



No firm is an island? How industry conditions shape firms' expectations[☆]



Philippe Andrade^a, Olivier Coibion^b, Erwan Gautier^c, Yuriy Gorodnichenko^{d,*}

^a Research Department, Federal Reserve Bank of Boston - 600 Atlantic Avenue, Boston, MA 02210, United States

^b Department of Economics, University of Texas - Austin, 2225 Speedway, Austin, TX 78712, USA

^c Directorate of Statistics, Economics and International Relations - Banque de France - 31, rue Croix des Petits Champs, 75039 Paris Cedex 1, France

^d National Bureau of Economic Research, 1050 Massachusetts Ave, Cambridge, MA 02138, USA

ARTICLE INFO

Article history:

Received 13 November 2020

Revised 16 May 2021

Accepted 17 May 2021

Available online 28 May 2021

JEL:

E2

E3

E4

Keywords:

Expectations

Rational inattention

Inflation

ABSTRACT

Using a survey of French manufacturing firms, we study how firms' expectations and actions are affected by both aggregate and industry-specific conditions. In response to industry-level shocks that have no aggregate effects, firms' aggregate expectations respond persistently. This is consistent with "island" models in which firms use the local prices they observe to make inferences about broader aggregate conditions. These patterns are related to observable characteristics of firms and the industries in which they reside. Finally, we extend the analysis to firms' expectations over their own future price changes and document how these respond to both industry and aggregate variation.

© 2021 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

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),¹ we clearly

[☆] We are grateful to Mark Bils, an anonymous referee, and our discussants Isabelle Salle and Mirko Wiederholt for very helpful comments as well as George-Marios Angeletos, James Bullard, Gianluca Violante, and conference participants at the Banque de France conference on Heterogeneity in Macroeconomics for suggestions and at the Carnegie-Rochester-NYU 2021 conference. We also thank seminar participants at Berkeley and the Federal Reserve Bank of Cleveland. We thank Sylvie Tarrieu for excellent research assistance. We are also grateful to INSEE for providing the access to the micro data and to the CASD (Centre d'Accès Sécurisé Distant) for the distant data access. 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 in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Boston, the Federal Reserve System, the Banque de France or the Eurosystem.

* Corresponding author.

¹ Bernanke (2007) observed, "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...?". The full speech is available at <https://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm>

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. This paper provides new evidence documenting how conditions in a firm's industry play an important role in shaping their view of broader macroeconomic conditions.

Specifically, the paper shows that — consistent with “island” models, in which firms form beliefs about the macroeconomy 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 signals to form beliefs about broader economic conditions. To the best of our knowledge, this is the first firm-level evidence confirming a central prediction of this class of models pioneered by Lucas (1972). Firms' reliance on industry conditions, which are notoriously volatile, to help form aggregate beliefs provides one potential explanation behind the high levels of disagreement about aggregate conditions observed among firms.

We establish these facts using a little-used survey of French manufacturers. This survey has both a large cross-section of firms (~3000) 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.

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 but also aggregate production growth and industry production growth. This provides a tractable approach to address a number of related questions.

We first 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 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 La'O, 2013, Nimark, 2014, Afrouzi, 2016, Gaballo, 2018). 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, this paper is 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 (2015b) emphasize the role played by gasoline prices in shaping households' inflation expectations. Cavallo, Cruces and Perez-Truglia (2017) and D'Acunto et al. (2021) 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 has been missing.

We also consider the response of firms' expectations about their own prices and production to both industry and aggregate variation, expectations which are available in *quantitative* form in the survey. We find that firms' expectations about their future price changes respond gradually to each type of variation: the forecast errors of firms are positively serially correlated after changes in both industry and aggregate inflation. This finding therefore extends the evidence for inattention documented in Coibion and Gorodnichenko (2012) to industry conditions and is broadly supportive of imperfect information

models. Furthermore, the response of firms' expectations is broadly consistent with their subsequent decisions. For example, firms expect to raise prices quickly after increases in industry inflation and they do indeed raise their prices quickly, more rapidly than they do after changes in aggregate inflation. This is consistent with prior work that prices tend to respond much more rapidly to industry-level shocks than aggregate shocks (e.g., Boivin, Giannoni and Mihov, 2009, Mackowiak, Moench and Wiederholt, 2009). One of our contributions is therefore to show that this extends to firms' expectations as well.

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, 2015a), or the characteristics of forecast revisions (e.g., Andrade and Le Bihan, 2013). This paper differs from this earlier line of work by studying in particular how firms' aggregate expectations respond to industry conditions.

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 are 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 4000 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 sample period, with on average 2500 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 4000 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 \times firm \times product) and approximately 270,000 firm-level observations (time \times 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 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%. Fig. 1 (Panel B) plots the average price change reported

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

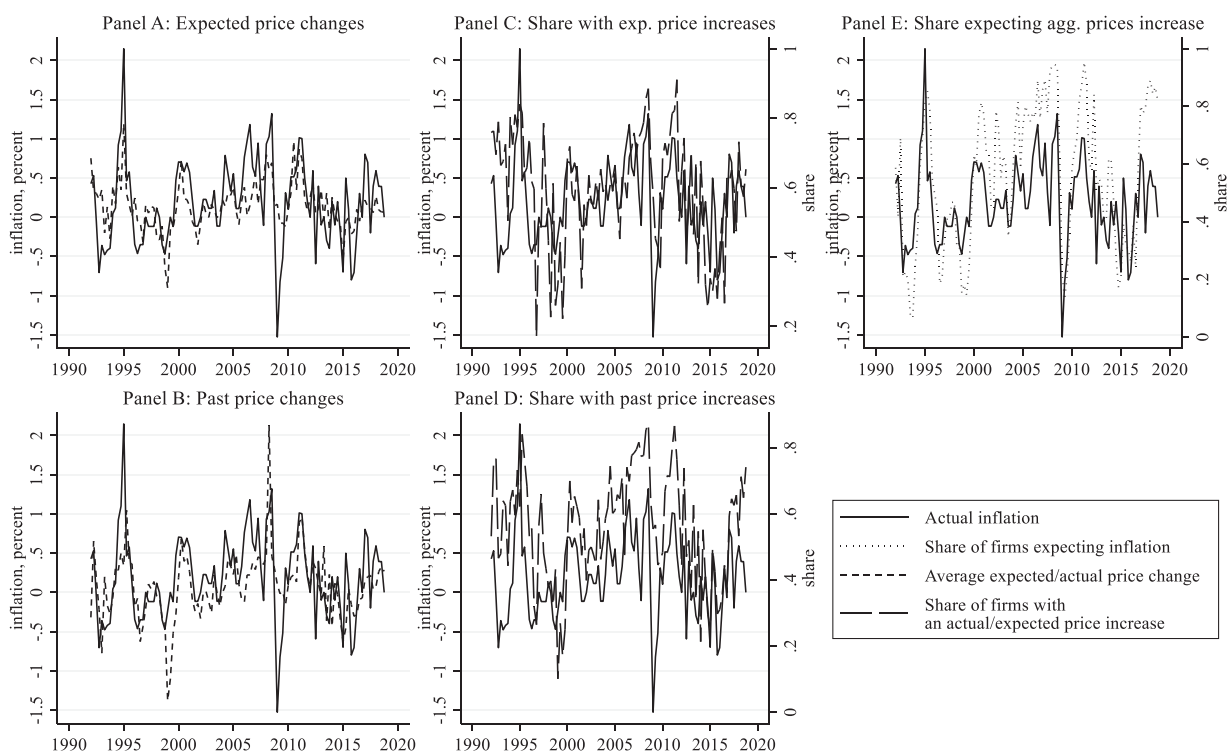
³ The French survey is part of the harmonized European Commission framework of business surveys since 2004. A majority of questions asked are common over different surveys across EU countries.

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.

Table 1
Descriptive statistics.

	Firm specific outcomes		Aggregate Expectations (3)
	Expectations (1)	Past (2)	
<i>Panel A. Prices</i>			
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
<i>Panel B. Production</i>			
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.

**Fig. 1.** Expected and actual price changes by firms, expected aggregate inflation, actual 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 presents the PPI inflation and the average price changes over the last three months reported by firms. Panel C compares actual PPI inflation to the average share of firms reporting that they expect to raise prices in next three months. Panel D reports the fraction of firms that claim to have increased prices over the preceding three months in the survey. Panel E plots actual PPI inflation and 3-month ahead expected aggregate inflation (the share of respondents expecting prices at the aggregate level to increase).




Fig. 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. Vertical bars correspond to the 95%-confidence interval.

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 series 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 Fig. 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. Fig. 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, these shares are computed 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. Fig. 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 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) closely lining up with ex-post price changes.

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.

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 Fig. 1), 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 Fig. 1). We find similar results for the share of firms expecting an increase in aggregate inflation (Panel E in Fig. 1), output, wages and export growth (Appendix Fig. 2). This close alignment between official statistics and the survey results again confirms that expectations data (even if they are qualitative) from this survey are informative.

3. Firms' aggregate beliefs after industry shocks

The unique characteristics of this survey data, namely a long, representative panel of French firms that includes measures of expectations and firm characteristics, make it an ideal setting to study 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 new 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), posit that *rational* firms face information constraints: firms are 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-specific factors, leading to a muted reaction of their own prices.⁶ Are 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?

Since firms in the survey are asked about what they expect to happen to aggregate prices and production along with those for their own firms, we can 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 regress ex-post changes in the aggregate expectations of firms on innovations to both industry and aggregate inflation and output. Specifically, we regress for each time horizon h :

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} \pi_{t+k+1}^{agg}\} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \delta_h \Delta y_t^{agg} + \theta_h \Delta y_t^j + \mu \mathbf{X}_{t-1} + \varepsilon_{t,h,i}, \quad (1)$$

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. In our baseline regression, the industry 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).⁷ Thus, the left-hand side of specification (1) characterizes the dynamics of firms' 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 of $\beta, \gamma, \delta, \theta$ do not have a direct interpretation.

The right-hand side of specification (1) includes measures of both aggregate and industry-specific inflation (π_t^{agg} and π_t^j) respectively. This allows us to distinguish between variations in prices that are common across industries versus those that

⁶ While Lucas (1972) and similar models treat firms as isolated islands, one obtains similar results if firms operate in clusters/industries and hence confound aggregate and industry-level shocks. In other words, it is not materially important for our tests whether firms are islands or archipelagos: the testable implication of this theory is the same, i.e., "local" shocks influence aggregate expectations.

⁷ All product-level price series are available from INSEE website (<https://insee.fr/en/statistiques/series/108665892>).




Fig. 3. Response of firms' aggregate inflation expectations to industry vs. aggregate inflation shocks.

Notes: the figure plots the cumulative response of firms' expectations of aggregate inflation to innovations in aggregate inflation vs. changes in industry inflation; specification (1). 90% confidence intervals (CI) are indicated by dashed lines in each panel. The horizontal axis shows the impulse-response horizon measured in quarters.

are specific to industries. Thus, γ_h measures the cumulative h -period effect of industry-level inflation changes on a firm's aggregate inflation expectations. We also similarly control for aggregate output growth and industry-specific growth (Δy_t^{agg} and Δy_t^i) respectively, which allows us to distinguish between industry and aggregate variation using real quantities as well as prices. Finally, additional controls (\mathbf{X}_{t-1}) are included. In our baseline specification, these controls include two lags of each of the independent variables. By including lags of these, we can interpret coefficients like $\{\beta_h\}$ and $\{\gamma_h\}$ as identifying the impulse response to *innovations* in aggregate inflation and industry-level inflation respectively.

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 aggregate expectations to industry-level inflation shocks that are orthogonal to contemporaneous innovations in aggregate inflation, which makes our analysis similar in spirit to the factor decomposition in [Boivin, Giannoni, and Mihov \(2009\)](#). The baseline specification also includes firm-specific fixed effects to capture the unobserved time-invariant firm characteristics that can affect average pricing behavior. Standard errors are clustered by time and firm to account for cross-sectional and time series correlation of the error term. We later consider a wide range of robustness checks to this baseline specification.

The results of these regressions are plotted in Panel A of [Fig. 3](#). This panel includes the estimated $\{\beta_h\}$, which indicate the response of firms' aggregate inflation expectations to aggregate inflation variation, and the estimated $\{\gamma_h\}$, which indicate the response of firms' aggregate inflation expectations to industry inflation variation. A shock to aggregate inflation is followed by a pattern of gradually increasing aggregate inflation expectations, as one would expect. With industry-level inflation shocks, we also find a gradually increasing response of aggregate inflation expectations.

This finding is notable for several reasons. First, it provides a new test and rejection of the full-information rational expectations hypothesis. Since specification (1) explicitly controls for aggregate inflation, innovations to industry-specific inflation are orthogonal to aggregate dynamics. This means they have no effect on aggregate inflation and therefore firms' expectations of aggregate inflation should be unaffected. We find that they are, a clear violation of the full-information rational expectations hypothesis. Second, this rejection is precisely what one would expect under the island models pioneered by [Lucas \(1972\)](#). Firms are unable to separately identify innovations that are specific to their industry from those that are

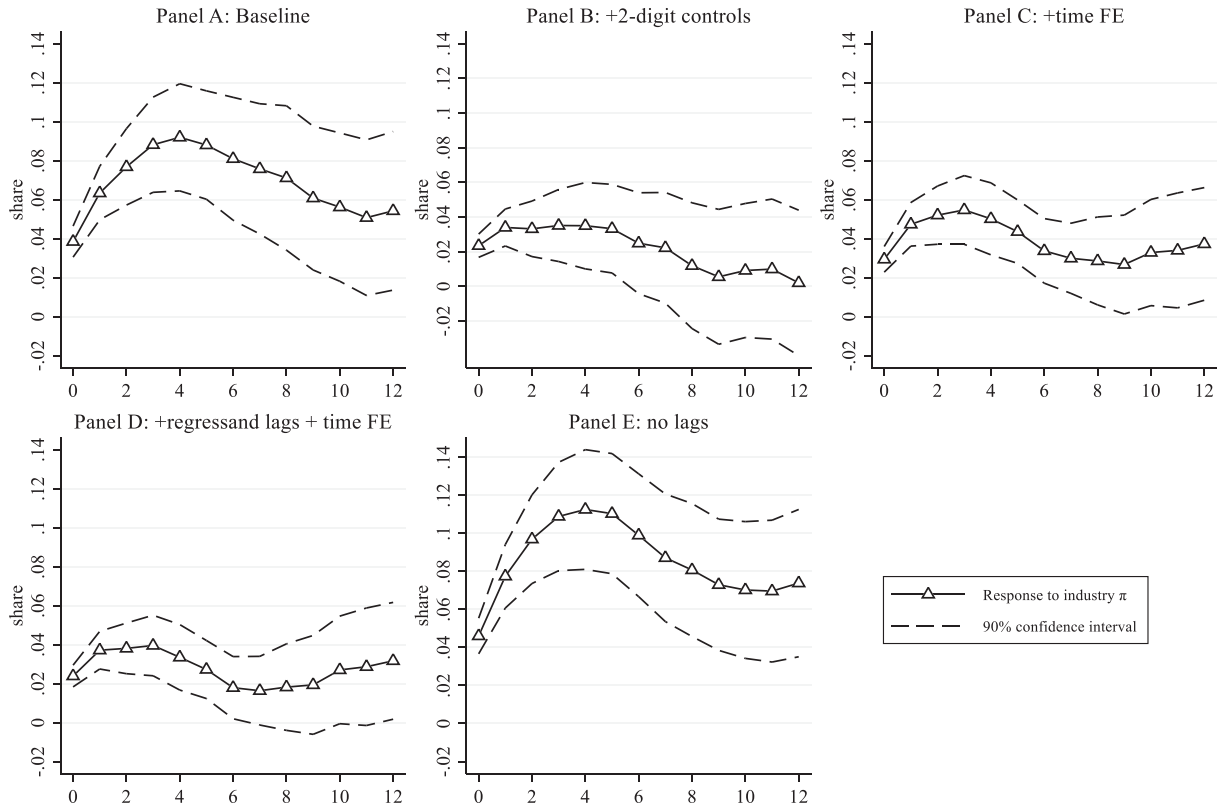


Fig. 4. Robustness: responses of aggregate inflation expectations to industry-level inflation shocks.

Notes: Each panel plots the impulse response of firms' aggregate inflation expectations to industry variation in inflation. Panel A is the baseline specification (specification 1), panel B includes two-digit industry inflation and output growth rate (current values and lags) as additional controls, panel C adds time fixed effects as controls, panel D includes lags of endogenous variable and time fixed effects, panel E includes only current values of four-digit and aggregate industry production growth rate and inflation. The horizontal axis shows the impulse-response horizon measured in quarters.

common to all industries, so they adjust their aggregate expectations even when innovations are industry-specific. To the best of our knowledge, this is the first firm-level evidence that directly validates this mechanism.

This result is quite robust and holds under a number of alternative specifications and identifications. For example, Appendix Fig. 3 plots the equivalent responses of firms' aggregate inflation expectations to aggregate vs. industry shocks but identifying the latter using output measures. We find equivalent results (i.e., estimated $\{\theta_h\}$ in specification (1)): changes in industry output that are orthogonal to aggregate output affect firms' aggregate inflation expectations. Panel C in Appendix Figs. 5 and 6 present equivalent results for firms' expectations of future aggregate output: these respond to changes in industry inflation orthogonal to aggregate inflation (Appendix Fig. 5) and to changes in industry output that are orthogonal to aggregate output changes (Appendix Fig. 6).

These results are also robust to a variety of alternative specifications with additional controls. For example, Panel B of Fig. 4 presents estimates of γ_h in Eq. (1) augmented to also include industry inflation and output measured at the 2-digit level, such that γ_h measures response to narrowly defined industry-level variation. The results are qualitatively unchanged. Panel C of Fig. 4 reports results of estimating Eq. (1) with time fixed effects (dropping all aggregate variables). We continue to find that industry innovations lead to persistent effects on firms' aggregate inflation expectations. Panel D of Fig. 4 reports results when including both time fixed effects and lags of the dependent variable, again yielding the same qualitative results. Finally, Panel E of Fig. 4 reports results when lags are not included in Eq. (1): again, the results are qualitatively unchanged.

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:

$$\sum_{k=0}^h \pi_{t+k}^{agg} = \alpha_j + \delta_{hj} \pi_t^j + \eta_{hj} \pi_t^{agg} + error, \quad (2)$$

$$\sum_{k=0}^h \Delta y_{t+k}^{agg} = \alpha_j + \delta_{hj} \Delta y_t^j + \eta_{hj} \Delta y_t^{agg} + error. \quad (3)$$

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.⁸ At short horizons, there is some evidence of predictability of subsequent aggregate inflation from industry-level variation in inflation: one 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 Fig. 5. These 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 : $\pi_t^{(s)} = b_0^{(s)} + b_1^{(s)} \pi_t^j + error$ 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 .⁹ Then we add this other industry's inflation rate $\pi_t^{s^*}$ as well as production growth $\Delta y_t^{s^*}$ to the baseline specification (1):

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} \pi_{t+k+1}^{agg}\} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \psi_h \pi_t^{s^*} + \delta_h \Delta y_t^{agg} + \theta_h \Delta y_t^j + \kappa_h \Delta y_t^{s^*} + \mu \mathbf{X}_{t-1} + \varepsilon_{t,h,i}, \quad (4)$$

where \mathbf{X} includes also lags of $\pi_t^{s^*}$ and $\Delta y_t^{s^*}$ in addition to standard controls of specification (1). Panel B of Fig. 5 plots the impulse response of firms' aggregate inflation expectations to innovations in the inflation of these other industries s^* . 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.

4. Firms' beliefs and actions in response to industry and aggregate shocks


While firms' expectations about aggregate prices in the survey are only qualitative, firms provide quantitative expectations about their expected prices over the following three months. These can also be used to study the nature of firms' expectation formation process. Specifically, one can use a similar empirical approach as before to characterize how these expectations respond to different types of fluctuations. The first step is to 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:

$$\sum_{k=0}^h E_{t+k}^{i,j} dp_{t+k+1}^{i,j} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \delta_h \Delta y_t^{agg} + \theta_h \Delta y_t^j + \mu \mathbf{X}_{t-1} + \varepsilon_{t,h,i} \quad (5)$$

where $E_{t+k}^{i,j} dp_{t+k+1}^{i,j}$ is now the quantitative expectation at time $t+k$ of firm i in industry j for their price changes over the next quarter. 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. Since these expectations of price changes are based on quantitative questions, the response provides a quantitative estimate of cumulative expected price changes. As before, we do not attempt to identify the structural sources of innovations to either aggregate or industry inflation and instead focus on reduced form innovations to both industry and aggregate fluctuations. Results are presented in Panel A of Fig. 6. In response to aggregate inflation, a relatively muted and delayed response of firms' expectations of

⁸ The results are reported in Appendix Table 2.

⁹ 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).

**Fig. 5.** Additional robustness checks.

Notes: Panel A plots the 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 (1). Panel B again plots the 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 (4). See [Section 3](#) for construction of placebo industries. The horizontal axis shows the impulse-response horizon measured in quarters.

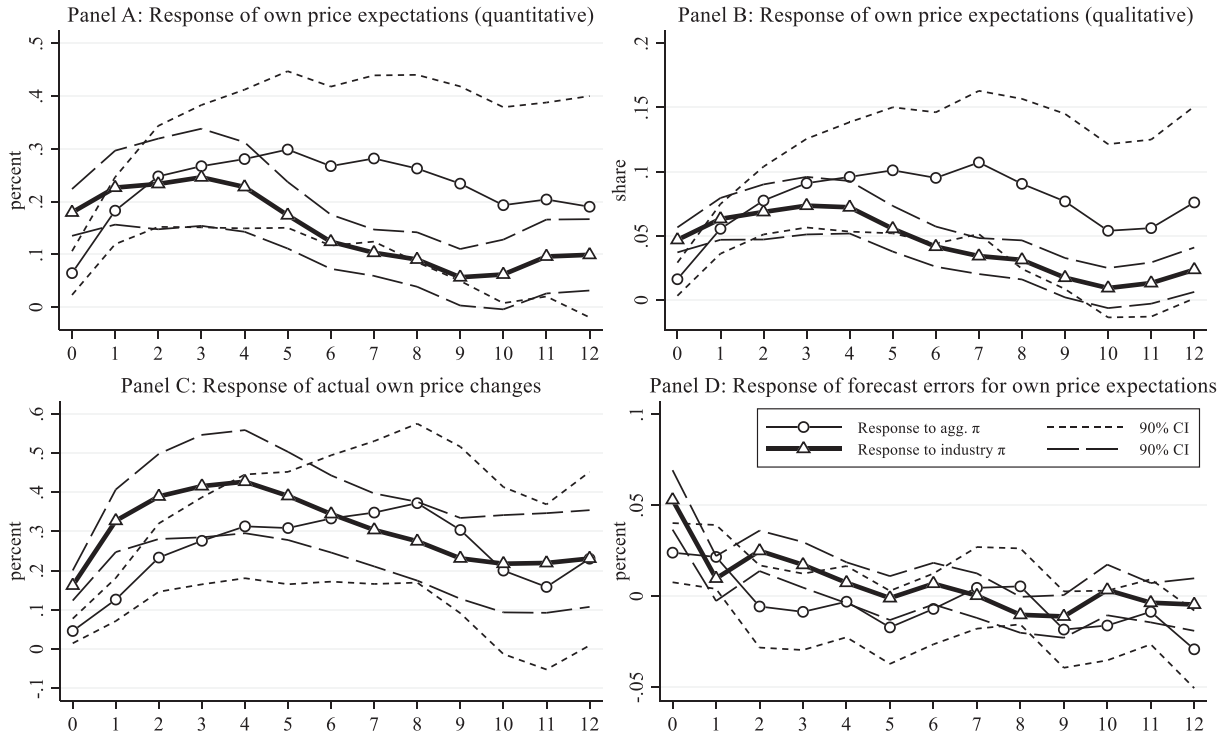


Fig. 6. Response of expectations, actual changes and forecast errors for own prices to inflation shocks.

Notes: Panel A plots the cumulative response of firms' expectations of their own price changes to changes in aggregate and industry inflation shocks; specification (5). Panel B plots the corresponding impulse responses for qualitative expectations of their own price changes specification (5'). Panel C plots the corresponding impulse responses for actual firm-level ("own") price changes to aggregate and industry-specific inflation shocks, specification (6). Panel D plots the corresponding impulse responses for forecast errors for expected changes in own prices, specification (7). The horizontal axis shows the impulse-response horizon measured in quarters.

their subsequent price changes. The pricing response is gradually increasing over the first four quarters and continues on another few quarters. Panel A 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.

Because responses about firm-level price changes are both qualitative and quantitative, one can assess the quality of qualitative responses as well as what roles are played by intensive and extensive margins. To this end, we estimate the following modification of specification (5):

$$\sum_{k=0}^h \mathbb{I}\{E_{t+k}^{i,j} dp_{t+k+1}^{i,j}\} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \delta_h \Delta y_t^{agg} + \theta_h \Delta y_t^j + \mu X_{t-1} + \varepsilon_{t,h,i} \quad (5)$$

where $\mathbb{I}\{\cdot\}$ takes values $\{-1, 0, 1\}$ for firm-level prices expected to decrease, stay the same, and increase respectively. The qualitative responses (Panel B, Fig. 6) largely follow the patterns of quantitative responses thus validating our analysis of aggregate inflation expectations in the previous section. Appendix Figs. 7 and 8 present the responses of the extensive and intensive margins of adjustment for expected own price changes. While both margins are active, the intensive appears to be somewhat more important quantitatively.

Importantly, the survey also allows us to assess the speed of firms' actual prices to the same innovations. In particular, we run the following specification

$$\sum_{k=0}^h dp_{t+k+1}^{i,j} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \delta_h \Delta y_t^{agg} + \theta_h \Delta y_t^j + \mu X_{t-1} + \varepsilon_{t,h,i} \quad (6)$$

which provide cumulative responses of firms' prices to both industry and aggregate variation in inflation. Panel C of Fig. 6 plots the resulting estimates. We again find a much faster response after industry-level shocks than aggregate shocks. This confirms a central finding of Boivin, Giannoni and Mihov (2009). Specifically, using a factor decomposition of industry-level prices, Boivin, Giannoni and Mihov (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). Our results therefore confirm this feature of the data in French data. Furthermore, we are able to show that the same pattern in firms' expectations about their price changes: these also evolve more gradually in response to aggregate fluctuations than after industry fluctuations.¹⁰

This result also provides additional evidence that firms' expectations are reflected in their decisions: when they report higher expectations for their future prices, they ultimately tend to raise their prices. This complements prior work that has found a strong unconditional correlation between firms' expectations of their future price changes and ex-post price changes (e.g., Coibion, Gorodnichenko and Kumar, 2018). In this case, however, we show a similar finding *conditional* on either aggregate fluctuations or industry-specific variation. This is only possible because of the long panel dimension available in this survey data.

The fact that both expectations of prices and actual prices respond more rapidly to industry shocks than aggregate ones could reflect different factors. For example, one possibility would be if industry shocks were less persistent than aggregate ones. Another is if firms pay less attention to aggregate shocks than industry 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 that 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.

Unlike the differential shock persistence explanation, the rational inattention explanation relies on information frictions and agents being less than fully informed about aggregate fluctuations, a feature consistent with island models discussed in Section 3. We can provide further evidence on this by comparing the ex-post response of prices to aggregate and industry variation to the ex-ante predictions of firms about those prices. This is in the same spirit as Coibion and Gorodnichenko (2012), who show the persistence of forecast errors can be used to test for the presence of information frictions.

To do so, we estimate the following specification

$$dp_{t+k+1}^{i,j} - E_{t+k}^{i,j} dp_{t+k+1}^{i,j} = \alpha_{i,h} + \beta_h \pi_t^{agg} + \gamma_h \pi_t^j + \delta_h \Delta y_t^{agg} + \theta_h \Delta y_t^j + \mu X_{t-1} + \varepsilon_{t,h,i} \quad (7)$$

and plot the estimated impulse responses in Panel D of Fig. 6. There are two key features of these responses to note. First, forecast errors are positively serially correlated after each type of innovation. This means that firms are consistently underestimating by how much they are ultimately going to raise prices after both industry and aggregate shocks. Using qualitative responses to construct forecast errors, we find similar results for firms' expectations error on their own prices but also on their output growth (see Appendix Fig. 9). Like the result in Section 3, these findings reject the null of full-information rational expectations in precisely the direction predicted by models of imperfect information. To the best of our knowledge, we are the first to do so with firms' own expectations of their future price changes and the first to show that imperfect information applies to industry-level shocks, rather than just aggregate shocks. The presence of these information frictions for both industry and aggregate conditions is consistent with the mechanism suggested by Mackowiak et al. (2009).

The second feature to note is that the forecast errors are largely overlapping and converge toward zero at broadly similar speeds. Coibion and Gorodnichenko (2012) show that, in noisy information models, the speed of convergence of forecast errors can identify the underlying degree of information rigidity after normalizing by the convergence speed of the variable being forecasted. The similar speed of convergence found here might suggest that the degree of information rigidity is therefore similar for industry and aggregate shocks, a feature which would be at odds with the rational inattention motives emphasized in Mackowiak et al. (2009). However, because the variable that firms are forecasting is their own price rather than an aggregate variable as in Coibion and Gorodnichenko (2012), one can show analytically (see Appendix C) 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 directly informative about the degree of information rigidity as is the case with forecasts of aggregate variables. Nonetheless, predictable forecast errors are still a key indicator consistent with the presence of information rigidity.¹¹

Jointly, these results support a growing body of evidence documenting pervasive information rigidities on the part of firms and households. Relative to this earlier work, our key contribution is to document that firms' aggregate expectations respond to industry fluctuations, even though the latter have no aggregate effects. This result provides direct evidence in

¹⁰ We reach the same conclusions when we study the responses of expected and actual price changes to production growth shocks at the industry and aggregate levels (Appendix Fig. 4). Appendix Figs. 5 and 6 present results on the dynamic responses of firms' expected and actual output growth (qualitative responses) to industry and aggregate inflation and production growth shocks. Again, we reach similar conclusions.

¹¹ We are very grateful to Mirko Wiederholt for making this suggestion.

Table 2
Industry heterogeneity, descriptive statistics.

	Mean (1)	Huber-robust mean (2)	25 th percentile (3)	50 th percentile (4)	75 th percentile (5)	St.Dev. (6)
Responses of expectations						
Average response of $E_t\pi^{agg}$ to $\pi^{industry}$, $\frac{1}{13} \sum_{h=0}^{12} \hat{\gamma}_{j,h}$	0.078	0.065	−0.007	0.050	0.142	0.134
Average response of $E_t\pi^{agg}$ to π^{agg} , $\frac{1}{13} \sum_{h=0}^{12} \hat{\beta}_{j,h}$	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 $dlog(p^{Comm})$	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 $E_t\pi^{agg}$ to $\pi^{industry}$) or aggregate inflation (Average response of $E_t\pi^{agg}$ 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.

favor of “island” models pioneered in Lucas (1972). The next section delves into whether the strength of this effect varies across firms depending on their observable characteristics.

5. Heterogeneity

If firms are optimally choosing either the quantity or type of information to acquire and process, these choices should likely depend on the characteristics of their industry, their product mix, etc. For example, under rational inattention models in which firms choose how much attention to allocate to industry vs. aggregate shocks (Mackowiak and Wiederholt, 2009), 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. Pasten and Schoenle (2016) and Yang (2019) similarly emphasize how the number of goods produced by a firm can affect their information choices.

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 Eq. (1) 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 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. One 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.

Fig. 7 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 to industry-specific shocks. 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 of

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.

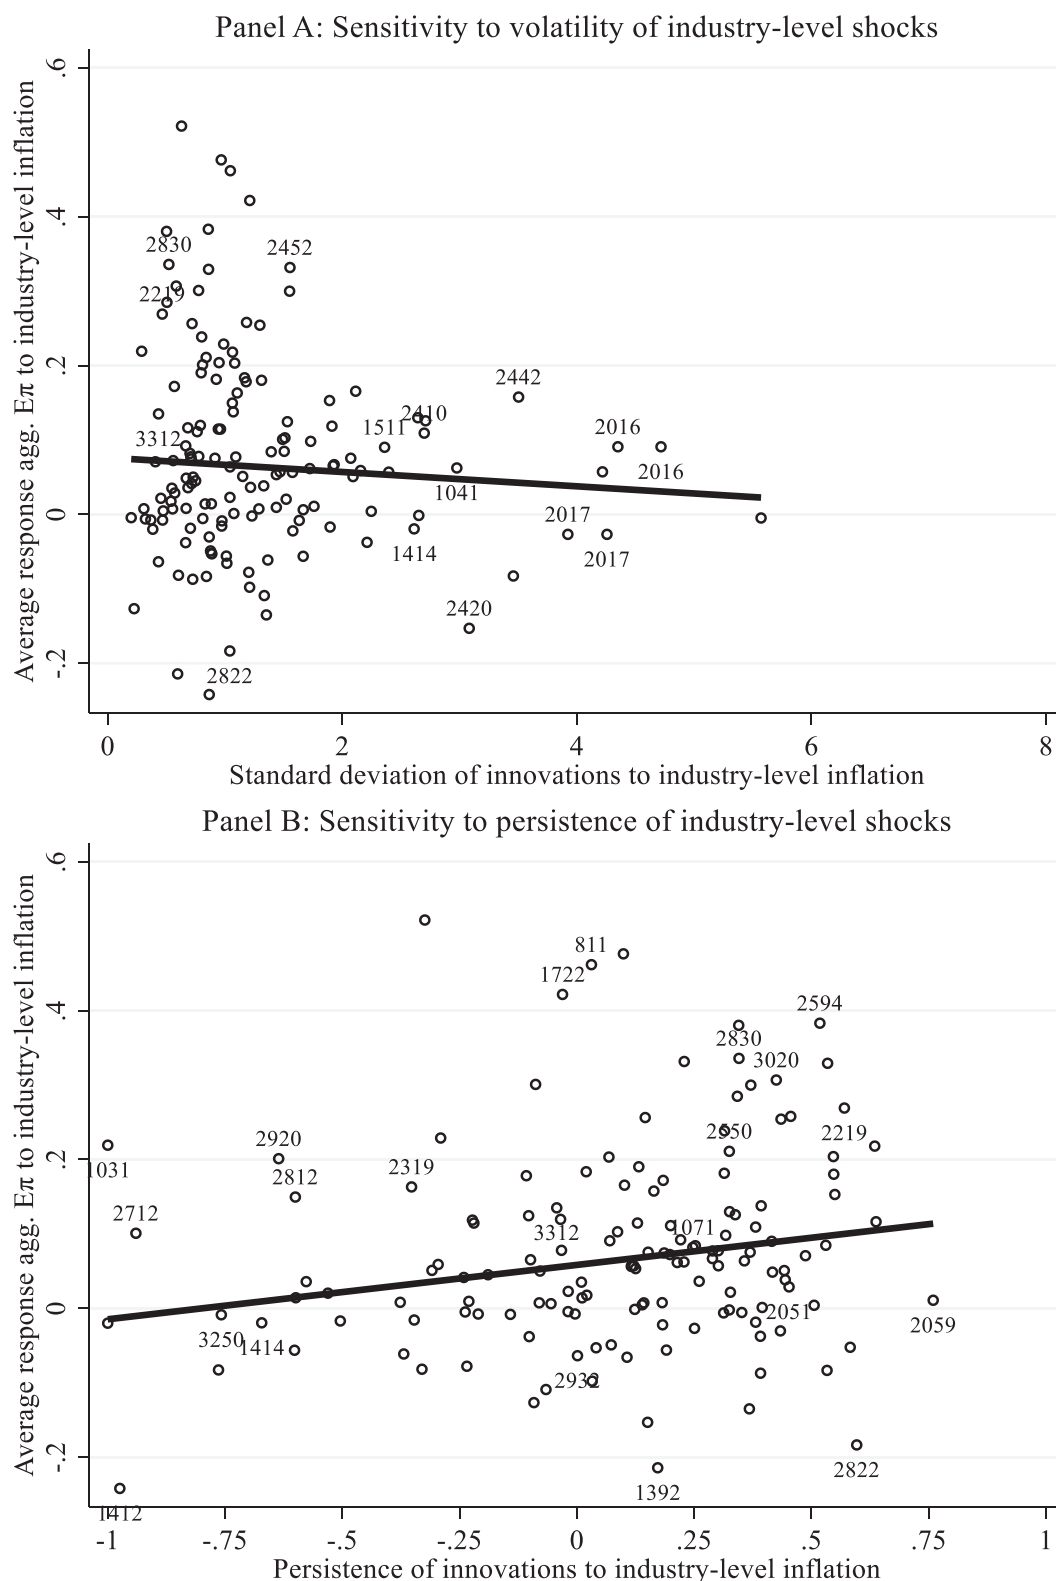


Fig. 7. Sensitivity of aggregate inflation expectations to volatility and persistence of industry-level inflation.

Notes: The top figure is a scatter plot of each industry's average impulse response of firms' expectations of aggregate price changes to changes in their industry's prices (vertical axis) versus the standard deviation of innovations to that industry's inflation rate. The bottom figure is a scatter plot of each industry's average impulse response of firms' expectations of aggregate price changes to changes in their industry's prices (vertical axis) versus the persistence of that industry's inflation rate.

Table 3

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

Regressors: Industry characteristics	Dependent variable: Ave. response of $E_t\pi^{agg}$ to $\pi^{industry}$, $\frac{1}{13}\sum_{h=0}^{12}\hat{\gamma}_{j,h}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
St.Dev. of innovations to $\pi^{industry}$	−0.012		−0.012	−0.023	−0.025	−0.026	−0.025
	(0.013)		(0.012)	(0.025)	(0.024)	(0.025)	(0.025)
Persistence of $\pi^{industry}$		0.069**	0.071**	0.066*	0.066*	0.070*	0.069*
		(0.032)	(0.032)	(0.036)	(0.035)	(0.036)	(0.037)
Ave. response of $E\pi^{agg}$ to π^{agg}				−0.010	−0.001	−0.001	−0.001
				(0.072)	(0.004)	(0.004)	(0.004)
Loading of $\pi^{industry}$ on π^{agg}				−0.004	−0.014	−0.019	−0.020
				(0.058)	(0.057)	(0.060)	(0.061)
Loading of $\pi^{industry}$ on $d\log(P^{Comm})$				0.399	0.438	0.469	0.383
				(0.507)	(0.490)	(0.517)	(0.529)
Labor cost share					0.109	0.115	0.105
					(0.113)	(0.119)	(0.122)
Material cost share					0.088*	0.089*	0.099*
					(0.048)	(0.049)	(0.051)
Export share						0.018	0.016
						(0.060)	(0.060)
Number of products						0.026	0.015
						(0.065)	(0.069)
Log(Median firm size)						−0.004	−0.006
						(0.020)	(0.020)
Capacity utilization rate							0.002
							(0.003)
Observations	126	126	126	125	126	126	126
R-squared	0.007	0.036	0.044	0.048	0.094	0.096	0.100

Notes: The table plots results from cross-sectional regressions, specification (8). 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.

motives, the former is not. Of course, there could be many other firm and industry characteristics that affect firms' incentives and capacity 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:

$$\frac{1}{13}\sum_{h=0}^{12}\hat{\gamma}_{j,h} = \phi\mathbf{Z}_j + error_j \quad (8)$$

where j indexes industries and \mathbf{Z} is a vector of industry characteristics such the volatility of industry-level inflation, labor shares, etc. For example, column 1 of Table 3 reports regressions of the response of aggregate expectations on the volatility of innovations to that industry's price level. The results confirm the scatterplot in Fig. 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. There is 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 does not increase 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 is observed: 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.¹³ 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.

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). Building on this literature, we document a new stylized fact: firms' aggregate expectations respond to industry shocks that have no aggregate effects. This implies that firms confound underlying sources of volatility: they attribute idiosyncratic industry variation they observe to aggregate forces. Our results provide the first direct micro-level empirical evidence supporting the mechanism of “island” models pioneered by 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 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 3, 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.

In addition, while the survey data we exploit cover a long period of time and a large sample of firms, we cannot assess whether the confusion between sectoral and aggregate conditions that we underline results from cognitive biases, costs and incentives to processing both types of information, or media focus that makes firms mix the two (see e.g. Chahrour et al., 2020). The specific channel could be established by designing appropriate randomized controlled experiments or lab experiments which would therefore be a useful complement to distinguish between these two potential explanations of our results. Future survey and experimental work can also help to refine our understanding of information rigidities at various levels of decision making within a firm.

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. This friction is important as it implies that, in contrast to professional forecasters and financial markets, individual firms that make pricing/investment/employment decisions pay limited attention to future macroeconomic conditions, which mitigates the efficacy of stabilization policies relying on macroeconomic expectations like forward guidance or average inflation targeting (e.g., Kiley 2016). Our results also 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.

¹³ Because the survey caps the number of products at 4, we conjecture that our data are not sufficiently detailed to find differences for multiproduct firms. Consistent with this conjecture, Bhattacharai and Schoenle (2014) show that some pricing moments behave noticeably different when the number of products is greater than 5.

Declaration of Competing Interest

None.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jmoneco.2021.05.006](https://doi.org/10.1016/j.jmoneco.2021.05.006).

References

- Afrouzi, H., 2016. Strategic Inattention, Inflation Dynamics, and the Non-Neutrality of Money. manuscript. Available at https://afrouzi.com/strategic_inattention.pdf.
- Andrade, P., Le Bihan, H., 2013. Inattentive professional forecasters. *J. Monet. Econ.* 60 (8), 967–982.
- Angeletos, G.-M., La'O, J., 2013. Sentiments. *Econometrica* 81 (2), 739–779.
- Berardi, N., Gautier, E., Le Bihan, H., 2015. More facts about prices: France before and after the great recession. *J. Money Credit Bank.* 47 (8), 1465–1502.
- Bils, M., Klenow, P.J., 2004. Some evidence on the importance of sticky prices. *J. Polit. Econ.* 112, 947–985.
- Bhattarai, S., Schoenle, R., 2014. Multi-product firms and price-setting: theory and evidence from U.S. producers prices. *J. Monet. Econ.* 66, 178–192.
- Boivin, J., Giannoni, M., Mihov, I., 2009. Sticky prices and monetary policy: evidence from disaggregated US data. *Am. Econ. Rev.* 99 (1), 350–384.
- Cavallo, A., Cruces, G., Perez-Truglia, R., 2017. Inflation expectations, learning and supermarket prices: evidence from survey experiments. *Am. Econ. J.: Macroecon.* 9 (3), 1–35.
- Chahrour, R., Nimark, K., Pitschner, S., 2020. Sectoral Media Focus and Aggregate Fluctuations, manuscript. Available at http://www.kris-nimark.net/pdf/papers/chahrour_nimark_pitschner_sectoral_media_focus.pdf.
- Christiano, L., Eichenbaum, M., Evans, C., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *J. Polit. Econ.* 113 (1), 1–45.
- Coibion, O., Gorodnichenko, Y., 2012. What can survey forecasts tell Us about information rigidities? *J. Polit. Econ.* 120 (1), 116–159.
- Coibion, O., Gorodnichenko, Y., 2015a. Information rigidity and the expectations formation process: a simple framework and new facts. *Am. Econ. Rev.* 105 (8), 2644–2678.
- Coibion, O., Gorodnichenko, Y., 2015b. Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation. *Am. Econ. J.: Macroecon.* 7 (1), 197–232.
- Coibion, O., Gorodnichenko, Y., Kumar, S., 2018. How do firms form their expectations? New survey evidence. *Am. Econ. Rev.* 108 (9), 2671–2713.
- Coibion, O., Gorodnichenko, Y., Ropele, T., 2020. Inflation expectations and firm decisions: new causal evidence. *Q. J. Econ.* 135 (1), 165–219.
- Coibion, O., Gorodnichenko, Y., Kumar, S., Pedemonte, M., 2020. Inflation expectations as a policy tool? *J. Int. Econ.* 124, 103297 Article.
- D'Acunto, F., Malmendier, U., Ospina, J., Weber, M., 2021. Exposure to daily price changes and inflation expectations. *J. Polit. Econ.* 129 (5), 1615–1639.
- Dovern, J., Müller, L.S., Wohlrabe, K., 2020. How do firms form expectations of aggregate growth? New Evidence from a Large-Scale Business Survey. CESifo Working Paper Series 8179.
- Gabaix, X., 2011. The granular origins of aggregate fluctuations. *Econometrica* 79, 733–772.
- Gaballo, G., 2018. Price dispersion, private uncertainty, and endogenous nominal rigidities. *Rev. Econ. Stud.* 85 (2), 1070–1110.
- Gautier, E., 2008. The behaviour of producer prices. evidence from French PPI micro-data. *Empir. Econ.* 35 (2), 301–332.
- Kiley, M., 2016. Policy paradoxes in the New Keynesian model. *Rev. Econ. Dyn.* 21, 1–15.
- Kumar, S., Afrouzi, H., Coibion, O., Gorodnichenko, Y., 2015. Inflation targeting does not anchor inflation expectations: evidence from firms in New Zealand. *Brook. Pap. Econ. Activ.* 46 (2 (Fall)), 151–225.
- Lorenzoni, G., 2009. A theory of demand shocks. *Am. Econ. Rev.* 99 (5), 2050–2084.
- Lucas, R., 1972. Expectations and the neutrality of money. *J. Econom. Theory* 4 (2), 103–124.
- Mackowiak, B., Wiederholt, M., 2009. Optimal sticky prices under rational inattention. *Am. Econ. Rev.* 99 (3), 769–803.
- Mackowiak, B., Moench, E., Wiederholt, M., 2009. Sectoral price data and models of price setting. *J. Monet. Econ.* 56 (S), 78–99.
- Nimark, K., 2014. Man-bites-dog business cycles. *Am. Econ. Rev.* 104 (8), 2320–2367.
- Pasten, E., Schoenle, R., 2016. Rational inattention, multi-product firms and the neutrality of money. *J. Monet. Econ.* 80 (C), 1–16.
- Vermeulen, P., Dias, D., Dossche, M., Gautier, E., Hernando, I., Sabbatini, R., Stahl, H., 2012. Price setting in the Euro area: some stylised facts from individual producer price data. *J. Money Credit Bank.* 44 (8), 1631–1650.
- Yang, C., 2019. Rational Inattention, Menu Costs, and Multi-Product Firms: Micro Evidence and Aggregate Implications. manuscript. Available at https://choongryulyang.github.io/research/Choongryul_Yang_JMP.pdf.