

NBER WORKING PAPER SERIES

NO FIRM IS AN ISLAND? HOW INDUSTRY CONDITIONS SHAPE FIRMS' AGGREGATE
EXPECTATIONS

Philippe Andrade
Olivier Coibion
Erwan Gautier
Yuriy Gorodnichenko

Working Paper 27317
<http://www.nber.org/papers/w27317>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2020

This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the “Investissements d’Avenir” program (reference: ANR-10-EQPX-17 – CASD). Coibion and Gorodnichenko thank NSF (SES # 1530467) for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, the Banque de France, the Eurosystem, the Federal Reserve Bank of Boston, or the Federal Reserve System.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Philippe Andrade, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

No Firm is an Island? How Industry Conditions Shape Firms' Aggregate Expectations
Philippe Andrade, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko
NBER Working Paper No. 27317

June 2020

JEL No. E2,E3,E4

ABSTRACT

We study how firms' expectations and actions are affected by both aggregate and industry-specific conditions using a survey of French manufacturing firms. We document two novel features. First, the adjustment of firms' expectations is more rapid after industry-specific shocks than aggregate shocks. This is consistent with rational inattention models which predict that firms should pay more attention to industry variation than aggregate conditions. Second, in response to industry shocks that have no aggregate effects, firms' aggregate expectations respond. This is consistent with "island" models in which firms use the specific prices they observe to make inferences about broader aggregate conditions. We also study how these results vary across industries.

Philippe Andrade
Federal Reserve Bank of Boston
philippe.andrade@bos.frb.org

Olivier Coibion
Department of Economics
University of Texas at Austin
2225 Speedway
Austin, TX 78712
and NBER
ocoibion@gmail.com

Erwan Gautier
Banque de France
31 rue Croix des Petits-Champs
75049 PARIS Cedex 01
France
and University of Nantes
erwan.gautier@banque-france.fr

Yuriy Gorodnichenko
Department of Economics
530 Evans Hall #3880
University of California, Berkeley
Berkeley, CA 94720-3880
and IZA
and also NBER
ygorodni@econ.berkeley.edu

“Information on the price expectations of businesses who are, after all, the price setters in the first instance (...) is particularly scarce. ... How do changes in various measures of inflation expectations feed through to actual pricing behavior? ... What factors affect the level of inflation expectations...?” Ben Bernanke (2007)¹

1. Introduction

Most modern business cycle models are built on the idea that economic agents have full-information rational expectations (FIRE). While most humans lack the ascribed FIRE abilities, one might think that firms’ executives are much more informed and educated than a typical consumer or worker in the economy so that the central theoretical tenet of the current business cycle research program still serves as a reasonable approximation. As observed by Bernanke (2007), we clearly need more facts to establish empirical support for this notion but the growing evidence appears to be discouraging. For example, disagreement among firms about future aggregate conditions is pervasive and large, much larger than disagreement among professional forecasters. Coibion et al. (2020) report results from a U.S. survey of firms’ inflation expectations that reveals a level of disagreement which is close to the high levels observed for households and far greater than anything observed among informed professional. Where does this disagreement stem from? In principle, firms should be observing similar aggregate statistics and therefore forming similar beliefs about the future, much like professional forecasters. In this paper, we provide new evidence documenting how conditions in a firm’s industry play an important role in shaping their view of broader macroeconomic conditions.

This evidence is two-fold. First, we show that consistent with a prediction of rational inattention models (Mackowiak and Wiederholt 2009), firms’ expectations respond much more rapidly to industry-specific shocks than they do to aggregate shocks. This indicates that firms are deliberately choosing to be less informed about aggregate conditions. With less precise information, we should therefore naturally expect substantial disagreement among them about future aggregate conditions. Second, we show that consistent with island-type models (in which firms form beliefs about the aggregate using the industry-specific information they are exposed to), firms’ expectations about aggregate economic conditions respond to shocks to their industry even though these shocks have no aggregate effects. This indicates that firms treat the signals they receive about their industries as informative about the aggregates and, in part, rely on these to form beliefs about broader economic

¹ <https://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm>

conditions. Jointly, these two facts provide a simple explanation for high and volatile disagreement among firms about aggregate beliefs. They have a strong incentive to focus on information about their industry as this affects their profits more and rely on this information to form beliefs about the broader economy. But because industry volatility is so high, this will contribute to significant and volatile disagreement among firms about the state of the economy.

We establish these facts using a little-used survey of French manufacturers. This survey has both a large cross-section of firms (~3,000) that are repeatedly surveyed over time as well as a long time series duration: it has been running quarterly since 1992. Firms are asked both about their own conditions and actions (e.g. recent price and production changes) as well as their expectations about the future (both about their own decisions and the broader economic outlook). This makes it an ideal survey to study how firms' expectations and decisions respond to both industry-specific and aggregate shocks. We verify the quality of firms' answers by comparing averages across the survey to broader macroeconomic aggregates and find striking overlap between the two. We can also verify that firms' reported expectations of their future decisions, on average, line up closely with their subsequent actions. In other words, the quality of the expectations responses appear to be high as well.

To characterize how firms' decisions and expectations respond to industry and aggregate shocks, we use a local projections specification which jointly estimates the dynamic response of firms' expectations and decisions to variation in industry and aggregate conditions. The latter two are measured using aggregate inflation and industry inflation, although similar results obtain when we use production data instead. This provides a tractable approach to address a number of related questions.

We first consider the response of firms' expected prices and production to each type of shock. It has been well-documented in prior work that prices tend to respond much more rapidly to industry-level shocks than aggregate shocks (Boivin et al. 2009, Mackowiak et al. 2009). We find a similar pattern. Importantly, we show that firms' expectations about their future price changes also respond more rapidly. After an industry-specific shock, firms immediately anticipate that they will raise prices, whereas very little adjustment in expectations takes place after an aggregate shock. This supports the rational inattention mechanism emphasized in Mackowiak and Wiederholt (2009) and Mackowiak et al. (2009). These studies argue that because aggregate shocks play a relatively small role in firm profits compared to idiosyncratic or industry variation, firms choose to devote little attention to tracking aggregate shocks. The result of this mechanism is that one should observe rapid price adjustment after industry shocks that firms pay attention to but much more gradual adjustment

of prices to aggregate shocks since firms only gradually learn about them. Our results support this mechanism and provide direct evidence for this differential learning behavior using surveys of firms' expectations.

We then consider how firms' aggregate expectations respond to industry vs aggregate shocks. The survey includes qualitative questions to firms about whether they expect broader prices and production to increase/decrease or stay the same. While qualitative questions prevent us from drawing clear quantitative conclusions about the magnitudes of firms' responses, we can still characterize the qualitative patterns in their expectations. We find a striking result: industry-level variation that is orthogonal to aggregate conditions has a pronounced and persistent effect on firms' aggregate expectations. This is striking because our empirical specification includes aggregate variables, so our identifying variation in industry variables is one that has no aggregate effects. Yet firms' beliefs about the aggregate respond to this industry-specific variation.

One interpretation of this result could be that firms are correctly anticipating that contemporaneous industry variation may have delayed aggregate effects, e.g. through input-output structures, even though they have no contemporaneous aggregate effects. We control for this possibility in a number of ways: dropping sectors for which we can reject the null that industry shocks have zero aggregate effects at different horizons, including a more extensive array of lags in the empirical specification to control for dynamic effects, including time fixed effects to soak up all aggregate variation, etc. Our result is impervious to these changes. We find systematic evidence that firms' aggregate expectations respond to industry shocks that have no aggregate effects, a clear violation of the full-information rational expectations (FIRE) hypothesis.

While at odds with FIRE, this result is consistent with a long line of "island" models in macroeconomics in which firms observe only a subset of prices in the economy with which they transact and use these prices to inform their beliefs about aggregate shocks (e.g. Lucas 1972, Lorenzoni 2009, Angeletos and Lao 2013, Nimark 2014, Afrouzi 2016). These models have been influential in providing potential explanations for monetary non-neutrality or expectations-driven shocks, but there has been little to no empirical evidence on the mechanism underlying these models. To the best of our knowledge, we are the first to provide direct evidence of this type of learning taking place among firms.

Evidence of learning from observed prices does exist for households. For example, Coibion and Gorodnichenko (2015) emphasize the role played by gasoline prices in shaping households' inflation expectations. Cavallo et al. (2017) and D'Acunto et al. (2019) study how the prices faced by

households on a frequent basis in their shopping affect their inflation expectations. They find that frequently purchased (salient) goods' price changes map clearly into households' beliefs about broader price movements. Kumar et al. (2015) similarly note that some firm managers in New Zealand identify the prices faced in their shopping as being a primary determinant of their inflation expectations. But direct evidence of firms' learning from their industries is missing.

Our paper builds on a much broader literature studying the expectations formation of economic agents and how those expectations affect their decisions. This literature has primarily focused on characterizing how expectations respond to aggregate shocks (e.g. Coibion and Gorodnichenko 2012), the predictability of expectations (e.g. Coibion and Gorodnichenko 2015), or the characteristics of forecast revisions (e.g. Andrade and Le Bihan 2013). We differ from this earlier line of work by studying in particular how firms' expectations respond to industry conditions. We document a significantly lower degree of inattention to industry conditions using the conditional response of forecasts to industry variation. In addition, we show that firms' aggregate expectations respond not just to aggregate shocks but also to industry ones.

2. Data

Our analysis exploits a unique survey of French firms known as the Enquete Trimestrielle de Conjoncture dans l'Industrie (ETCI; the English translation is "Quarterly Survey of Economic Conditions in the Industry"). This survey is managed and implemented by the French economic statistics institute (Institut National de la Statistique et des Etudes Economiques (INSEE)) and it is a part of surveys conducted by national statistical offices for the European Commission. Micro-data from the survey is available to researchers after approval from the INSEE and via a restricted access to a secure data hub (Secure Data Access Center – CASD).

This specific survey has been ongoing on a quarterly basis since 1992 and is conducted via postal mail or internet.² It covers firms in the French manufacturing sector, which accounted for 17% of total employment in France (on average between 2010 and 2018). The sample of firms is meant to be nationally representative, excluding small firms of less than 20 employees. Every quarter about 4,000 firms are sampled but firms with more than 500 employees or firms with revenues higher than 150 million euros are all surveyed. Approximately 10,000 firms participated in the survey over our

² <https://www.insee.fr/en/metadonnees/source/operation/s1498/presentation> for a full description of the methodology of the survey.

sample period, with on average 2,500 firms reporting per quarter. While participation in the survey is not mandatory, response rates are very high, more than 60% on average (since the sample size is about 4,000 firms). Larger firms are over-represented in the sample: the average number of employees by firm is about 450 whereas the median is only 150. Sampling weights are available to ensure that the sample is representative. Total employment by firms in the survey is approximately 1 million, which represents about one third of total employment in the manufacturing sector. The long *panel* dimension of the survey allows us to follow firms over extended durations of time: on average, a firm is present in the sample over a period of 7 years. This is especially true for larger firms. In addition, the survey asks firms questions about the overall firm but also about their main products. Our data set contains about 16,000 different products over the sample period and the median number of products for a given firm is 2. We also have information on the share of revenues coming from exports: on average, the mean share of exports is about 25% but for about one third of products the share of exports over sales is less than 5% (see Appendix Table 1). Overall, our data set contains more than 360,000 individual product-specific observations (time×firm×product) and approximately 270,000 firm-level observations (time×firm).

Surveys are meant to be filled out by top executives in the firm. To ensure high response rates, the survey is deliberately designed to be easy for these executives to fill out. Respondents are asked a variety of mostly qualitative questions about their firm and broader economic conditions.³ The survey questionnaire is reported in Appendix B. The scope of questions is quite extensive, covering areas such as prices, employment, production, wages, factors constraining production, the economic outlook, etc. In contrast to other firm surveys (e.g., IFO in Germany) in the European Commission framework, this French survey contains *qualitative* questions not only about firm-level outcomes and projections but also *quantitative* questions on firm-level variables (e.g., percent changes in prices) and *qualitative* questions on aggregate expectations (price, production, export and wages), a critical element for our analysis.⁴

While most questions are qualitative, the survey does include several quantitative questions. For example, firms are asked whether/how they changed their prices over the last three months, including in both qualitative and quantitative form. In the survey, prices can be provided for different products among the main products sold by the firm. All firm products are classified in the CPF/CPA

³ The French survey is part of the harmonized European Commission framework of business surveys since 2004. A majority of questions asked are common over different surveys across EU countries.

⁴ A recent exception is Dovern, Muller, and Wohlrabe (2020) in which they combine German business survey info (IFO) with new questions on aggregate expectations on GDP, but this survey covers only three quarters between 2018 and 2019.

2008 classification at level 4, there is a direct mapping of this product classification with the classification of firms into sectors (NACE classification).

Table 1 provides descriptive statistics on the answers to price questions. Overall, the average quarterly price change is 0.08%, implying an annual rate of about 1% whereas the average PPI quarterly inflation is 0.18% (when excluding energy and food prices). In a typical quarter, about one third of firms adjust their prices, which is consistent with Gautier (2008) and Vermeulen et al. (2012) documenting frequency of price changes for French and euro area PPI and with Berardi, Gautier and Le Bihan (2015) documenting the frequency of price changes for French CPI. Among price changes, two thirds of price changes are increases with the average price change being about 3%. Figure 1 (Panel B) plots the average price change reported over the last three months across firms in each wave of the survey since 1992 (with sampling weights), as well as the official PPI inflation rate for France for comparison. The two line up quite closely, indicating that the survey is indeed fairly representative and firms are providing factually correct answers about their price changes.

Respondents are also asked about their expected price changes over the next three months. The average expected price change is higher than the average past price change but the main average statistics are in line with what we obtain on past price changes. The time series for average responses for expected price changes are also plotted in Figure 1 (Panel A) with PPI inflation. The two also line up quite closely. The main swings in PPI inflation over time are well-captured in the survey answers. There is greater volatility in reported price changes from the survey than in PPI inflation. Figure 1 (Panels C and D) also reports the comovement of PPI inflation with the share of firms reporting that they increased prices in the last three months or expect to increase prices over the next three months. Here and henceforth, we compute these shares as the number of firms reporting an actual (expected) price increases divided by the number of firms reporting any actual (expected) price change.⁵ The correlation with PPI inflation is very high, indicating that using qualitative responses is informative.

Because firms' reported expected price changes will play an important role in our analysis, we want to ensure that the quality of these expectations data is high. The strong correlation between these average expectations and the time series of PPI inflation is consistent with this. Figure 2 presents additional evidence supporting the quality of these data. It plots the binned scatter plot of answers to expected price changes versus the actual price changes by firms. There is a strong positive relationship between firms' anticipated price changes and ex-post actual price changes. This is consistent with

⁵ The results are similar when we use the balance, i.e., the share of firms reporting an actual (expected) price increase minus the share of firms reporting an actual (expected) price decrease.

other survey evidence on firms' expected price changes (e.g. Coibion, Gorodnichenko and Kumar (2018) for firms in New Zealand, Coibion, Gorodnichenko and Ropele (2020) for firms in Italy) and how closely these line up with ex-post price changes.

Another quantitative question asked of firms is by how much did hourly wages of their workers change over the last quarter. While we will not focus on firms' wage changes, this provides another metric for assessing whether firms are providing high-quality answers to survey questions. When the average wage growth reported by firms in the survey is compared to official estimates of wage growth in manufacturing (Panel C in Appendix Figure 1), we find that both series display the same strong seasonal pattern, as well as very similar time series variation at lower frequencies. Again, this supports the notion that survey answers are in general of high quality. Another set of questions for firms focuses on their production levels, both past and future. Unlike pricing questions however, these are only qualitative in nature with firms being asked to state whether their production is "higher", "lower" or "about the same". Since few firms report declining levels of production, we also examine the time series for the average fraction of firms reporting that their production increased over the last three months as well as the fraction of firms reporting that they expect their production to increase over the next three months. Both series track the aggregate measure of manufacturing production in France very closely, even though the quantities in the two series cannot be directly compared (Panels A and B in Appendix Figure 1). We find similar results for the share of firms expecting an increase in aggregate inflation, output, wages and export growth (Appendix Figure 2). This close alignment between official statistics and the survey results again confirm that expectations data (even if they are qualitative) from this survey are informative.

3. Firms' Beliefs and Actions in Response to Industry and Aggregate Shocks

A key stylized fact about pricing behavior documented in previous work (e.g., Boivin et al. 2009) is that price rigidity is high after aggregate shocks but is much weaker after industry-specific shocks. Specifically, using a factor decomposition of industry-level prices, Boivin et al. (2009) show that aggregate shocks have very persistent effects on prices for most industries, consistent with the gradual response of the aggregate price level to monetary shocks documented using aggregate time series (e.g., Christiano, Eichenbaum and Evans 2005). However, shocks that are specific to industries, identified using factor decompositions of the panel of industry price levels, are followed by a much more rapid adjustment of prices, consistent with the micro evidence on the frequency of price

adjustment (Bils and Klenow 2004). Because industry-level shocks explain most of the volatility in prices over time, this finding can reconcile the micro-level evidence on rapid price adjustment with the macro-level evidence of gradual price adjustment.

This fact still requires an explanation however: why would prices respond so differentially to industry vs. aggregate shocks? Building on earlier work in Mackowiak and Wiederholt (2009), Mackowiak et al. (2009) provide one potential explanation for this using a model of rational inattention in which firms optimally choose how much of their limited information processing capacity to devote to learning about industry shocks vs aggregate shocks. In their frameworks, both shocks affect a firm’s ideal price (under full-information) but information frictions prevent firms from learning fully about these shocks. Firms must instead optimally allocate their information across the different shocks in whatever way maximizes profits. Mackowiak et al. (2009) note that since industry level shocks are so much more volatile empirically, firms should optimally choose to devote more of their information capacity to learning about shocks to their industries than about aggregate shocks. In other words, firms should rationally choose to be inattentive to aggregate shocks. Such a division of attention by firms would then naturally imply that prices should respond more rapidly to industry shocks than to aggregate shocks. They provide empirical evidence using pricing dynamics of different sectors consistent with rational inattention motives.

What remains to be established however is whether firms really do track information about industry shocks more closely than they do about the aggregate. Doing so requires direct measures of firms’ expectations to assess the relative attention that they devote to different types of shocks. The extensive panel of firms’ expectations in the French survey of firms therefore presents an ideal setting to directly test the hypothesis that inattention is greater after aggregate shocks than industry ones.

To assess how firms’ expectations respond, we use local projections to trace out the dynamic response of firms’ expectations of their own future price changes to both industry and aggregate shocks. Specifically, we regress:

$$\sum_{k=0}^h E_{t+k}^{i,j} dp_{t+k+1}^{i,j} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + error, \quad (1)$$

where $E_{t+k}^{i,j} dp_{t+k+1}^{i,j}$ is the expectation at time $t + k$ of firm i in industry j for their price changes over the next quarter, π_t^{agg} denotes aggregate inflation (aggregate producer price index for manufacturing goods (production sold in France)), and π_t^j is inflation in industry j for each horizon h (in our baseline regression the sector is defined at the 4-digit level of the product classification (CPF/CPA) which is the most disaggregated level for PPI indices (a little more than 150 different sectors), for each product,

we compute sectoral inflation using 4-digit product price indices (domestic market).⁶ The dependent variable is the cumulative sum of (expectations of) price changes over time, and therefore the coefficients β_h and γ_h trace out the response of expected level of prices over time. Note that these expectations of price changes are based on quantitative questions, so the response provides a quantitative estimate of cumulative expected price changes. We do not attempt to identify the structural sources of innovations to either aggregate or industry inflation. Instead, we view this simple specification as providing a tractable, “model-free” approach to characterizing the dynamics of expectations after each type of innovation, regardless of the fundamental source of these innovations. The fact that aggregate inflation is included as a regressor implies that γ_h identifies the response of expectations to industry-level inflation that is orthogonal to contemporaneous changes in aggregate inflation, which is similar in spirit to the factor decomposition used in Boivin et al. (2009). The baseline specification includes firm-specific fixed effects to capture the unobserved time-invariant firm characteristics that can affect average pricing behavior. To account for cross-sectional and time series correlation of the error term, we cluster standard errors by time and firm. We later consider a wide range of robustness checks to this baseline specification.

We present results in Panel A of Figure 3. In response to aggregate inflation, we find a relatively muted and delayed response of firms’ expectations of their subsequent price changes. The pricing response is gradually increasing over the first four quarters and continues on another few quarters. In other words, we find a similar delayed response of expected price changes in response to aggregate inflation as was documented by Boivin et al. (2009) for aggregate prices. This gradual adjustment of firms’ expectations to aggregate innovations in inflation is also consistent with empirical evidence that documents pervasive information rigidities after aggregate shocks, be it for professional forecasters, firms, households or central bankers (Coibion and Gorodnichenko 2012).

Panel A of Figure 3 also plots the dynamic response of firms’ expected price changes to industry level variation in inflation. In contrast to what we find with aggregate inflation, the response of firms’ expected price changes to industry-level variation is much more rapid. There is very little building up of firms’ price changes over time: instead, large price changes happen contemporaneously with the shock, continue for a few quarters, then rapidly begin to dissipate. We view this as providing direct evidence for the mechanism described in Mackowiak et al. (2009): firms’ expectations adjust

⁶ All product-level price series are available from INSEE website (<https://insee.fr/en/statistiques/series/108665892>).

more rapidly to industry shocks than to aggregate shocks, which provides a direct explanation for the well-known empirical facts reported in Boivin et al. (2009).⁷

We can also identify industry versus aggregate shocks using quantities produced rather than prices. We do so using the following specification:

$$\sum_{k=0}^h E_{t+k}^{i,j} dp_{t+k+1}^{i,j} = \alpha_{i,h} + \beta_h \Delta y_t^{agg} + \gamma_h \Delta y_t^j + error. \quad (2)$$

This specification is identical to equation (1) except for the fact that we now use aggregate production (Δy_t^{agg}) (measured using the aggregate French industrial production index (seasonally and working days adjusted)) and industry production (Δy_t^j) growth rates (measured using sectoral industrial production indices at the 4-digit NACE classification (seasonally and working days adjusted)) as regressors.⁸ Results are in Panel B of Figure 3 and are qualitatively similar to what we find in Panel A, with expected prices responding relatively more rapidly in response to industry fluctuations in output than aggregate ones, although the difference is somewhat more muted than when identification is based on prices.

We find similar results when we use firms' qualitative responses to price expectations. Specifically, before being asked to provide quantitative answers about their future price changes, respondents were asked whether their prices were going to rise, fall or stay the same. We assign values to each of 1, -1 and 0 respectively and use the same cumulative sum approach for the dependent variable as with prices. Checking whether these qualitative responses are consistent with quantitative ones is useful because many other expectations in the survey are measured only using qualitative questions. Since price expectations are asked both ways, they provide a particularly useful way to assess whether the two approaches provide similar results. Specifically, we now estimate

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} dp_{t+k+1}^{i,j}\} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + error \quad (1')$$

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} dp_{t+k+1}^{i,j}\} = \alpha_{i,h} + \beta_h \Delta y_t^{agg} + \gamma_h \Delta y_t^j + error \quad (2')$$

where $\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} dp_{t+k+1}^{i,j}\}$ is the cumulative sum between time t and $t+k$ of firms' qualitative answers about their expected price changes. We report results in panels A and B of Figure 4 for equations (1') and (2') respectively. Strikingly, the results are indistinguishable from those in Figure

⁷ We find similar results when we consider the (qualitative) response of firms' expected and actual output change to industry versus aggregate shocks (Appendix Figures 3 and 4).

⁸ All industrial production indices are available on the INSEE website <https://insee.fr/en/statistiques/series/108695026>

3, with only the scales being different. Qualitatively, we find the same faster reaction of firms' price expectations to industry variation relative to aggregate variation whether expectations are measured using quantitative or qualitative questions.

While these results focus on the speed of the response of expectations, we can also assess whether similar results obtain when looking at actual price changes. This can help verify two issues. The first is whether changes in expectations of future price changes are reflected in ex-post price changes, not just unconditionally (as in Figure 2) but also conditional on industry and aggregate shocks. Second, it allows us to verify that the stylized facts of Boivin et al. (2009), namely that industry prices respond more rapidly to industry shocks than aggregate shocks, also hold in France during our time sample. We implement this test using an equivalent empirical specification applied to actual price changes using either inflation or production to measure industry vs. aggregate variation:

$$\sum_{k=0}^h dp_{t+k}^{i,j} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + error, \quad (3)$$

$$\sum_{k=0}^h dp_{t+k}^{i,j} = \alpha_{i,h} + \beta_h \Delta y_t^{agg} + \gamma_h \Delta y_t^j + error. \quad (4)$$

We report results in Figure 5. Results are qualitatively similar to those found for expectations. Hence, we find that ex-post price changes are consistent with the ex-ante beliefs reported by firms not just unconditionally but also conditional on identifying industry and aggregate fluctuations. These findings also confirm the results of Boivin et al. (2009) that industry shocks lead to more rapid responses of prices than aggregate ones.⁹

Jointly, these results confirm a central prediction of rational inattention models, namely that firms should be more attentive to conditions that affect their profits most, which in this case consist of industry over aggregate conditions as proposed in Mackowiak and Wiederholt (2009). This

⁹ Since we can estimate the response of both actual prices and expectations to industry and aggregate shocks, we could also estimate the response of ex-post forecast errors to these same innovations, as in Coibion and Gorodnichenko (2012). They argue that the response of forecast errors to identified shocks can identify the degree of information rigidity in expectations. In principle, one could implement a similar test here to assess whether information frictions are higher for aggregate variation than industry variation, as would be the case in the rational inattention setting of Mackowiak and Wiederholt (2009). However, there is an important distinction between the two settings. In Coibion and Gorodnichenko (2012), the variable being forecasted is an aggregate one that is taken as evolving exogenously. Here, the variable being forecasted is the firm's own price. One can show analytically that the response of forecast errors with respect to one's own price, once normalized by the speed of the response of the variable being forecasted as in Coibion and Gorodnichenko (2012), is no longer a monotonic function of the degree of information rigidity but instead becomes a highly-nonlinear function of it. As a result, the speed of the response of forecast errors is not as directly informative about the degree of information rigidity as is the case with forecasts of aggregate variables. However, predictable forecast errors are consistent with information rigidity. We discuss this in more detail in Appendix C. We are very grateful to Mirko Wiederholt for making this suggestion.

differential attention across shocks implies that firms should respond more rapidly to industry-level shocks, as was documented empirically by Boivin et al. (2009). Furthermore, it also explains why firms might choose to be uninformed about aggregate conditions, a feature which is consistent with a wealth of empirical data on the aggregate expectations of firms and other economic agents (Coibion and Gorodnichenko 2012, 2015, Coibion et al. 2018).

4. Firms Aggregate Beliefs after Industry Shocks

The previous section focused on how rapidly firms learn about and respond to different shocks, particularly as viewed through the prism of rational inattention theories. A broader issue is understanding how firms *interpret* different kinds of innovations. For example, can they distinguish between shocks that have aggregate effects and those that don't? In this section, we provide additional evidence that firms set prices under an imperfect knowledge of their underlying fundamentals, in particular their aggregate or sectoral nature.

A large class of island models, following the celebrated model of Lucas (1972), posits that firms are inherently unable to distinguish between different shocks and are therefore likely to confound industry and aggregate shocks. In these models, agents are located on separate islands and trade with only a subset of islands in the economy. From these trades, they can in general observe prices on the other islands but they cannot observe the entirety of what is happening in the economy. These models capture the intuitive conundrum of firms that observe some other firms raising their prices: is this happening because of an aggregate shock or because of something specific to these few firms. The uncertainty about underlying forces in such an environment induces firms to put some weight on the possibility of an aggregate shock and some weight on the alternative possibility of idiosyncratic or industry factors, leading to a muted reaction of their own prices. Is the data consistent with this type of confusion or is it more consistent with the view that firms understand the nature of different shocks but are simply unable to observe each of them fully at all times?

The French survey ETCI is also remarkably well-suited to answer this closely related question. As firms are asked about what they expect to happen to aggregate prices and production along with those for their own firms, we can in principle distinguish what they believe will happen to them vs what they think is happening to the aggregate economy. Specifically, we can assess whether their expectations about the aggregate respond to both aggregate shocks as well as industry-specific ones *that have no aggregate effects*. Finding that firms' beliefs about the aggregate change in response to

industry-specific shocks would be direct evidence for the type of confusion about underlying shocks that is the key mechanism in island models.

To implement this test, we run regressions of the same type as before. Specifically, we regress for each time horizon h :

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} \pi_{t+k+1}^{agg}\} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + error \quad (5)$$

where $E_{t+k}^{i,j} \pi_{t+k+1}^{agg}$ represents the expectations at time $t+k$ of firm i in industry j over aggregate inflation (π) over the subsequent quarter, $\mathbb{I}\{\cdot\}$ takes values $\{-1,0,1\}$ for aggregate prices expected to decrease, stay the same, and increase respectively. Thus, rather than tracking firms' expectations of their own price changes as in section 3, we now characterize the dynamics of their expectations of *aggregate* prices in response to changes in industry vs aggregate inflation. As noted in section 2 however, the expectations of aggregate inflation are only qualitative in nature. The cumulative summation in the LHS can still be interpreted as speaking to the degree to which expectations respond to each form of inflation, but the quantitative values do not have a direct interpretation.

The results of these regressions are plotted in Panel A of Figure 6. This panel includes the estimated β_h , which indicate the response of firms' aggregate inflation expectations to aggregate inflation variation, and the estimated γ_h , which indicate the response of firms' aggregate inflation expectations to industry inflation variation. An increase in aggregate inflation is followed by a pattern of gradually increasing aggregate inflation expectations, as one would expect. With industry-level inflation, we also find a gradually increasing response of inflation expectations. Note that because we explicitly control for aggregate inflation, the changes in industry-specific inflation that are identified from γ_h are ones that are orthogonal to aggregate inflation and therefore should not be associated with changes in the aggregate beliefs of agents under full-information. Instead, we find precisely the result that underlies island models like Lucas (1972): the beliefs of agents about aggregates respond even to innovations that have no aggregate effects.

This result is quite robust and holds under a number of alternative specifications and identifications. For example, Panel B plots the equivalent responses of firms' aggregate inflation expectations to aggregate vs industry shocks but identifying the latter using output measures. That is, we replace aggregate inflation and industry inflation in equation (5) above with aggregate output and industry output:

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} \pi_{t+k+1}^{agg}\} = \alpha_{i,h} + \beta_h \Delta y_t^{agg} + \gamma_h \Delta y_t^j + error \quad (6)$$

We find equivalent results: changes in industry output that are orthogonal to aggregate output affect firms' aggregate inflation expectations. Appendix Figure 5 presents equivalent results for firms' expectations of future aggregate output. Panel A documents that these respond to changes in industry inflation orthogonal to aggregate inflation, while Panel B documents that firms' expectations of aggregate output respond also to changes in industry output that are orthogonal to aggregate output changes.

These results are robust to a variety of alternative specifications and additional controls. For example, in Panel B of Figure 7, we present estimates of equation (5) augmented to included controls for four-digit industry and aggregate production growth rates: we find almost identical effects. In Panel C of Figure 7, we plot the same impulse responses but also controlling for inflation and output at the two-digit industry. Again, we find the same qualitative dynamics of firms' aggregate inflation expectations to industry inflation. Panel D of Figure 7 includes time fixed effects along with four-digit industry production growth rate. Time fixed effects soak up all aggregate variation, not just that coming from aggregate inflation. We continue to find a significant response of firms' aggregate inflation expectations to industry inflation. Panel E of Figure 7 further augments the specification by also controlling for lags of four-digit industry inflation and production growth rates. Again, the results are qualitatively unchanged.¹⁰

Jointly, these results are strongly supportive of the mechanism underlying island models: firms confound shocks to industry and aggregate conditions. However, this interpretation of the empirical results hinges on whether industry variation in inflation really has no effect on aggregate inflation, making the apparent response of firms' aggregate inflation expectations at odds with underlying shocks. While our empirical specification restricts industry variation inflation to be orthogonal to aggregate conditions, it does not necessarily satisfy the restriction that contemporaneous variation in industry inflation has no effects on aggregate inflation in later periods, which could happen via e.g. input-output linkages.

To assess to what extent this possibility exists in our data, we run the following empirical tests on aggregate prices and output:

$$\sum_{k=0}^h \pi_{t+k}^{agg} = \alpha_j + \delta_{hj} \pi_t^j + \eta_{hj} \pi_t^{agg} + error \quad (7)$$

$$\sum_{k=0}^h \Delta y_{t+k}^{agg} = \alpha_j + \delta_{hj} \Delta y_t^j + \eta_{hj} \Delta y_t^{agg} + error \quad (8)$$

¹⁰ Appendix Figure 6 shows that our conclusions for Figure 3 are also robust to these variations.

These regressions assess whether subsequent changes in aggregate inflation and production are predictable using contemporaneous changes in *industry* inflation and output after conditioning on contemporaneous aggregate conditions. These regressions are run industry by industry for each forecasting horizon from 1 quarter to 8 quarters ahead. The results are reported in Appendix Table 2. At short horizons, there is some evidence of predictability of subsequent aggregate inflation from industry level variation in inflation: we can reject the null of no predictability for about one in four industries at the one-quarter horizon. Predictability is lower for output, with a little over ten percent of industries displaying predictability for subsequent aggregate changes in output at the one quarter ahead horizon. This predictability falls sharply at longer horizons: down to less than ten percent of industries at a horizon of two years for inflation and around two percent for output.

While this predictability in aggregates from industry-specific conditions is therefore limited, it nonetheless presents an alternative explanation for our empirical results. To assess whether this is behind the estimated responses of firms' aggregate expectations, we reproduce our baseline results dropping all industries for which we can reject the null of no predictability at the five percent level for a given horizon. The results for the response of firms' aggregate inflation expectations to industry inflation are presented in Panel A of Figure 8. The results are nearly indistinguishable from our baseline estimates.

We also verify the robustness of our empirical strategy by running a placebo test. Specifically, for a firm in industry j , we consider all other industries $s \neq j$ and regress, one by one, those industries' inflation rates on the inflation rate in industry j : $\pi_t^j = b_0^{(s)} + b_1 \pi_t^{(s)} + error$ where j, s index industries. We then identify the industry s^* that has the smallest value of $|b_1^{(s)}|$. This is the industry whose inflation rate has the least predictive power for inflation in industry j .¹¹ Given this other industry's inflation rate, we then reproduce our baseline estimates but replacing the inflation rate of the industry to which the firm belongs with the inflation rate of the industry with which it is least correlated:

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} \pi_{t+k+1}^{agg}\} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \psi_h \pi_t^{s^*} + error. \quad (9)$$

¹¹ We have similar results when we use $s^* = \arg \min \rho(\pi_t^j, \pi_t^{(s)})$. The advantage of the regression approach is that it does not depend on the variance of π_t^j and b_1 can be interpreted as the sensitivity (a unit increase in $\pi_t^{(s)}$ translates into b_1 unit increase in π_t^j).

We then plot the impulse response of firms’ aggregate inflation expectations to innovations in the inflation of these other industries in Panel B of Figure 8. What one would expect this placebo test to yield is an absence of predictive power on firms’ aggregate expectations: this is precisely what we find. This result illustrates that our finding of predictive power running from firms’ industries’ inflation to their aggregate inflation expectations is not an artifact of the empirical procedure but truly captures the fact that firms are forming beliefs about the aggregate based on what they observe in their own industries.

In short, we present novel direct evidence on two theoretical hypotheses. First, as described in section 3, firms devote relatively more information processing capacity to tracking industry conditions than aggregate conditions as displayed by the differential speed of response of their expectations to each. Second, this large weight that firms implicitly assign to variation originating in their industry shapes not just their expectations about their industry but also their expectations about the broader economy, as suggested by island models.

5. Heterogeneity

While these results largely conform to the predictions of rational inattention and island-type models in the aggregate, these models also make additional predictions in terms of cross-sectional heterogeneity. For example, under rational inattention models in which firms choose how much attention to allocate to industry vs aggregate shocks, attention to industry conditions should be greater when industry volatility is relatively high or when it is more persistent. By the same logic, one would expect firms’ extrapolation of industry conditions to aggregate conditions to be larger when their industry conditions are more volatile or more persistent.

In this section, we assess whether firms’ extrapolations of industry variation to their expectations about the aggregate vary along any observable characteristics.¹² To this end, we estimate equation (5) for each industry j separately and study variation in $\frac{1}{13} \sum_{h=0}^{12} \hat{\beta}_{j,h}$ and $\frac{1}{13} \sum_{h=0}^{12} \hat{\gamma}_{j,h}$ across industries. Table 2 documents that there is indeed quite a bit of cross-industry variation in how firms’ aggregate expectations respond to industry shocks on average. Then, we match the French survey with administrative balance-sheet data set (covering the universe of French firms), providing us with detailed annual information on the total wage bill, values of intermediate inputs (materials), and value

¹² We find very similar results when we study heterogeneity in the relative adjustment of firms’ expected prices to industry vs aggregate variation, so we focus on the response of aggregate expectations to industry shocks in the interest of space.

added to construct cost shares. As a result, the French survey and administrative data provide extensive information about firms and their industries, allowing for a relatively rich analysis of the amount of heterogeneity underlying our aggregated results in previous sections. For example, in addition to cost shares, we observe the importance of exports as a share of total sales, the number of employees firms have, how many products they sell, and their capacity utilization. Table 2 documents both average levels of these across all firms in the sample as well as some of the heterogeneity present in the data. We can also assess characteristics of the industries in which firms reside. For example, we can regress industry prices on aggregate prices to measure the degree of comovement of a specific industry with broader price movements. We also measure the degree to which an industry's price level comoves with commodity prices, providing a simple metric for the likelihood of more volatile prices. Finally, we also measure the volatility and persistence of an industry's price. We do so by running an AR(4) on each industry's inflation rate and measure persistence via the sum of the AR(4) coefficients and volatility via the standard deviation of residuals.

Figure 9 presents simple scatterplots that compare both industry price persistence and volatility to the average response within each industry of firms' aggregate price expectations to industry variation. Panel A shows that there is a weak unconditional negative relationship between the size of innovations to industry prices and the average response of firms' aggregate expectations. In contrast, Panel B documents a positive correlation between the persistence of industry prices with the response of aggregate expectations. While the latter is consistent with rational inattention type motives, the former is not. Of course, there could be many other firm and industry characteristics that affect firms' incentives to track aggregate vs industry conditions, and since the characteristics can be correlated with either persistence or volatility of industry prices, these unconditional correlations are only suggestive.

To assess the role of different characteristics in determining the speed of expectations adjustments, we run a sequence of cross-sectional regressions. Specifically, we first estimate the average response of aggregate expectations to industry variation for each industry (row 1 of Table 2 shows the distribution of the resulting cross-section). Second, we regress these industry level estimates of industry characteristics. For example, column 1 of Table 3 plots regressions of the response of aggregate expectations on the volatility of innovations to that industry's price level. The results confirm the scatterplot in Figure 3: there is a negative, albeit statistically weak, correlation between them. Column 2 presents the equivalent regression for industry price persistence, yielding a

strong positive relationship. When both variables are included in the same regression (column 3), the results for each are unchanged.

Column 4 augments this empirical specification with three additional control variables. The first is the average response of firms' aggregate expectations to aggregate inflation variation. Intuitively, this is to control for the possibility that firms in some industries face lower attention costs, and therefore pay more attention to all variables. This would make the interpretation of baseline regressions in columns 1-3 problematic since a higher response of firms' aggregate expectations to industry shocks could reflect not just an extrapolation property on the part of firms but also a more systematic higher elasticity of expectations to new information. We find no evidence for the latter: including the average response of firms' aggregate expectations to aggregate inflation does not affect the results. We also include the elasticity of industry prices to both aggregate inflation and commodity prices. We estimate these objects by regressing industry-level inflation on aggregate inflation and by regressing industry-level inflation on commodity price inflation. Including these additional controls adds no predictive power and does not affect the estimated coefficients on industry price persistence or volatility.

We then consider an additional set of industry characteristics, specifically the average cost shares of both labor and materials. Intuitively, a higher cost of materials could be indicative of more volatility in costs and prices, which could induce firms to allocate more attention to industry conditions relative to aggregates and therefore also induce to extrapolate more from their industry's prices to aggregate conditions. This is indeed what we observe: adding this variable (column 5 of Table 3) yields a positive and statistically significant coefficient. In addition, the coefficient on industry price volatility becomes statistically insignificant. In contrast, including the labor cost share does not add any predictive power for the average response of firms' aggregate expectations to industry price variation. As documented in columns 6 and 7 of Table 3, including the average share of exports, the number of products, firm size or capacity utilization has no effect on these results. The only characteristics that are robustly associated with the average response of firms' aggregate expectations to industry price variation are the persistence of industry prices and the share of materials in firms' costs. We interpret these results as broadly consistent with predictions of rational inattention models.

6. Conclusion

Recent work has increasingly turned to understanding how agents form their expectations and how those expectations affect economic decisions. Most of that work has focused on the aggregate expectations of agents and how those respond to aggregate shocks (e.g., Coibion and Gorodnichenko 2012). Much less effort has been devoted to understanding how expectations respond to more firm-specific factors. We provide two new empirical results in this spirit that should provide guidance to future research.

The first is that firms' expectations respond more strongly to industry shocks than aggregate shocks. This is consistent with rational inattention theories in which agents decide where to allocate their information processing capacity: aggregate shocks are much less volatile and affect profits less than industry or idiosyncratic shocks and therefore warrant less attention on the part of firms. The second is that firms' aggregate expectations respond to industry shocks that have no aggregate effects. This implies that firms confound underlying sources of volatility: they attribute some of the industry variation they observe to aggregate forces that in fact are not underlying their sectoral volatility. This is consistent with island models in the spirit of Lucas (1972) in which firms observe signals which are combinations of idiosyncratic or industry shocks and aggregate shocks.

While macroeconomic models in this spirit are not uncommon (e.g., Lorenzoni 2009, Angeletos and La'O 2013, Nimark 2014), empirical work testing this learning mechanism has been, to the best of our knowledge, non-existent. As our results strongly support this type of inference problem on the part of firms, we hope it will stimulate additional work on these channels. In addition, our empirical results can potentially provide a novel set of empirical facts that can be used to discipline this class of models.

More generally, these results also provide a potential lens through which to explain the puzzling amount of disagreement among firms about aggregate economic conditions. Rational inattention motives can explain why firms would devote little attention to aggregate conditions, meaning the signals they receive about the aggregate are very noisy. This noise implies that aggregate beliefs should respond little to these signals, consistent with the gradual response of firms' beliefs to aggregate shocks, but it also implies that disagreement about the aggregate need not be high since firms' beliefs will not respond strongly to these signals. But they will respond strongly to signals from their industries. If they then attribute an aggregate component to this information, as they seem to in section 4, then this informational response combined with the high level of industry volatility can potentially deliver a powerful quantitative force to explain the magnitude of cross-sectional

disagreement observed among firms. While we are not able to quantify this mechanism here due to the qualitative nature of the aggregate expectations in this survey, we hope future work will ascertain the quantitative importance of this channel.

More broadly, our results speak to the large divide between full-information macroeconomic models and the growing empirical evidence of pervasive information rigidities on the part of economic agents. In these models, expectations about the future adjust immediately to shocks and can provide a powerful propagation mechanism for even small aggregate shocks into the decisions of very forward-looking agents. Evidence of information frictions suggests that these expectational effects are likely much weaker, at least when it comes to macroeconomic shocks. Indeed, our results support a growing literature focusing on granularity and network structures in the economy (e.g. Gabaix 2011). This work has emphasized the potential importance of idiosyncratic shocks to specific firms in the economy who play a disproportionate role either through their size or network linkages. Our results imply that expectational forces are likely to be much stronger in response to these types of “local” shocks than they are in response to aggregate ones.

References

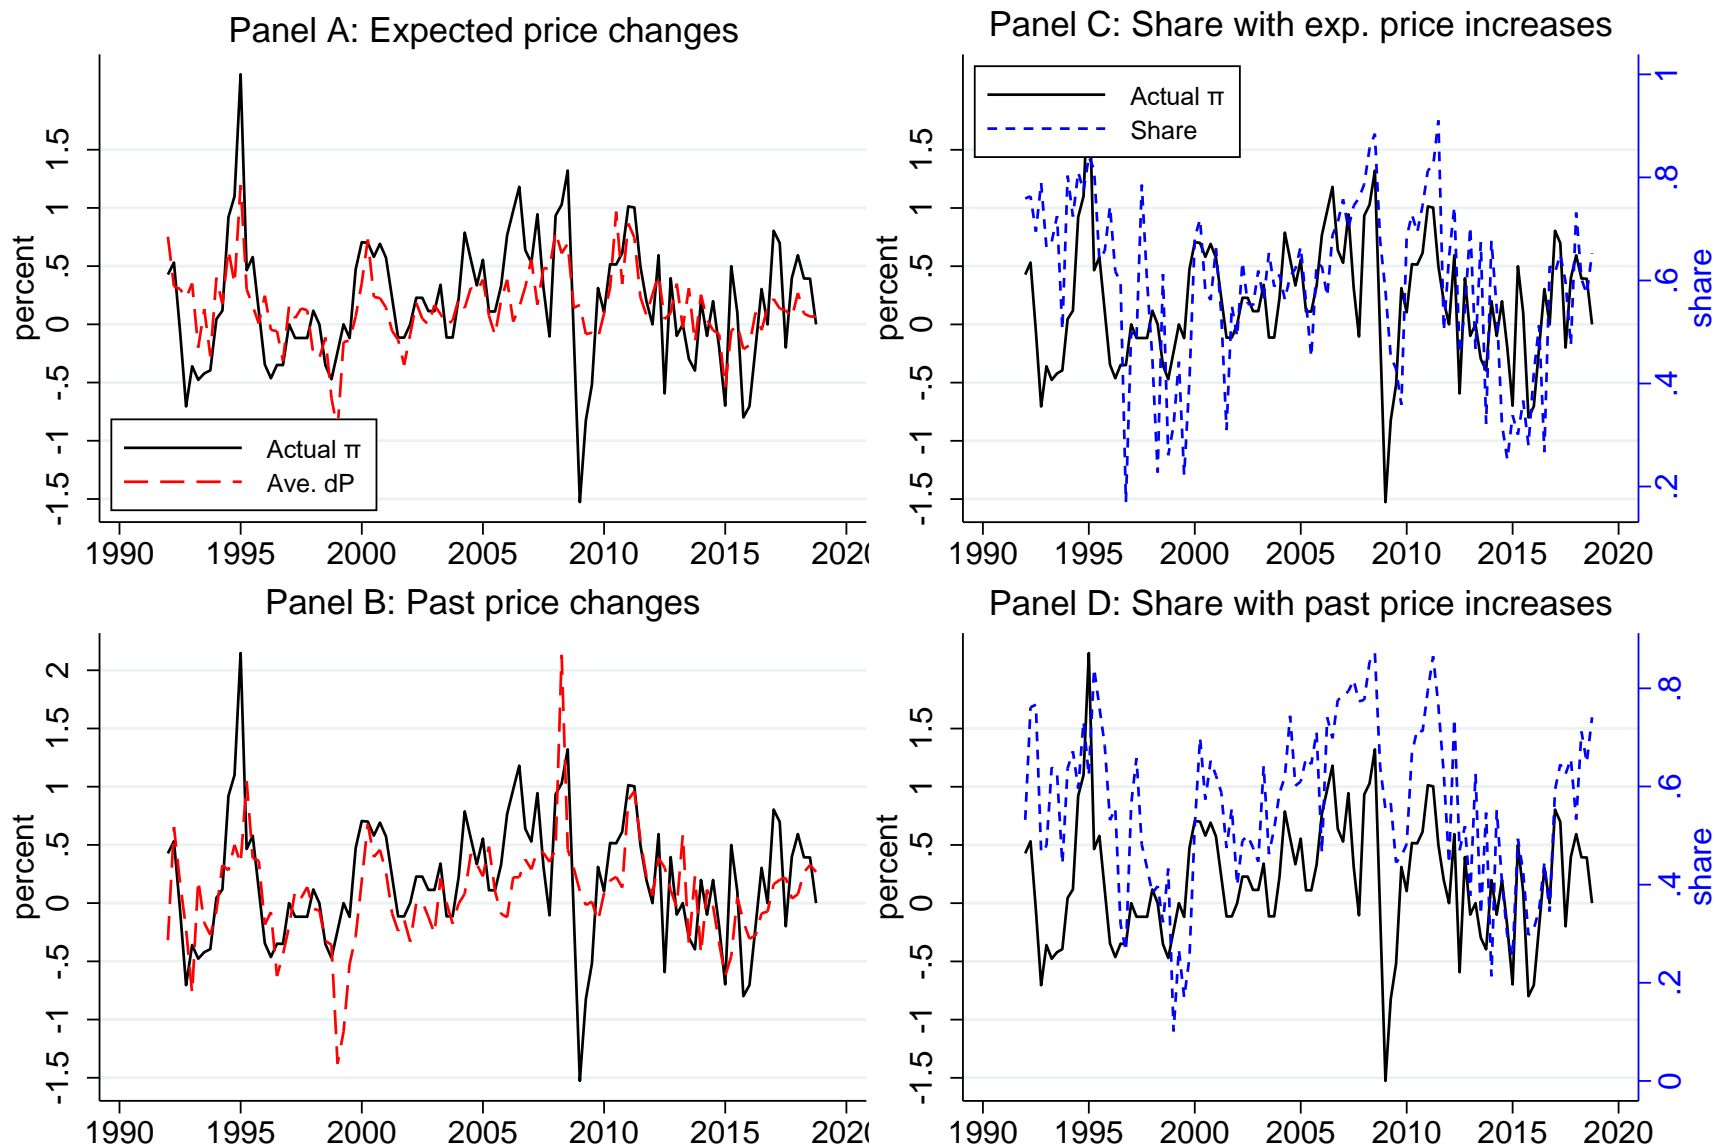
- Afrouzi, Hassan, 2016. “Strategic Inattention, Inflation Dynamics, and the Non-Neutrality of Money,” manuscript.
- Andrade, Philippe, and Hervé Le Bihan, 2013. “Inattentive professional forecasters,” *Journal of Monetary Economics* 60(8): 967-982.
- Angeletos, Georges-Marios, and Jennifer Lao, 2013. “Sentiments,” *Econometrica* 81(2): 739–779.
- Berardi, Nicoletta, Erwan Gautier, and Herve Le Bihan, 2015. “More Facts about Prices: France Before and After the Great Recession,” *Journal of Money, Credit and Banking* 47(8): 1465-1502.
- Bils, Mark, and Pete J. Klenow, 2004. “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947-985.
- Boivin, Jean, Marc Giannoni, and Ilian Mihov, 2009. “Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data,” *American Economic Review* 99(1): 350-84.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia, 2017. “Inflation Expectations, Learning and Supermarket Prices: Evidence from Survey Experiments,” *American Economic Journal: Macroeconomics* 9(3): 1-35.
- Chahrour, Ryan, Kristoffer Nimark, and Stefan Pitschner, 2019. “Sectoral Media Focus and Aggregate Fluctuations”, manuscript.

- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*.
- Coibion, Olivier, and Yuriy Gorodnichenko, 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120(1): 116-159.
- Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105(8): 2644-78.
- Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar, 2018. "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review* 108(9): 2671-2713.
- Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele, 2020. "Inflation Expectations and Firm Decisions: New Causal Evidence," *Quarterly Journal of Economics* 135(1): 165–219.
- Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte, 2020. "Inflation Expectations as a Policy Tool?" forthcoming in *Journal of International Economics*.
- D’Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, 2019. "Exposure to Daily Price Changes and Inflation Expectations," NBER Working Paper No. 26237.
- Dovern, Jonas, Lena Sophia Müller, and Klaus Wohlrabe, 2020. "How Do Firms Form Expectations of Aggregate Growth? New Evidence from a Large-Scale Business Survey," CESifo Working Paper Series 8179.
- Gabaix, Xavier, 2011. "The Granular Origins of Aggregate Fluctuations," *Econometrica* 79: 733-772.
- Gautier, Erwan, 2008. "The Behaviour of Producer Prices. Evidence from French PPI Micro-data," *Empirical Economics*, 35(2): 301-332.
- Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko, 2015. "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand," *Brookings Papers on Economic Activity* (Fall), 151-225.
- Lorenzoni, Guido, 2009. "A Theory of Demand Shocks," *American Economic Review* 99(5): 2050-84.
- Lucas, Robert Jr., 1972. "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4(2): 103-124.
- Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention," *American Economic Review* 99(3): 769-803.
- Mackowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt. 2009. "Sectoral Price Data and Models of Price Setting," *Journal of Monetary Economics* 56(S): 78-99.

Nimark, Kristoffer, 2014. "Man-Bites-Dog Business Cycles," *American Economic Review* 104(8): 2320-2367.


Vermeulen, Philip, Daniel Dias, Maarten Dossche, Erwan Gautier, Ignacio Hernando, Roberto Sabbatini, and Harald Stahl. 2012 "Price Setting in the Euro Area: Some Stylised Facts from Individual Producer Price Data," *Journal of Money, Credit and Banking* 44(8):1631-1650.

Figure 1. Expected and Actual Price Changes by Firms Compared to Aggregate Inflation.




Notes: Panel A plots actual quarterly PPI inflation (excluding food and energy) in France and the average expected price changes reported by firms in the survey. Panel B compares actual PPI inflation to the average share of firms reporting that they expect to raise prices in next three months. Panel C presents the PPI inflation and the average price changes over the last three months reported by firms. Panel D reports the fraction of firms that claim to have increased prices over the preceding three months in the survey.

Figure 2. Past and Expected Future Price Change Distribution.



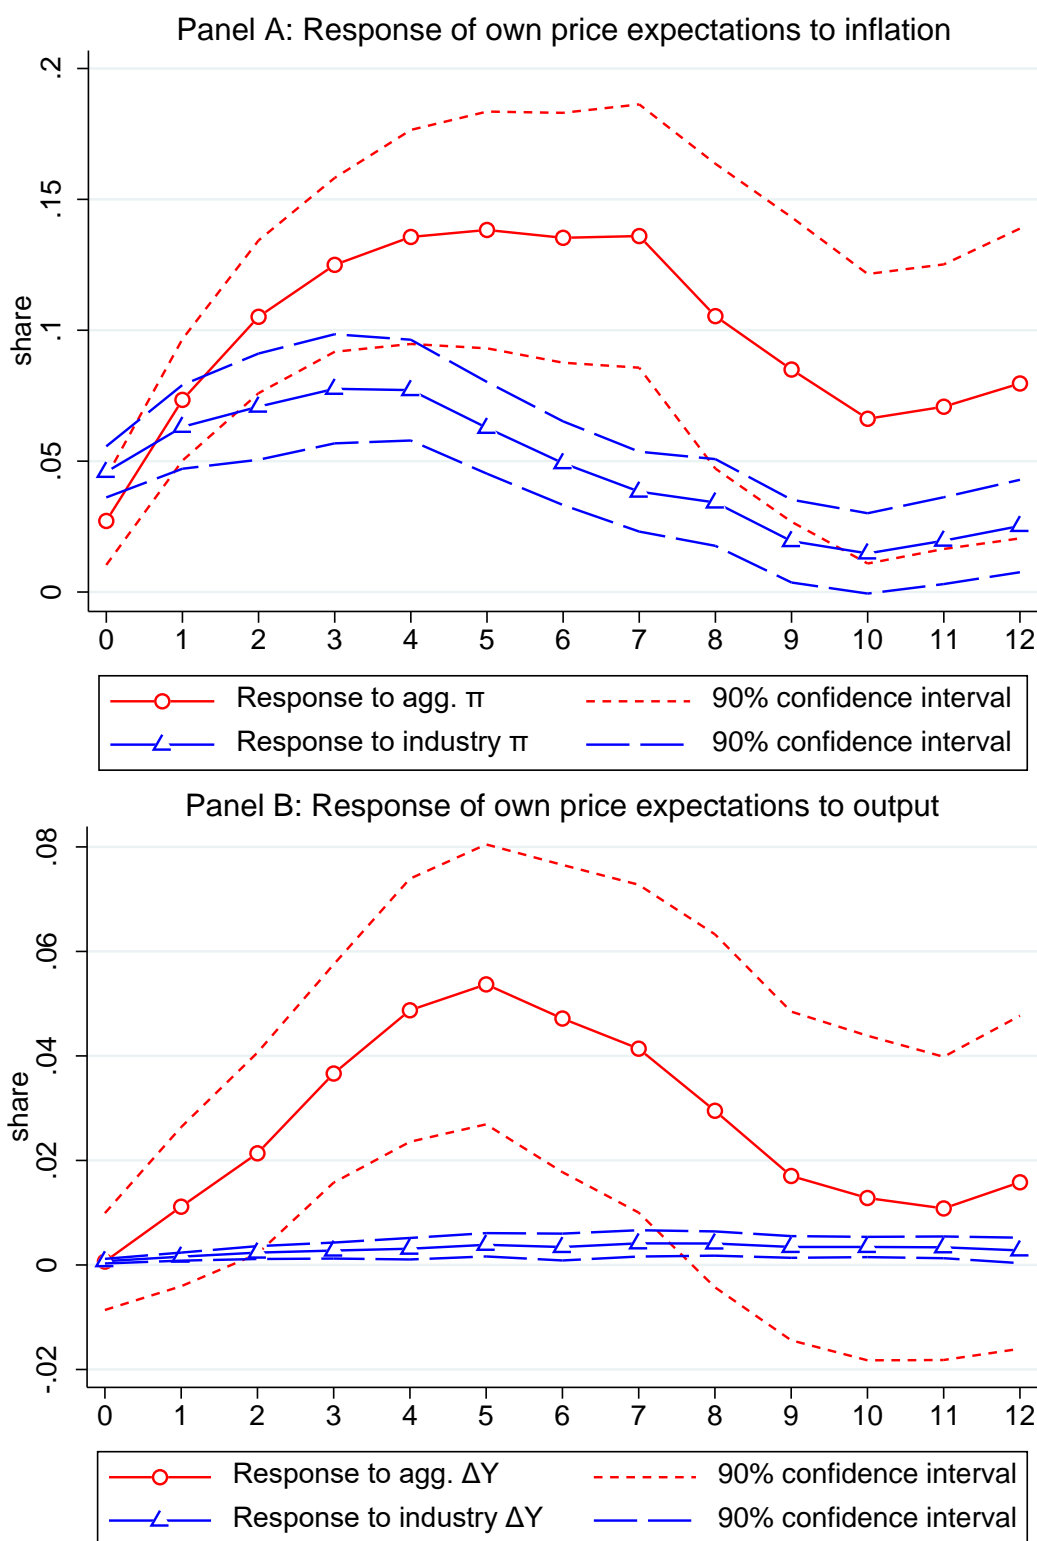
Notes: The figure plots a binscatter of expected future price changes over the following three months reported by firms across all quarters of the survey against ex-post actual price changes over the previous three months reported by firms across all quarters of the survey.

Figure 3. Response of Firms' Expected Price Changes to Industry vs. Aggregate Inflation and Output.



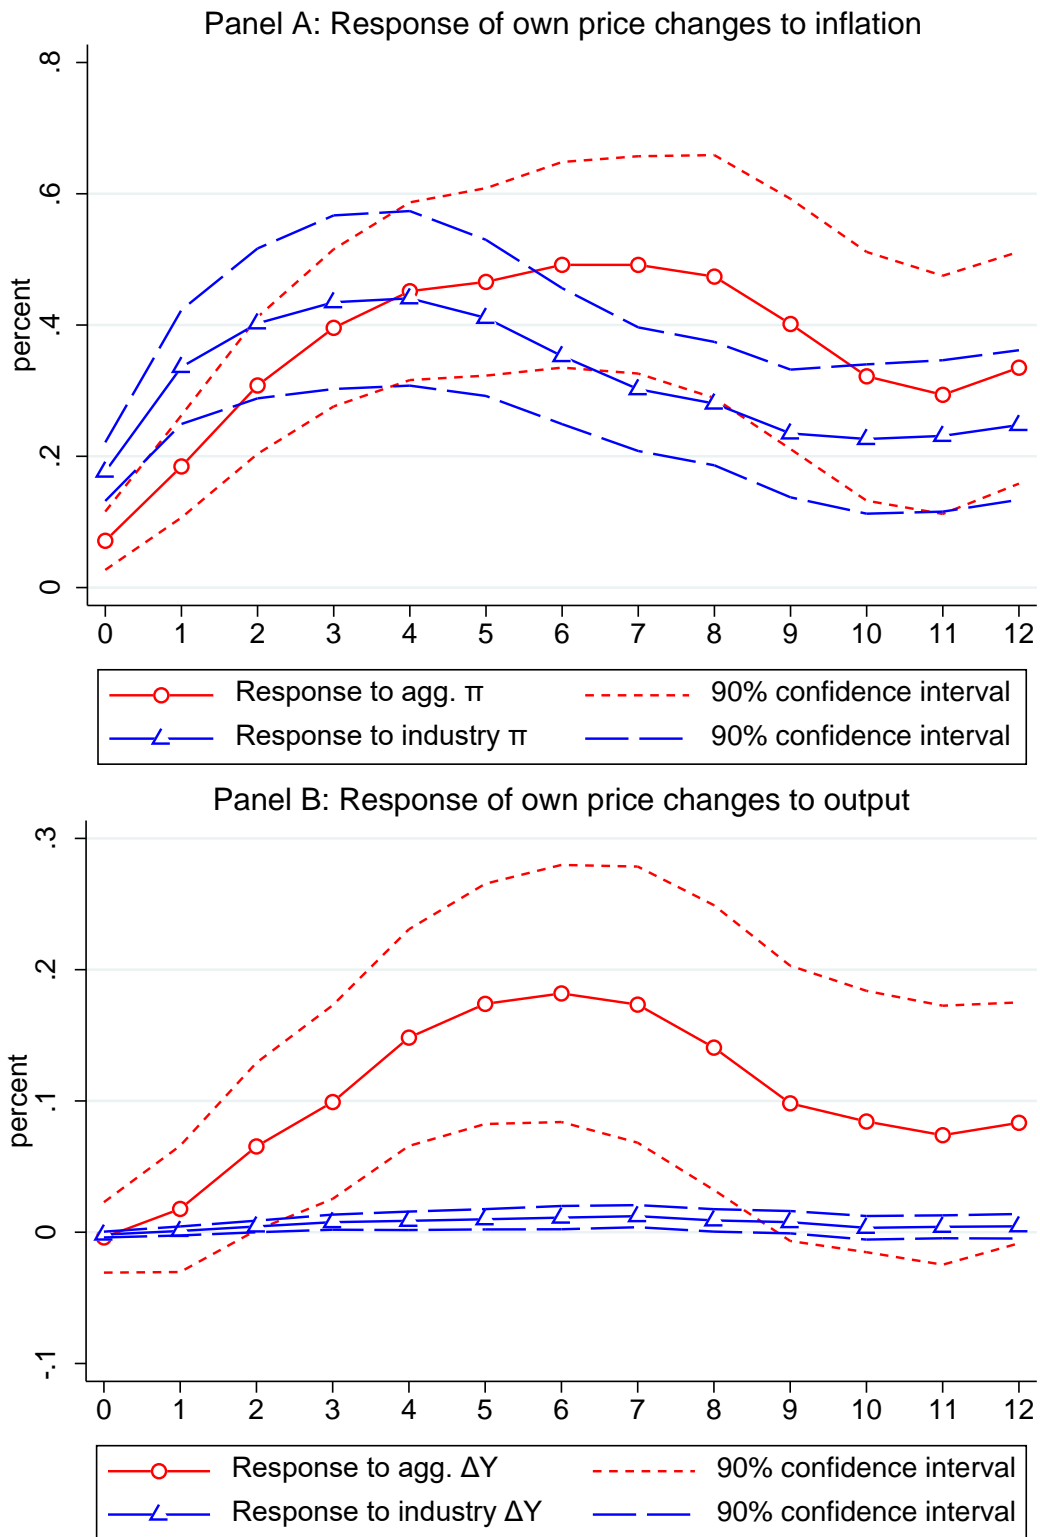
Notes: Panel A plots the cumulative response of firms' expectations of their own price changes to changes in aggregate inflation (blue line) and industry inflation (red lines); specification (1). Panel B plots the corresponding impulse responses for firms' expectations their own price changes to changes in aggregate output (blue line) and industry output (red lines); specification (2). The horizontal axis shows the impulse-response horizon measured in quarters.

Figure 4. Response of Firms' Expected (Qualitative) Price Changes to Industry vs. Aggregate Inflation and Output.




Notes: Panel A plots the cumulative response of firms' qualitative expectations of their own price changes to changes in aggregate inflation (blue line) and industry inflation (red lines); specification (1'). Panel B plots the corresponding impulse responses for firms' qualitative expectations their own price changes to changes in aggregate output (blue line) and industry output (red lines); specification (2'). The horizontal axis shows the impulse-response horizon measured in quarters.

Figure 5. Response of Firms' Actual Price Changes to Industry vs. Aggregate Inflation and Output.



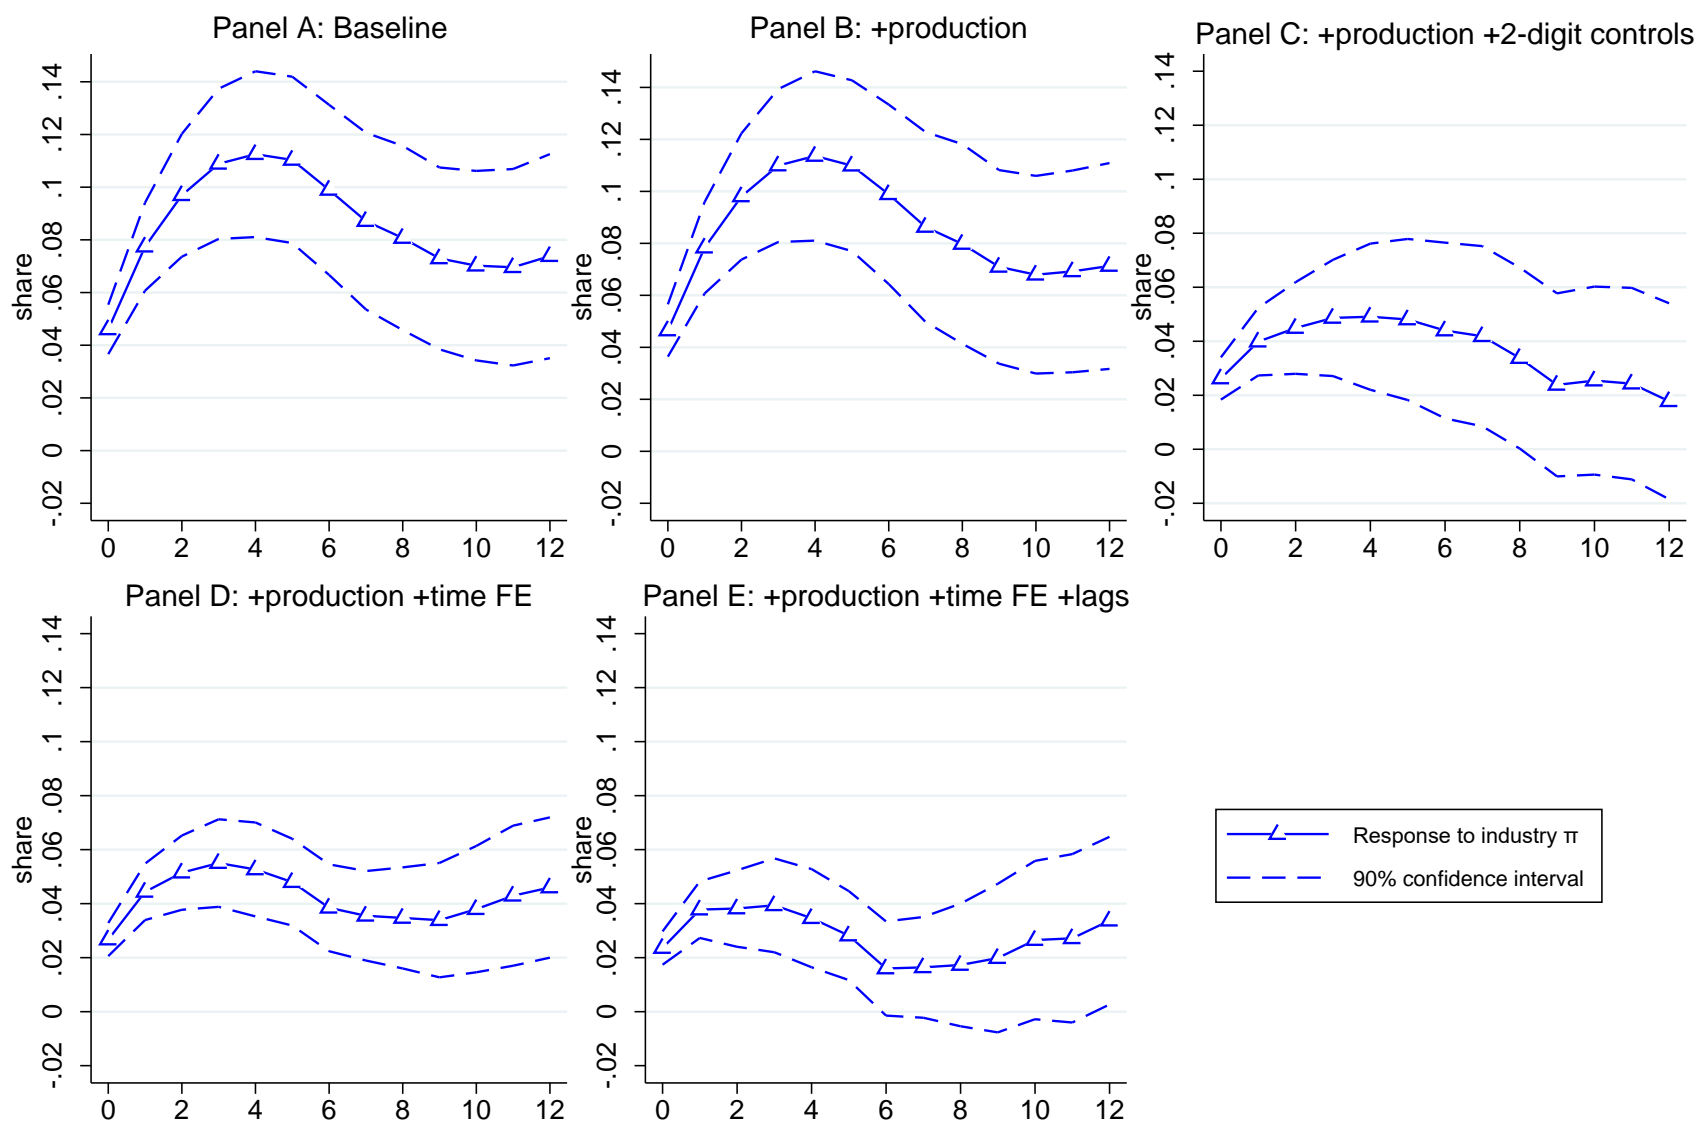
Notes: Panel A plots the cumulative response of firms' actual price changes to changes in aggregate inflation (blue line) and industry inflation (red lines); specification (3). Panel B plots the corresponding impulse responses for firms' actual price changes to changes in aggregate output (blue line) and industry output (red lines); specification (4). The horizontal axis shows the impulse-response horizon measured in quarters.

Figure 6. Response of Firms' Aggregate Inflation Expectations to Industry vs Aggregate Inflation and Output.




Notes: Panel A plots the cumulative response of firms' expectations of aggregate inflation to changes in aggregate inflation vs. changes in industry inflation; specification (5). Panel B plots the equivalent responses to changes in aggregate vs. industry output; specification (6). 95% confidence intervals are indicated by dotted lines in each panel. The horizontal axis shows the impulse-response horizon measured in quarters.

Figure 7. Robustness: responses of aggregate inflation expectations to industry-level inflation shocks.



Notes: Each panel plots the impulse response of firms' aggregate inflation expectations to industry variation in inflation. Panel A is the baseline specification (specification 5), panel B includes four-digit industry and aggregate production growth rate as an additional control, panel C adds two-digit industry inflation and output growth rate, panel D includes four-digit industry production and time fixed effects as controls, while panel E includes time fixed effects as well as current values and lags of four-digit industry production growth rate and inflation. The horizontal axis shows the impulse-response horizon measured in quarters.

Figure 8. Additional robustness checks.



Notes: Panel A plots the average response of all firms' aggregate inflation expectations to industry inflation variation as well as the response for firms that are specifically in industries for which industry inflation has no predictive power for subsequent periods of aggregate inflation; specification (5). Panel B again plots the average response of all firms' aggregate inflation expectations to industry inflation variation as well as a response of firms' aggregate inflation expectations to a placebo industry's inflation; specification (9). See section 4 for construction of placebo industries. The horizontal axis shows the impulse-response horizon measured in quarters.

Figure 9. Sensitivity of Aggregate Inflation Expectations to Volatility and Persistence of Industry-Level Inflation.

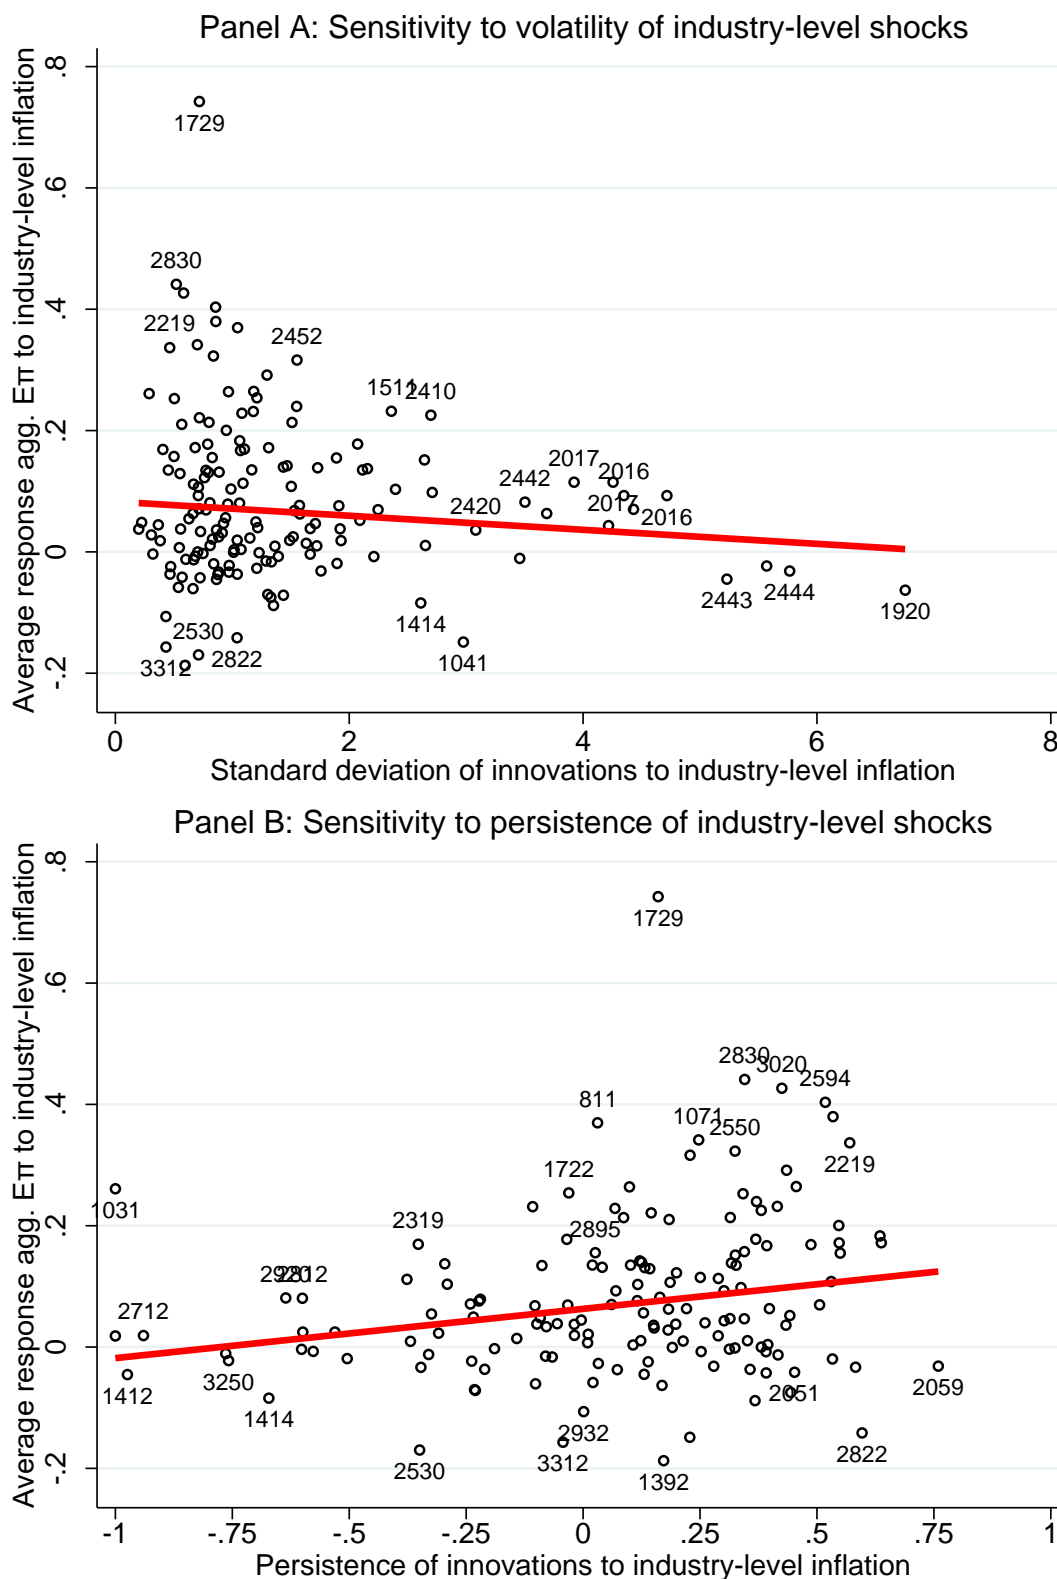


Table 1. Descriptive Statistics.

	Firm specific outcomes		Aggregate
	Expectations	Past	Expectations
	(1)	(2)	(3)
Panel A. Prices			
Average price change, %	0.15	0.08	-
Average non-zero price change, %			
Increase	3.04	3.15	-
Decrease	-3.21	-3.55	-
Price change, share (%)			
All	34.1	32.7	35.7
Increases	19.5	20.2	21.1
Decreases	14.6	12.5	14.6
Panel B. Production			
Change in production			
All	52.5	44.7	42.4
Increases	28.8	25.6	22.3
Decreases	23.4	19.1	19.1

Notes: The table reports descriptive statistics from the survey of French firms. Panel A focuses on price statistics, Panel B on statistics about production. For Panel A, statistics are provided for quantitative responses from firms about their price changes over the previous three months as well as their expected price changes over the next three months. All other statistics are based on qualitative responses regarding whether they expect variables to increase, decrease or stay the same.

Table 2. Industry heterogeneity, descriptive statistics.

	mean	Huber-robust mean	25 th percentile	50 th percentile	75 th percentile	St.Dev.
	(1)	(2)	(3)	(4)	(5)	(6)
Responses of expectations						
Average response of $E_i\pi^{agg}$ to $\pi^{industry}$	0.078	0.065	-0.007	0.050	0.142	0.134
Average response of $E_i\pi^{agg}$ to π^{agg}	0.281	0.265	0.159	0.271	0.376	0.182
Industry characteristics						
St.Dev. of innovations to $\pi^{industry}$	1.351	1.050	0.724	1.016	1.521	1.064
Persistence of $\pi^{industry}$	0.131	0.171	-0.048	0.173	0.377	0.349
Loading of $\pi^{industry}$ on π^{agg}	0.318	0.101	0.018	0.105	0.347	0.606
Loading of $\pi^{industry}$ on $d\log(P^{Comm})$	0.023	0.005	-0.004	0.005	0.025	0.068
Median labor cost share	0.493	0.502	0.443	0.503	0.567	0.104
Median material cost share	0.797	0.844	0.685	0.874	0.986	0.232
Median export share	0.297	0.290	0.110	0.262	0.456	0.209
Median number of products	1.312	1.227	1.089	1.217	1.413	0.332
Log(Median firm size)	4.920	4.907	4.636	4.910	5.225	0.606
Median capacity utilization rate	81.397	81.599	78.755	81.734	83.642	5.087

Notes: The first two rows are the average impulse response for each industry of firms' aggregate inflation expectations to industry inflation (Average response of $E_i\pi^{agg}$ to $\pi^{industry}$) or aggregate inflation (Average response of $E_i\pi^{agg}$ to π^{agg}). St.Dev. of innovations to $\pi^{industry}$ and Persistence of $\pi^{industry}$ are estimated using Huber-robust regression of AR(4) processes. Loadings are estimated Huber-robust regressions. Industry-level characteristics such as labor share, number of products, etc. are taken from the EITC survey. These characteristics are computed using sampling weights.

Table 3. Determinants of sensitivity of aggregate inflation expectations to industry-level inflation.

	Dependent variable: Ave. response of $E_i\pi^{agg}$ to $\pi^{industry}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
St.Dev. of innovations to $\pi^{industry}$	-0.015*		-0.017*	-0.015	-0.027	-0.024	-0.025
	(0.009)		(0.009)	(0.018)	(0.017)	(0.018)	(0.018)
Persistence of $\pi^{industry}$		0.087***	0.094***	0.089***	0.069**	0.068**	0.064**
		(0.031)	(0.031)	(0.032)	(0.030)	(0.031)	(0.031)
Ave. response of $E\pi^{agg}$ to π^{agg}				0.013	0.074	0.072	0.072
				(0.060)	(0.059)	(0.061)	(0.061)
Loading of $\pi^{industry}$ on π^{agg}				0.030	0.030	0.027	0.026
				(0.040)	(0.037)	(0.039)	(0.039)
Loading of $\pi^{industry}$ on $d\log(P^{Comm})$				-0.298	-0.208	-0.205	-0.223
				(0.279)	(0.260)	(0.279)	(0.277)
Labor cost share					0.055	0.060	0.050
					(0.095)	(0.101)	(0.101)
Material cost share					0.159***	0.159***	0.168***
					(0.044)	(0.045)	(0.046)
Export share						-0.022	-0.027
						(0.051)	(0.051)
Number of products						0.004	-0.005
						(0.055)	(0.057)
Log(Median firm size)						0.000	0.001
						(0.017)	(0.017)
Capacity utilization rate							0.002
							(0.002)
Observations	134	134	134	134	134	134	134
R-squared	0.021	0.057	0.088	0.092	0.192	0.190	0.196

Notes: The table plots results from cross-sectional regressions. The dependent variable is each industry's average response of firms' aggregate inflation expectations to variation in that industry's inflation rate after conditioning on aggregate inflation.

Appendix A
Additional Tables and Figures

Appendix Table 1. Sample characteristics.

	Number of employees	Sales (million €)	Share of exports in total sales, %	Duration of participation in the survey, years
Average	435	332.6	19.9	6.75
Percentiles				
P10	32	6.6	0	0.50
P25	60	17.5	0	1.75
P50	146	51.3	8.6	5.25
P75	350	165.5	31.8	10.0
P90	810	477.0	60.8	15.5

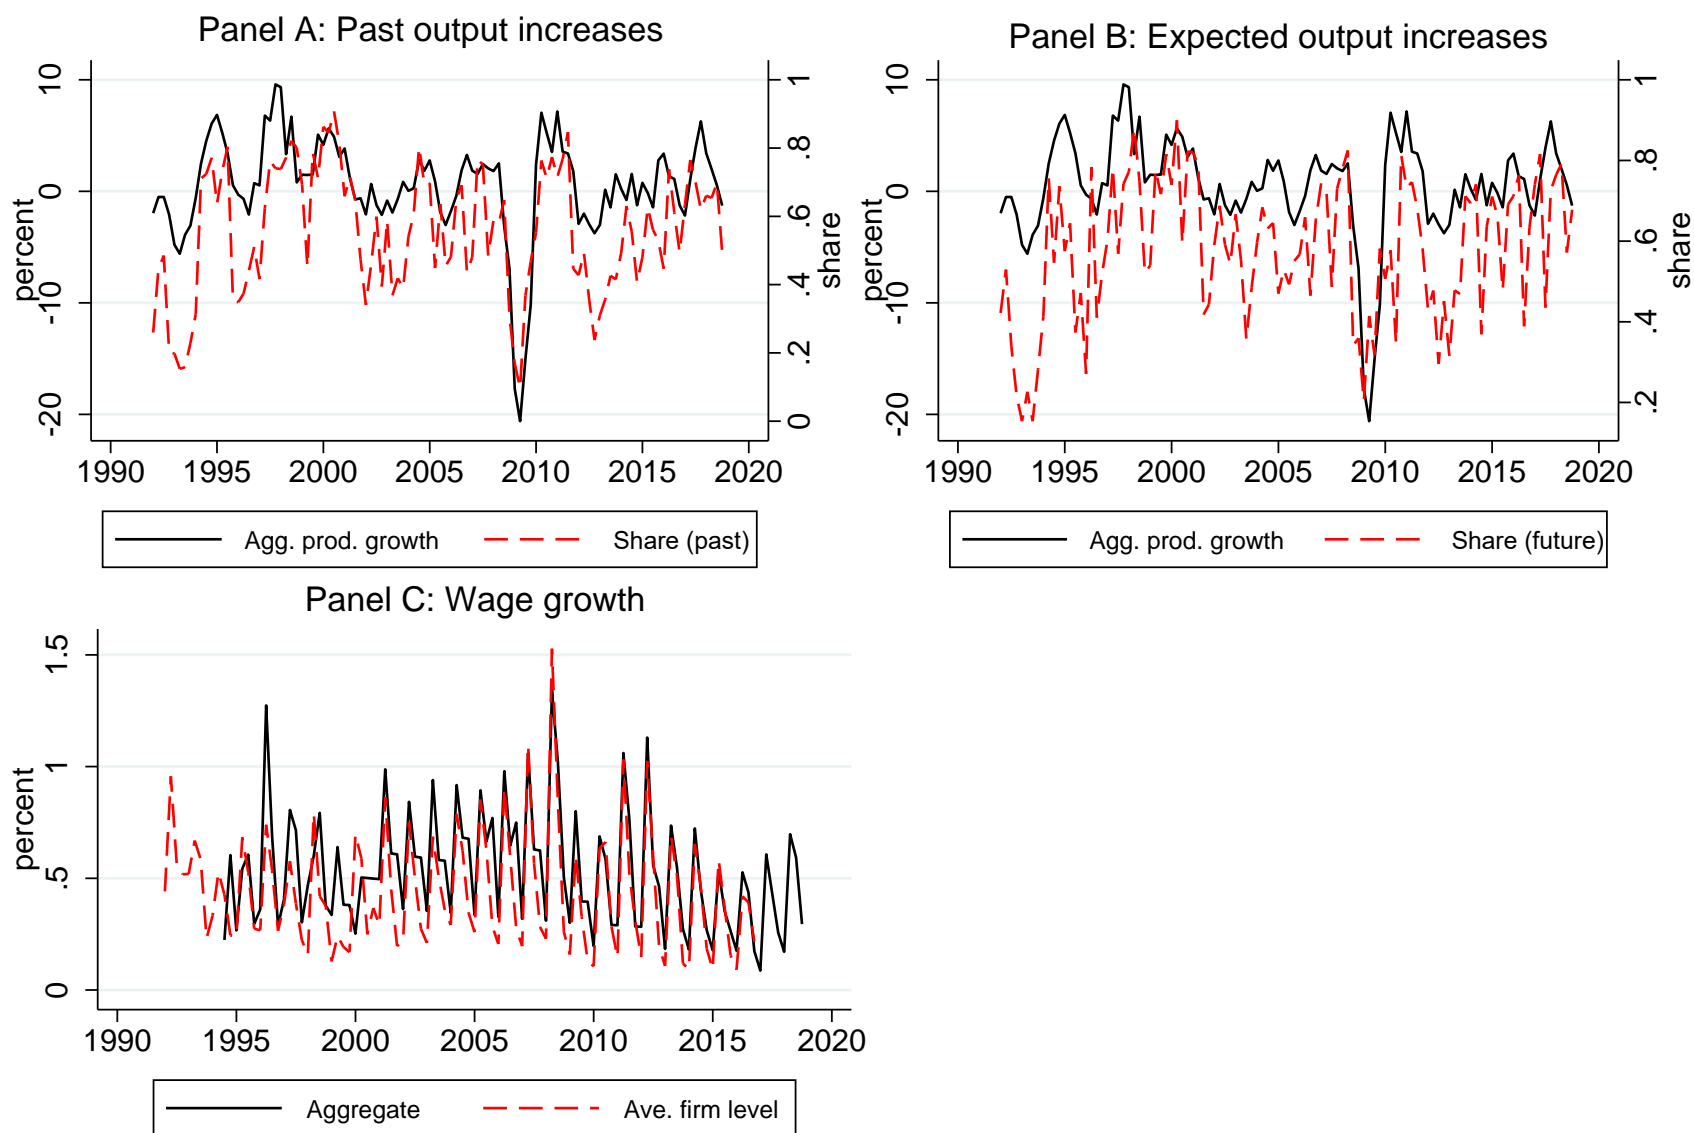
Notes: the table reports sample unweighted statistics. The number of employees is reported by firms when they answer for the first time to the survey and then this number is updated every year. Sales corresponds to total sales of a firm in a given year (for all products sold by this firm). Share of exports is calculated as the ratio between export sales (in millions euros) over total sales (for all products). Duration is calculated as the number of quarters the firm answers to the questionnaire divided by 4.

Appendix Table 2. Share of industries with industry-level inflation predicting aggregate inflation at a given horizon.

HORIZON	PRICES	PRODUCTION
H=1	0.27	0.13
H=2	0.22	0.09
H=3	0.15	0.06
H=4	0.13	0.04
H=5	0.10	0.03
H=6	0.09	0.03
H=7	0.09	0.02
H=8	0.09	0.02

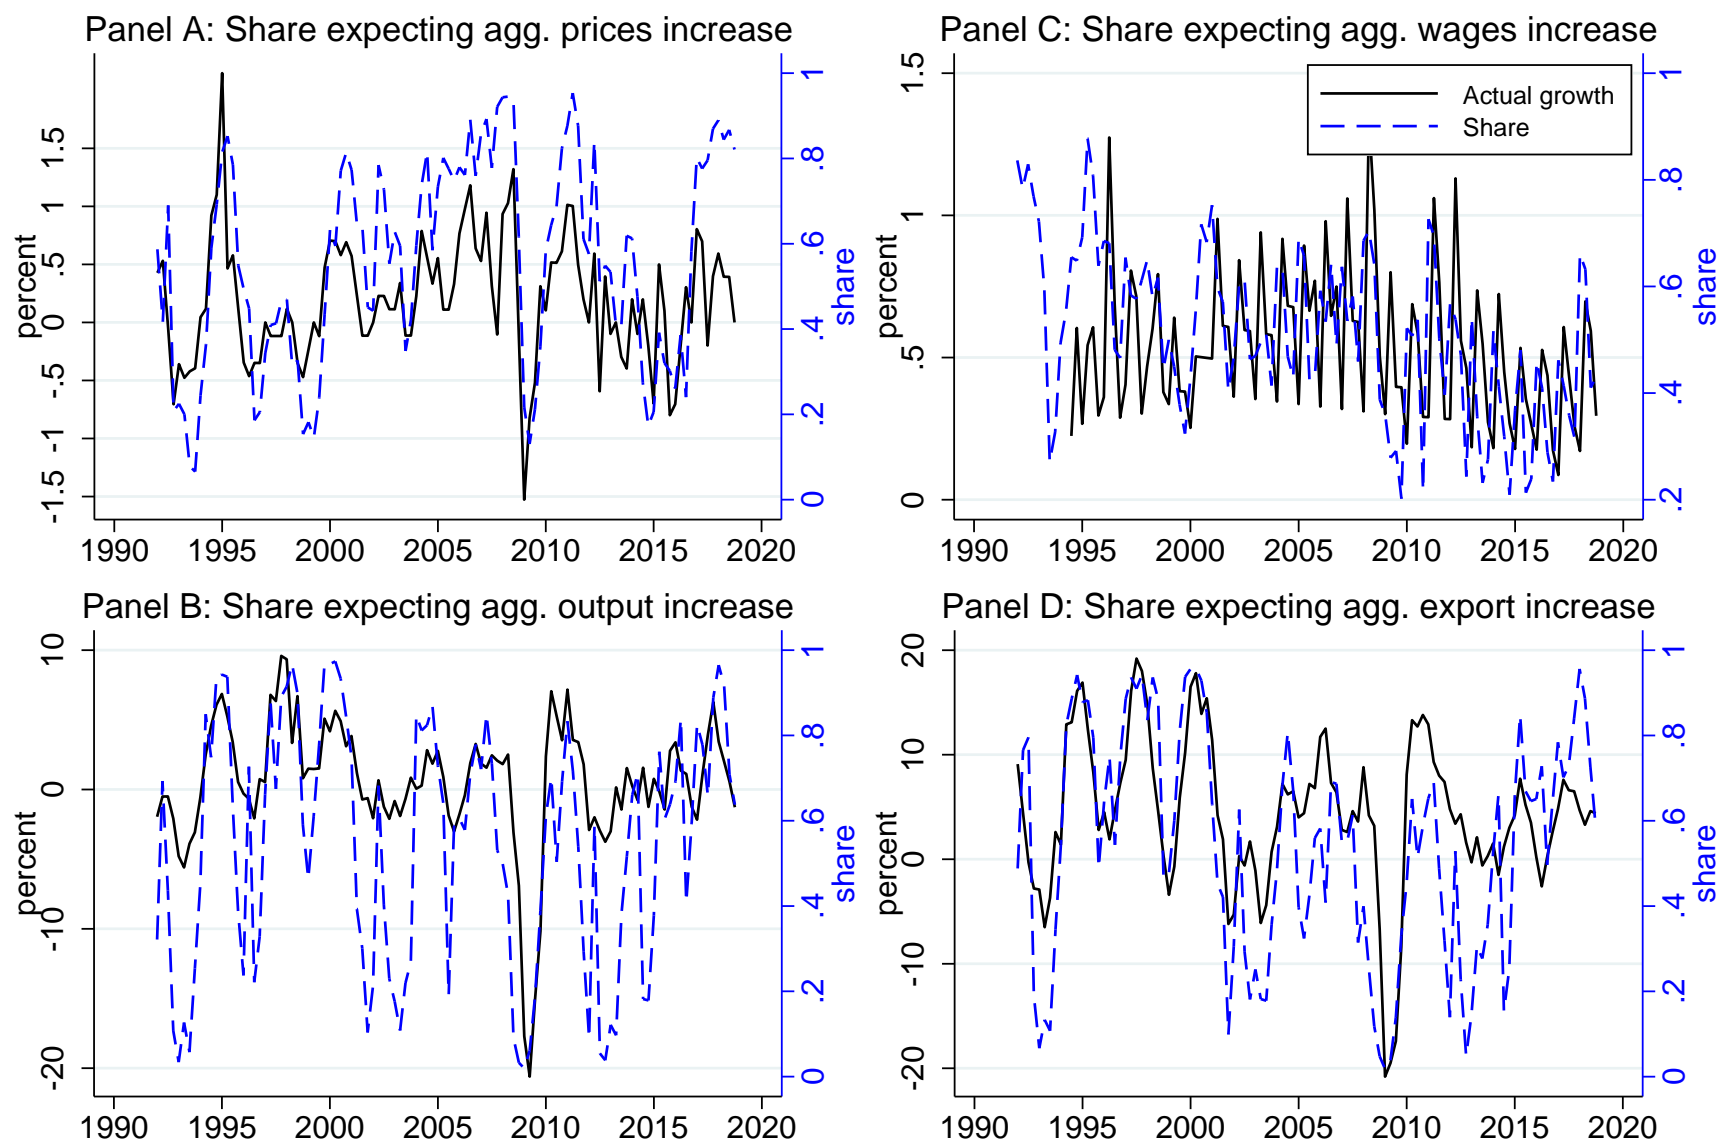
Notes: The table reports the fraction of industries for which one can reject the null (at the 95% level) that industry-level variation in either prices (column 1) or production 2 (column 2) has predictive content for subsequent aggregate price changes (column 1) or subsequent aggregate production changes (column 2) after conditioning on either current and lagged aggregate prices or production. H indicates the horizon over which we test for predictive power of contemporaneous industry level variation. See section 4 for details.

Appendix Figure 1. Firms' Expected Production and Wage Changes vs. Aggregate Production and Wage Changes.




Notes: Panel A plots the fraction of firms in the survey who report that they increased their production over the previous three months as well as a seasonally adjusted measure of industrial production in France. Panel B plots the fraction of firms in the survey who expect to increase their production over the following three months. Panel C plots the average growth in wages reported by firms in the survey versus a measure of wage growth in industry in France.

Appendix Figure 2. Firms' Aggregate Price and Output Expectations vs Aggregate Price and Output Changes.




Notes: Each panel compares the share of firms in the survey reporting increases for specific variables (Prices in panel A, Wages in Panel B, Production in Panel C, and Exports in Panel D) to growth rates of corresponding measures for total industry in France.

Appendix Figure 3. Response of Firms' Expected Output Changes to Industry vs. Aggregate Inflation and Output.




Notes: Panel A plots the cumulative response of firms' expectations of their own output changes to changes in aggregate inflation (blue line) and industry inflation (red lines); specification (1') with output expectations as the regressand. Panel B plots the corresponding impulse responses for firms' expectations their own output changes to changes in aggregate output (blue line) and industry output (red lines); specification (2') with output expectations as the regressand. The horizontal axis shows the impulse-response horizon measured in quarters.

Appendix Figure 4. Response of Firms' Actual Output Changes to Industry vs. Aggregate Inflation and Output.



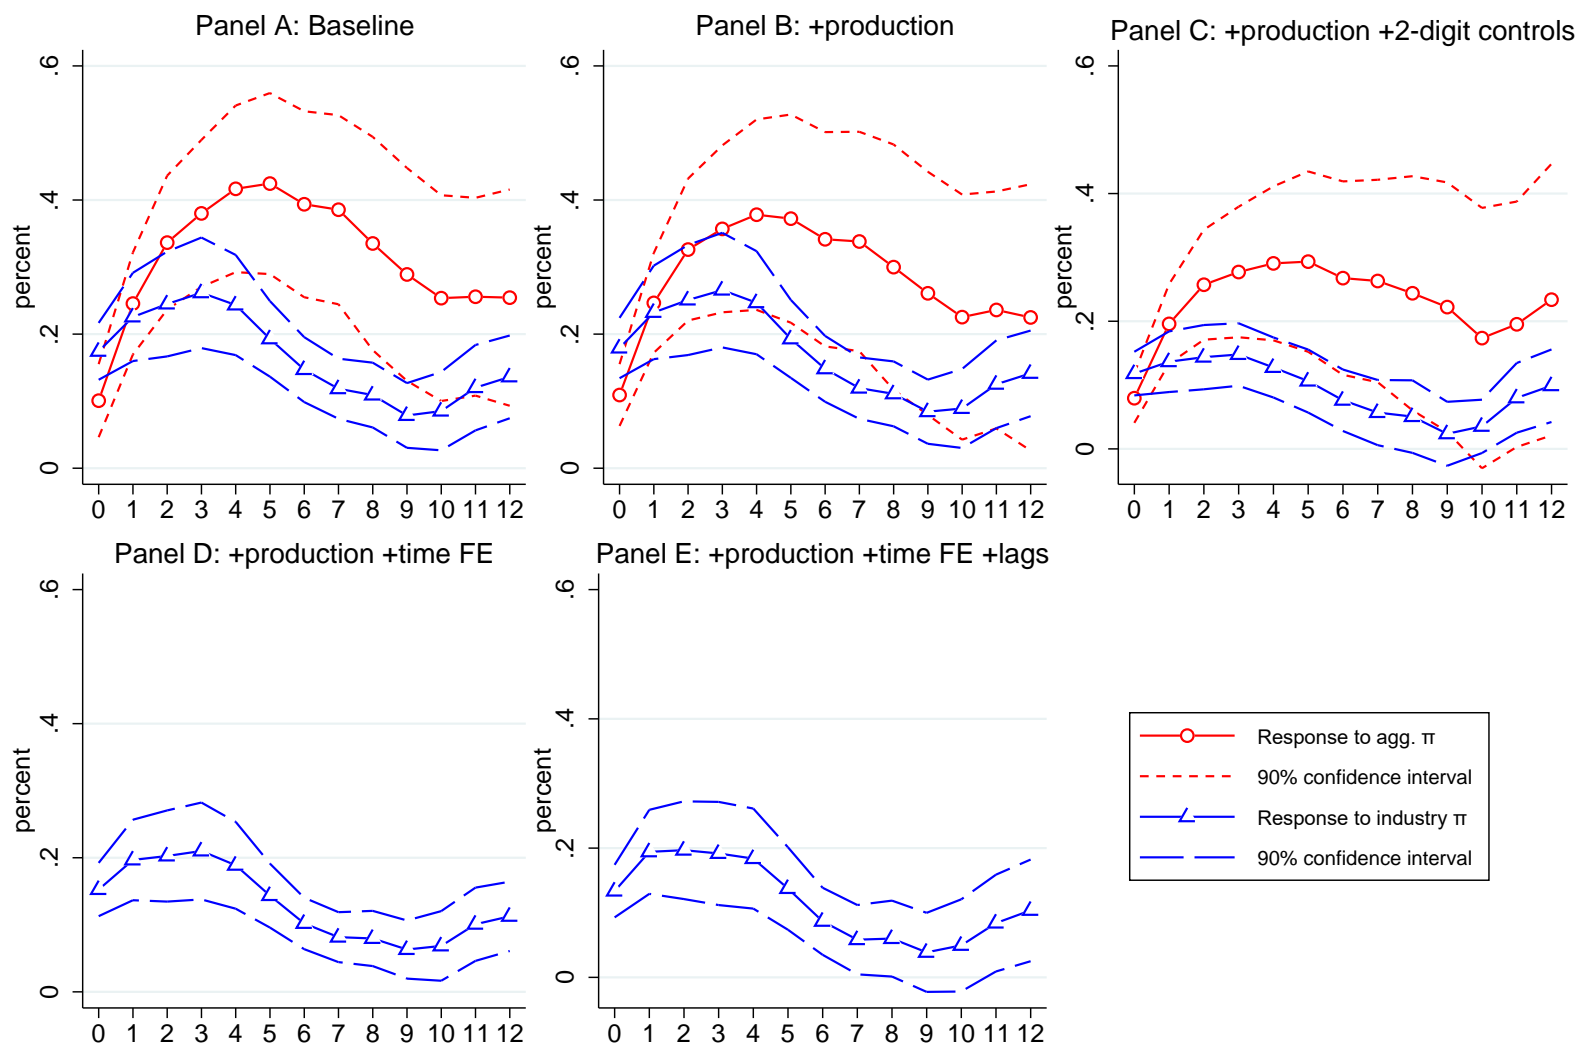
Notes: Panel A plots the cumulative response of firms' expectations of firm's own output to changes in aggregate inflation vs. changes in industry inflation; specification (3) with output expectations as the regressand. Panel B plots the equivalent responses to changes in aggregate vs. industry output; specification (4) with output expectations as the regressand. 95% confidence intervals are indicated by dotted lines in each panel. The horizontal axis shows the impulse-response horizon measured in quarters.

Appendix Figure 5. Response of Firms' Aggregate Output Expectations to Industry vs Aggregate Inflation and Output.



Notes: Panel A plots the cumulative response of firms' expectations of aggregate output to changes in aggregate inflation vs. changes in industry inflation; specification (5) with aggregate output expectations as the regressand. Panel B plots the equivalent responses to changes in aggregate vs. industry output; specification (6) with aggregate output expectations as the regressand. 95% confidence intervals are indicated by dotted lines in each panel. The horizontal axis shows the impulse-response horizon measured in quarters.

Appendix Figure 6. Robustness for Response of Firms' Expected Price Changes to Industry vs. Aggregate Inflation and Output.



Notes: This figure demonstrates robustness of responses in Figure 3 to variation in the set of controls. Panel A is the baseline specification (specification 5), panel B includes four-digit industry and aggregate production growth rate as an additional control, panel C adds two-digit industry inflation and output growth rate, panel D includes four-digit industry production and time fixed effects as controls, while panel E includes time fixed effects as well as current values and lags of four-digit industry production growth rate and inflation. The horizontal axis shows the impulse-response horizon measured in quarters.

Appendix B

Translation of the Questionnaire sent to firms *Enquête Trimestrielle de Conjoncture dans l'Industrie*

I. YOUR FORECASTS ON THE FRENCH INDUSTRY AS A WHOLE

This is your opinion on the entire French industry. Please circle the arrow corresponding to your answer

PROBABLE EVOLUTION IN THE NEXT 3 MONTHS:

- | | | | |
|----------------------------------------------------------|-----------------------------------------------|--------------------------|--------------------------|
| 1. The volume of industrial production | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| 2. The volume of exports of products manufactured abroad | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| 3. General level of prices of industrial products | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| 4. Hourly wages | significant increase // low rise // stability | | |

II. QUESTIONS ABOUT THE PRODUCTS OF YOUR COMPANY (if necessary, update the list of pre-printed products, please)

1. PRODUCT DESIGNATION (several products here, in different columns, separate answers for each product)

Please tick the appropriate box or circle the arrow corresponding to your reply.
All the questions asked below concern your production units located in France:

Approximate amount of total sales in France and abroad in 2013 (excluding taxes)
..... thousands of euros

Approximate amount of sales abroad in 2013 thousands of euros

- | | | | |
|----------------------------------------------------|--------------------------|--------------------------|--------------------------|
| a. Evolution during the last 3 months | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| b. Probable evolution over the next 3 months | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

2. GLOBAL ORDERS (OR DEMAND) (from all sources)

- | | | | |
|----------------------------------------------------|--------------------------|--------------------------|--------------------------|
| a. Evolution during the last 3 months | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| b. Probable evolution over the next 3 months | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

c. On the basis of the recorded orders still to be delivered and the current rate of manufacturing, for how many weeks do you think your business is guaranteed?
about weeks

d. Do you consider that given the season, your order book (or your addressed demand) is currently higher than usual normal lower than usual

3. FOREIGN ORDER (S)

- | | | | |
|---------------------------------------------|--------------------------|--------------------------|--------------------------|
| a. Evolution during the last 3 months | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
|---------------------------------------------|--------------------------|--------------------------|--------------------------|

b. Probable evolution over the next 3 months



c. Do you consider that, given the season, your foreign order book (or addressed foreign demand) is currently higher than average / normal / lower than average

4. YOUR COMPETITIVE POSITION

Evolution in the last 3 months:

a. On the national market



b. In foreign markets within the European Union



c. Outside the European Union



5. DELIVERY TIMES

Evolution during the last 3 months



6 YOUR FINAL PRODUCT STOCKS (products ready for sales

If the nature of your production is that you are still working without stock of manufactured products, check the box opposite

a. Evolution during the last 3 months



b. Do you consider that, given the season, your current level of inventories of manufactured products is higher than average / average / lower than average

7. YOUR SELLING PRICES

Evolution of your sales prices (excluding taxes) during the last 3 months



Please also indicate their approximate variation over the last 3 months in % ...

Evolution of your export sales prices expressed in euros during the last 3 months ...

Please also indicate their approximate variation in %



Probable evolution of your sales prices (excluding taxes) over the next 3 months



Please also indicate their approximate variation in %

III. NATURE AND IMPORTANCE OF THE COMPANY

The data below relates to your production in France:

1. Order of magnitude of your turnover (excluding taxes) in 2013 thousands of euros

2. Number of employees employed by the company as of December 31, 2013..... employees

3. Approximate amount of your foreign sales in 2013.... .. thousands of euros

IV. SOME INDICATIONS ON THE CURRENT SITUATION IN YOUR COMPANY

All the questions asked below concern your production units located in France:

1. Factors currently limiting your production (Place a cross in the appropriate box)

Are you currently prevented from developing your production as you would like because:

- insufficient demand?
- the inadequacy of your equipment or material?
- the inadequacy of a staff that you have difficulties to increase?
- financial constraints?
- supply difficulties?
- other factors (specify)?
- not applicable (you are currently able to develop your production as you wish)

2. Bottlenecks and use of production capacities

If you receive more orders, could you produce more, with your current means? YES - NO

If YES what could be the increase of your production with the existing capital and without hiring additional staff? about %

Could you produce more by hiring additional staff? YES ? NO ?

Your company currently operates at% of its overall capacity.

This is the ratio (in%) of your current production to the maximum production you could get by hiring possibly additional staff.

3. Based on your current order backlog and likely future orders over the next few years/months, do you consider that your current production capacity:
Is more than enough? is sufficient ? it's not enough ?

4. Are you currently experiencing cash flow difficulties? YES ? NO

V. QUESTIONS CONCERNING THE LABOR

Please tick the appropriate box or circle the arrow corresponding to your answer.




1. Are you currently experiencing recruitment difficulties? YES ? NO?

If YES, for which types of personnel? laborers and specialized workers? skilled workers and foremen? technicians or executives?

2. Total number of employees and weekly hours of work




Number of employees

a. Evolution during the last 3 months   

b. Probable evolution over the next 3 months   

Hours of work

a. Evolution during the last 3 months   

b. Probable evolution over the next 3 months   

3. Rate of pay (put 0 if they have not changed)

On average, how much did hourly wages vary in your business during the fourth quarter of 2014?%

Appendix C:

Forecast Errors for Firms' Own Price Changes

In this appendix, we characterize the properties of forecast errors for firms' own price changes.

Suppose firms have optimal price equal to marginal cost: $p_t^\# = mc_t$. Marginal cost follows an AR(1): $mc_t = \rho mc_{t-1} + \varepsilon_t$. Firms don't observe marginal cost but receive signal each period: $s_t = mc_t + v_t$

Beliefs about marginal cost follow: $E_t mc_t = G s_t + (1 - G) E_{t-1} mc_t = G s_t + \frac{(1-G)}{\rho} E_{t-1} mc_{t-1}$

Firms set prices for that period after receiving signal so $p_t = E_t mc_t$ and their expected price for the next period is $E_t p_{t+1} = E_t mc_{t+1} = \rho E_t mc_t = \rho p_t$. To know prices and expected prices, we just need to track evolution of beliefs about marginal costs.

These follow:

$$E_t mc_t = G s_t + \frac{(1-G)}{\rho} E_{t-1} mc_{t-1} = G mc_t + (1-G) \rho E_{t-1} mc_{t-1} + G v_t$$

Dynamics of perceived marginal costs after a shock to actual marginal costs are:

$$\begin{aligned} \frac{dE_t mc_t}{d\varepsilon_t} &= G \\ \frac{dE_{t+1} mc_{t+1}}{d\varepsilon_t} &= G\rho + (1-G)\rho G \\ \frac{dE_{t+2} mc_{t+2}}{d\varepsilon_t} &= G\rho^2 + (1-G)\rho[G\rho + (1-G)\rho G] = G\rho^2 + (1-G)\rho^2 G + (1-G)^2 \rho^2 G \end{aligned}$$

...

$$\frac{dE_{t+h} mc_{t+h}}{d\varepsilon_t} = G\rho^h (1 + (1-G) + \dots + (1-G)^h) = G\rho^h \left[\frac{1 - (1-G)^{h+1}}{G} \right] = (1 - (1-G)^{h+1}) \rho^h$$

And

$$\frac{dE_{t+h} mc_{t+h+1}}{d\varepsilon_t} = (1 - (1-G)^{h+1}) \rho^{h+1}$$

Since $p_t = E_t mc_t$ it follows that the impulse response of prices is given by

$$\frac{dp_{t+h}}{d\varepsilon_t} = \frac{dE_{t+h} mc_{t+h}}{d\varepsilon_t} = (1 - (1-G)^{h+1}) \rho^h$$

And the impulse response of expected prices is given by

$$\frac{dE_{t+h} p_{t+h+1}}{d\varepsilon_t} = \frac{dE_{t+h} mc_{t+h+1}}{d\varepsilon_t} = (1 - (1-G)^{h+1}) \rho^{h+1}$$

So the impulse response of forecast errors $FE_t \equiv p_t - E_{t-1} p_t$ follows


$$\begin{aligned} \frac{dFE_{t+h}}{d\varepsilon_t} &= \frac{dp_{t+h}}{d\varepsilon_t} - \frac{dE_{t-1} p_{t+h}}{d\varepsilon_t} = (1 - (1-G)^{h+1}) \rho^h - (1 - (1-G)^h) \rho^h = \rho^h [-(1-G)^{h+1} + (1-G)^h] \\ &= G(1-G)^h \rho^h = G[(1-G)\rho]^h \end{aligned}$$

This is the same result as in Coibion and Gorodnichenko (2012): forecast errors converge due to learning ($1-G$) and transitory nature of shock (ρ). In CG (2012), we then normalize by response of fundamental (here marginal cost), which captures the effect of ρ and leaves only $1-G$.

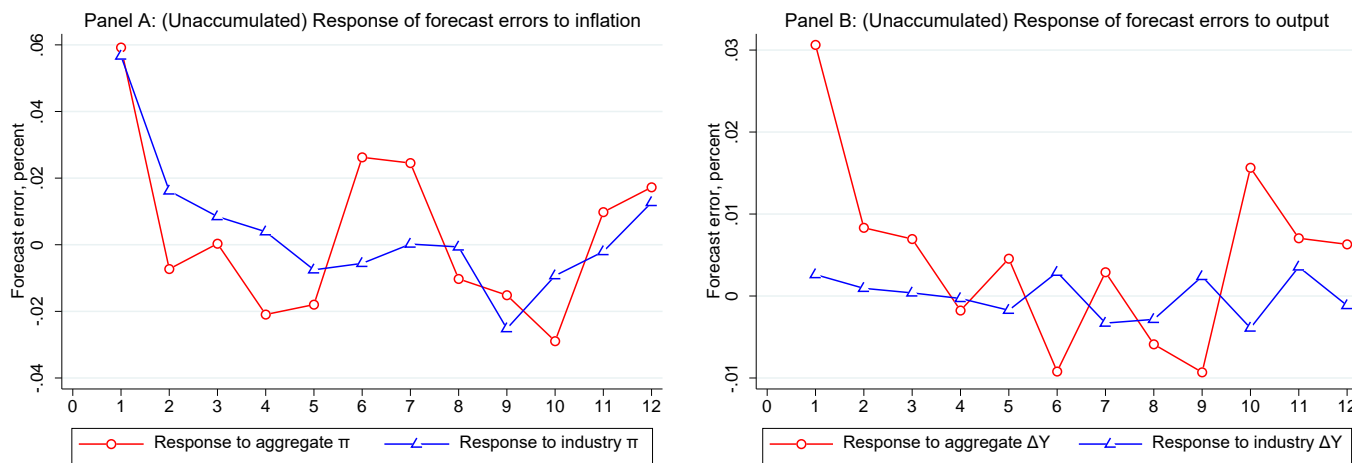
Here, we do not observe fundamental marginal cost. If we normalize forecast error response by response of actual prices:

$$\frac{\frac{dFE_{t+h}}{d\varepsilon_t}}{\left[\frac{dp_{t+h}}{d\varepsilon_t}\right]} = \frac{G[(1-G)\rho]^h}{(1-(1-G)^{h+1})\rho^h} = \frac{G(1-G)^h}{1-(1-G)^{h+1}}$$

Which is highly nonlinear in Kalman gain. For example, these are normalized impulse responses for different values of the Kalman gain:



In practice, if we estimate impulse responses of forecast errors for firms' own prices in the French survey of firms, we get the following IRFs:



When we estimate Kalman gain after normalizing by the IRF of own-price responses, results are exceedingly sensitive to all empirical choices (e.g. control variables, length of IRF etc) and do not point to any clear result for associated Kalman gains other than that they are systematically below 1, consistent with the presence of imperfect information.