertainty and investment dynamics. *Rev. Econ. Stud.* 74 (2), 391–415.
- Bloom, N., Davis, S.J., Foster, L., Lucking, B., Ohlmacher, S., Saporta-Eksten, I., 2020. Business-Level Expectations and Uncertainty. Tech. Rep., National Bureau of Economic Research.
- Bontempi, M.E., Golinelli, R., Parigi, G., 2010. Why demand uncertainty curbs investment: Evidence from a panel of Italian manufacturing firms. *J. Macroecon.* 32 (1), 218–238.
- Bottoni, M., 2025. Comparing survey measures of firms' expectations and uncertainty. *Empir. Econ.* 1–38.
- Buchetti, B., Parbonetti, A., Pugliese, A., 2022. Covid-19, corporate survival and public policy: The role of accounting information and regulation in the wake of a systemic crisis. *J. Account. Public Policy* 41 (1), 106919.
- Buchheim, L., Dovern, J., Krolage, C., Link, S., 2022. Sentiment and firm behavior during the COVID-19 pandemic. *J. Econ. Behav. Organ.* 195, 186–198.
- Bunn, P., Bloom, N., Crundwell, A., Khan, S., Menzies, C., Mizen, P., Sculthorpe, M., Shah, K., Thwaites, G., Yotzov, I., 2024. The decision maker panel: a user's guide. Bank of England Staff Working Paper.
- Caldara, D., Iacoviello, M., 2022. Measuring geopolitical risk. *Am. Econ. Rev.* 112 (4), 1194–1225.
- Carletti, E., Oliviero, T., Pagano, M., Pelizzon, L., Subrahmanyam, M.G., 2020. The COVID-19 shock and equity shortfall: Firm-level evidence from Italy. *Rev. Corp. Financ. Stud.* 9 (3), 534–568.
- Cerqueti, R., Coppier, R., Girardi, A., Ventura, M., 2022. The sooner the better: lives saved by the lockdown during the COVID-19 outbreak. The case of Italy. *Econom. J.* 25 (1), 46–70.
- Chen, C., Senga, T., Zhang, H., 2021. Measuring business-level expectations and uncertainty: survey evidence and the COVID-19 pandemic. *Jpn. Econ. Rev.* 72, 509–532.
- Coibion, O., Gorodnichenko, Y., Kumar, S., 2018. How do firms form their expectations? New survey evidence. *Am. Econ. Rev.* 108 (9), 2671–2713.
- Commission, E., 2022. Italy 2022 Country Report. SWD(2022) 640 Final, European Commission, Brussels, URL https://ec.europa.eu/info/publications/2022-european-semester-country-report-italy_en.
- Conteduca, F.P., Borin, A., 2022. A new dataset for local and national COVID-19-related restrictions in Italy. *Ital. Econ. J.* 8 (2), 435–470.
- Di, L., Papadia, G., et al., 2016. Using External Sources to Understand Sample Survey Bias: the Case of the Invind Survey. Tech. Rep., Bank of Italy, Economic Research and International Relations Area.
- Diether, K.B., Malloy, C.J., Scherbina, A., 2002. Differences of opinion and the cross section of stock returns. *J. Financ.* 57 (5), 2113–2141.
- Dovern, J., Müller, L.S., Wohlrabe, K., 2023. Local information and firm expectations about aggregates. *J. Monet. Econ.* 138, 1–13.
- Duncan, R.B., 1972. Characteristics of organizational environments and perceived environmental uncertainty. *Adm. Sci. Q.* 313–327.
- Enders, Z., Hünnekes, F., Müller, G., 2022. Firm expectations and economic activity. *J. Eur. Econ. Assoc.* 20 (6), 2396–2439.
- Engle, R.F., Campos-Martins, S., 2023. What are the events that shake our world? Measuring and hedging global COVOL. *J. Financ. Econ.* 147 (1), 221–242.
- Epstein, L.G., Wang, T., 2004. Intertemporal asset pricing under Knightian uncertainty. In: *Uncertainty in Economic Theory*. Routledge, pp. 445–487.
- Ercolani, V., Natoli, F., 2020. Forecasting US recessions: the role of economic uncertainty. *Econom. Lett.* 193, 109302.
- Fiori, G., Scoccianti, F., 2023. The economic effects of firm-level uncertainty: Evidence using subjective expectations. *J. Monet. Econ.* 140, 92–105.
- Gennaioli, N., Ma, Y., Shleifer, A., 2016. Expectations and Investment. *NBER Macroecon. Annu.* 30 (1), 379–431.
- Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I., 2022. Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages? *Am. Econ. Rev.* 112 (5), 1437–1474.
- Guiso, L., Parigi, G., 1999. Investment and demand uncertainty. *Q. J. Econ.* 114 (1), 185–227.
- Hassan, T.A., Hollander, S., Van Lent, L., Schwedeler, M., Tahoun, A., 2023. Firm-level exposure to epidemic diseases: Covid-19, SARS, and H1N1. *Rev. Financ. Stud.* 36 (12), 4919–4964.
- Jovanovic, B., 1982. Selection and the evolution of industry. *Econometrica* 649–670.
- Jurado, K., Ludvigson, S.C., Ng, S., 2015. Measuring uncertainty. *Am. Econ. Rev.* 105 (3), 1177–1216.
- Kellard, N.M., Kontonikias, A., Lamla, M., Maiani, S., 2022. Deal or no deal? Modelling the impact of brexit uncertainty on UK private equity activity. *Br. J. Manag.* 33 (1), 46–68.
- Knight, F.H., 1921. *Risk, Uncertainty and Profit*, vol. 31, Houghton Mifflin.
- Lamorgese, A., Patnaik, M., Linarello, A., Schivardi, F., 2024. Management practices and resilience to shocks: Evidence from COVID-19. *Manag. Sci.* 70 (12), 9058–9072.
- Lehmann, R., 2023. The forecasting power of the Ifo business survey. *J. Bus. Cycle Res.* 19 (1), 43–94.
- Lewin, A.Y., Sakano, T., Stephens, C.U., Victor, B., 1995. Corporate citizenship in Japan: Survey results from Japanese firms. *J. Bus. Ethics* 14, 83–101.
- Lucas, Jr., R.E., 1972. Expectations and the Neutrality of Money. *J. Econom. Theory* 4 (2), 103–124.
- Lucas, Jr., R.E., 1977. Understanding business cycles. *Carnegie-Rochester Conf. Ser. Public Policy* 5, 7–29.
- Maćkowiak, B., Wiederholt, M., 2009. Optimal sticky prices under rational inattention. *Am. Econ. Rev.* 99 (3), 769–803.
- Massenot, B., Pettinicchi, Y., 2018. Can firms see into the future? Survey evidence from Germany. *J. Econ. Behav. Organ.* 145, 66–79.
- Milliken, F.J., 1987. Three types of perceived uncertainty about the environment: State, effect, and response uncertainty. *Acad. Manag. Rev.* 12 (1), 133–143.
- Notarnicola, G., 2018. *Essays on Firm-Level Shocks, Expectations and Uncertainty* (Ph.D. thesis). Università Ca' Foscari Venezia.
- Olley, S., Pakes, A., 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64 (6), 1263–1297.

- Tanaka, M., Bloom, N., David, J.M., Koga, M., 2020. Firm performance and macro forecast accuracy. *J. Monet. Econ.* 114, 26–41.
- Tzouvanas, P., Kizys, R., Chatziantoniou, I., Sagitova, R., 2019. Can variations in temperature explain the systemic risk of European firms? *Environ. Resour. Econ.* 74 (4), 1723–1759.
- Vaona, A., Ascari, G., 2012. Regional inflation persistence: evidence from Italy. *Reg. Stud.* 46 (4), 509–523.
- Yotzov, I., Anayi, L., Bloom, N., Bunn, P., Mizen, P., Öztürk, Ö., Thwaites, G., 2023. Firm inflation uncertainty. In: *AEA Papers and Proceedings*, vol. 113, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 56–60.