Comment on "No Firm Is an Island? How Industry Conditions Shape Firms' Expectations" by Philippe Andrade, Olivier Coibion, Erwan Gautier and Yuriy Gorodnichenko

Isabelle SALLE*

June 24, 2021

Abstract

Andrade et al. analyze the data from a large panel of firms and show how sector-specific shocks erroneously influence firms' outlooks about the aggregate economy. This paper is part of a broader research agenda that has extensively documented the existence of information rigidities using data from large-scale surveys of households and firms. This comment discusses the implications of this research agenda for central banking in light of a complementary empirical method, namely laboratory experiments, and their use in informing macroeconomic policies.

JEL classification: E7, E5, E3, E4.

Keywords: Information rigidities, Expectations, Survey data, Experimental macroeconomics.

1 Introduction

Central banks (hereafter, CBs) are mainly concerned with controlling inflation, which itself ultimately stems from the price-setting behavior of a

*Bank of Canada, Universiteit van Amsterdam and Tinbergen Institute, the Netherlands (ISalle@bank-banque-canada.ca). I thank Jason Allen, Camille Cornand, Cars Hommes, Janet Jiang and Oleksiy Kryvtsov for their useful contributions to this discussion. The views expressed in this note are mine and do not necessarily reflect those of the Bank of Canada.

collection of heterogeneous firms operating in the distinct sectors that compose the aggregate economy. Therefore, understanding how firms *actually* set prices constitutes an essential, yet under-documented, ingredient of useful modeling and sound policymaking.

The paper by Philippe Andrade, Olivier Coibion, Erwan Gautier and Yuriy Gorodnichenko (hereafter, ACGG), which I summarize in Section 2, constitutes a major contribution towards achieving this goal. The ACGG paper falls within a prolific research agenda that aims to collect empirical evidence on how agents form expectations and make decisions, evidence which is used to provide empirical microfoundations for macroeconomic models. Section 3 discusses this research agenda in light of a complementary empirical method, namely controlled laboratory experiments. These are specifically designed to shed light on the behaviors of human subjects, particularly their expectation formation processes. To this end, I will contrast the 'survey route' with the 'laboratory route'. Section 4 focuses on the implications of this literature for CB research. Section 5 briefly concludes.

2 The ACGG paper in a nutshell

ACGG analyze the data from a quarterly survey of French firms to document a particular form of information friction, namely the influence of industry-specific conditions on firms' expectations about *aggregate* outlooks. Let me first say a few words on the survey.

The survey elicits expectations of both individual and aggregate variables and has been doing so for the past 30 years. As such, the historical data encompass a wide range of macroeconomic conditions. The respondents form a representative sample of about 2,500 medium to large-size firms in the French manufacturing sector. Taken together, these factors provide the survey with a rich panel dataset. On a firm level, the questions asked broadly cover company-specific production, prices and sales. On an industry level, questions relate to companies' expectations concerning the short-run evolution of production, exports and prices. Finally, views are also solicited regarding the outlook for aggregate inflation. Questions regarding expectations have both qualitative and quantitative elements.

Moreover, the survey frames questions simply, which keeps the cognitive load of the respondents reasonable and the response rate high. Large firms are also over-represented in the sample, which suggests a higher level of sophistication on average. From these features, one can expect these self-reported data to be sufficiently reliable for research purposes, which the authors verify before engaging further.

The main conclusion of the ACGG paper is that firms tend to wrongly treat industry-specific information as relevant for aggregate outlooks. This

finding, which I will refer to as a 'generalization bias', results in information rigidities: firms' expectations of aggregate conditions adjust to both industry-specific and aggregate shocks, while a fully-informed rational-expectation (FIRE) agent would discard industry-specific shocks, which are orthogonal to the economy as a whole, when predicting aggregate conditions. This main finding is robust to a wide range of specifications, including when accounting for potential transmission delays from sector-specific shocks to the whole economy.

While this class of biases and the resulting violations of the FIRE benchmark have been documented with respect to households and managers – including in some of the authors' previous works – ACGG find the first direct, cross-industry empirical evidence among firms. The authors argue that this new stylized fact provides an empirical foundation for macroeconomic models with information frictions; in particular, it provides evidence for the mechanism underlying nominal frictions in the class of macroeconomic models known as 'island models'. One could also think further about what this particular friction entails, for instance regarding welfare-costly price dispersion.

Additionally, ACGG show that this generalization bias is more pronounced among firms operating in relatively volatile industries. The authors interpret their findings in light of the information rigidities resulting from rational inattention theory: firms devote more of their limited resources to processing and reacting to industry-specific information than to aggregate information because industry conditions are more volatile and, hence, more relevant to their profits than aggregate conditions. Yet, the survey does not allow for the identification of the shocks perceived by the firms. Hence, it is not possible to take a clear stand on whether this bias arises from noise (i.e. from a friction in the firm's *access* to information; see, again, Fig. 1) or from confusion (i.e. from a friction at the stage of information *processing*).

ACGG further show how the firms' adjustment of their own inflation expectations and their own actual pricing decisions in response to changes in the industry price level is rapid, while these same adjustments are sluggish when it comes to changes in the aggregate price level. This asymmetry in responding to local versus global shocks constitutes an information friction that amplifies the persistence of aggregate shocks.

ACGG also discuss this generalization bias as an explanation of the observed level of heterogeneity in firms' aggregate expectations. Indeed, since conditions differ widely between industries, and since industry-specific information interferes with expectations about aggregate conditions, the resulting cross-sectional dispersion of aggregate expectations among all firms should be higher than if firms were to base these expectations only on aggregate shocks. However, the paper does not report on disagreement within- versus between-industries, which could strengthen their explanation.

The ACGG paper is part of a broader research agenda that I will now discuss. In particular, I will show how laboratory experiments may help complement survey-based work.

3 A roadmap to empirical evidence on expectations and behaviors

Figure 1 depicts the decision process of an economic agent. As the FIRE paradigm assumes no friction at any step of this process, it constitutes the theoretical benchmark against which one can assess how real-world expectations are formed and actual economic decisions are made. This allows for the identification of frictions and behavioral biases. In this vein, an influential strand of the macroeconomic literature has been concerned with shedding light on how agents process information to form expectations and make economic decisions by relying on empirical (micro-)data. Here I will specifically discuss the '*survey route*' of this literature – to which the ACGG paper belongs – alongside the '*laboratory route*' and show how the two may be viewed as complementary. I will focus on the expectation formation step, as pictured in Figure 1, because this is probably the most relevant one for CB research as discussed in Section 4.

The survey route The first route exploits large-scale survey data of households and firms – sometimes combined with randomized controlled trial (RCT) experiments – to empirically investigate i) how agents form their expectations, ii) how these expectations map into their economic and financial decisions and iii) how this empirical evidence may discipline theoretical models.[1] Crucially, the surveys must contain questions about both expectations and decisions. The combination of the recurring nature of surveys and their large numbers of respondents allows them to generate unique and representative panel data of real-world expectations and actual decisions, hence providing a high degree of external validity. This realism, however, comes at the cost not only of relying on self-reported data, but also of control.

Indeed, making inferences about agents' behaviors – and in particular expectation formation processes – in light of economic models requires the ability to at least somewhat control the three environmental factors pictured in Figure 1: i) the information that the agents use, ii) the incentives that the agents face and iii) the structure of the environment in which they operate. The use of survey studies may make it more difficult to control

[1] See, in particular, Coibion and Gorodnichenko (2012) and the dozen subsequent contributions by these authors, including the present ACGG paper.

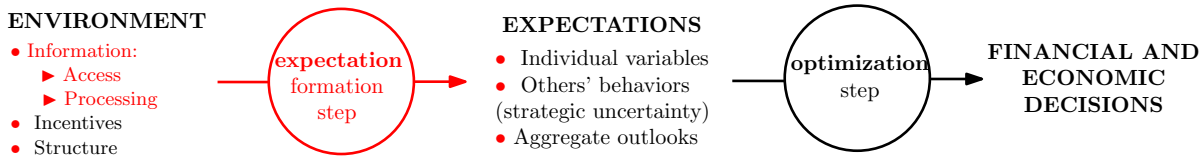

Figure 1: Schematic representation of the decision-making process of an economic agent

Note: The mechanism studied in ACGG and discussed in Section 3 in light of the experimental method is depicted in red. The FIRE benchmark assumes no friction in any step of the process.

these dimensions. As surveys are conducted during daily life, respondents may be influenced by unobservable and idiosyncratic factors, such as the most recent events in their environment. Because the quality of their answers does not affect their compensation, surveys do not guarantee their attention or cognitive effort. Regarding forecast elicitation in particular, it can be uncertain whether participants aim to minimize their forecast errors by answering with their true beliefs, or whether they are reporting what they think are consensus answers – or even what they heard from the media just before filling out the survey. Furthermore, the researcher cannot easily evaluate the participants' understanding of the questions.

Laboratory experiments offer a high degree of control on precisely these three key ingredients of economic decision-making. Before providing an example of how a laboratory experiment could build upon the results of the ACGG survey paper, let me first briefly characterize how the lab is used in the context of macroeconomic studies.

The laboratory route A macroeconomic experiment can be defined as one that tests either the predictions of a macroeconomic model or its underlying assumptions.[2] In the context of the decision-making process depicted in Figure 1, the lab technique allows us to break down and examine each step separately. So-called learning-to-forecast experiments are group experiments that focus on the expectation component of macroeconomic models, usually within well-defined inflation-targeting models, where the aggregate depends on individual expectations.[3]

Although both scholars and practitioners agree on the importance of expectations, there is considerable debate over how they should be modeled

[2]Since the first market experiments in the classroom by Chamberlin (1948), experimental macroeconomics has grown steadily as an area of research; see, e.g., Duffy (2016) for a survey of the field.

[3]A short, very much non-exhaustive list of these includes Hommes et al. (2019); Assenza et al. (2020); Kryvtsov and Petersen (2020).

and they are not easily observed in the field. In this controlled environment, the expectations are the only degree of freedom: they are elicited from the participants so as to take an agnostic stand on their formation process. Furthermore, the instructions and the interface of the game are the only source of information with which to complete the task. A straightforward payoff mechanism and monetary rewards complete the environment by controlling the incentives of the participants.

What one loses in terms of realism due to the stylized lab environment may thus be compensated for by the tight control of the environment in which expectations are elicited. The close link to the theoretical model facilitates causal inference, while the laboratory setting allows one to control for confounding factors and collect 'clean' data on expectations. In particular, these lab experiments make it possible to isolate the expectation formation step from the maximization problem, while survey data collect actual decisions that are the product of both steps taken simultaneously. Moreover, forecasts elicited in the lab have been found to share key statistical moments with those collected in surveys involving professional forecasters, individual households and firms (Cornand and Hubert, 2020).

Hence, lab data may help understand how information and policies influence expectations and, therefore, economic outcomes. I will now discuss the added value of lab experiments with respect to expectation formation in the context of the findings by ACGG.

To give a hypothetical example, one could design an 'island' lab economy in which subjects, distributed across distinct markets, would be tasked with selling an endowment. In the process of doing so, they would experience specific price shocks that are chosen so as to exactly cancel out when computing the aggregate price level across all markets. A simple data-generating process could be selected for this purpose. The experimenter could then elicit subjects' expectations of the aggregate price (possibly in combination with a price-setting task) and test whether subjects who experience inflationary (deflationary) shocks tend to have higher (lower) aggregate price expectations. The following are just three potential ways in which such a lab experiment could usefully complement the findings of ACGG.

First, while researchers using existing surveys are constrained by the established set of questions and their formulation, lab experiments allow us to design our own. In the present context, one could elicit quantitative point expectations and even probability distributions. This way, the relative size of the reaction of agents' expectations to sector-specific and aggregate shocks could be directly compared and any deviation with respect to the FIRE benchmark could be precisely estimated.

Second, one dimension of interest for us is how the generalization bias identified in the survey of French firms relates to the state of the economy. Is it more pronounced if the actual level or volatility of inflation is relatively

high? The control over fundamentals in the lab implies that the shocks need not be identified, as they are drawn *ex ante*. The distribution of the shocks could be a treatment variable to test this hypothesis. Additional treatment variables could be other dimensions that may affect this friction, such as the nature and quantity of information or the market structure.

Third, thanks to the control offered by the lab environment, one could explore the source of this generalization bias. The authors argue that their results support the theory of rational inattention. While this is a reasonable and appealing conclusion, one could equally argue that these results reflect bounded rationality and cognitive biases that may result in the use of simple heuristics, such as a form of the representativeness heuristic.[4] Under this class of explanations, firms would be naturally inclined to generalize from their own, industry-specific experience to draw inference about the aggregate economy. To measure the relative contribution of these two explanations to the observed information friction, one could design a payoff function that emphasizes participants' aggregate forecast errors rather than individual market outcomes.

Finally, it is worth noting that surveys are larger and more complex than lab experiments, which are comparatively easy to implement. These practical considerations matter when it comes to informing policy, which moves faster than academic research.

I will now discuss the relevance of laboratory experiments, when used to complement survey data, for CB research.

4 Implications for central bank research

In the context of CB research, understanding the behavior of agents in the real world – in particular, their expectation formation processes – is essential for at least three reasons. First, identifying frictions helps us gain an understanding of the observed economic dynamics, particularly the propagation of shocks and the effects (or lack thereof) of policies. An immediate example that comes to my mind is the attenuated effect of forward-guidance policies with respect to what FIRE models predict, which may be accounted for precisely by information friction and a resulting muted expectation channel; see, e.g., Baeriswyl et al. (2021) for the study of this mechanism in a lab experiment.

Second, being able to take our models seriously requires us to empirically test their hypotheses and incorporate behavioral features identified in real-world data, whether from surveys or from the lab. This feedback between

[4]The work on cognitive biases and the use of simple heuristics in individual decision-making was pioneered by the psychologists Amos Tversky and Daniel Kahneman beginning in the late 1960s.

theory and empirics enhances the ability of our models to make forecasts and predict the consequences of policy alternatives. In this respect, an empirical validation of macroeconomic models' expectation components – upon which most of the effects of macroeconomic policies hinge – is a pressing issue for CB research.

CB communication has helped manage expectations and, *in fine*, economic decisions. This began with the implementation of inflation targeting regimes and has in recent years complemented unconventional monetary policies at the effective lower bound. As such, targeted communication policies have become a key part of central banks' toolkit. This leads us to the third, broader, objective of this line of research, for which lab experiments may end up being particularly insightful: the systematic comparison of policy alternatives.

In this context, the lab can be seen as a kind of wind tunnel with which to test big ideas on a small scale.[5] We can envision the laboratory test as a necessary, but not sufficient, condition to give the green light to a policy: if a theory has no explanatory power or if a policy does not work in a stylized and controlled environment, we can reasonably cast doubt on whether it would work outside the lab, in the admittedly much more complex real-world economies.

For instance, coming back to the ACGG paper, the information friction identified is welfare-costly because it aggravates price dispersion, increases uncertainty and weakens the expectation channel of policies. Experimenting with various forms of information provision in the survey of French firms would be feasible but cost a great deal of time, while the amount of information that is communicated to the subjects or the way in which it is communicated would be easier to manipulate in the lab. If observed in the lab, one could experiment with different communication strategies to evaluate whether this generalization bias – and the resulting disagreement between agents – could be dampened and forecasts improved. Furthermore, as expectations are also policy-dependent, survey data may be prone to confounding factors, while independent observations in the lab circumvent this issue.

Another topical example is the design of so-called 'make-up strategies'. Indeed, over the last decade, persistently low inflationary pressures have resulted in price levels falling behind the paths consistent with the inflation target. One commonly discussed alternative is average-inflation targeting (hereafter, AIT), which aims to compensate past missed targets through higher current and future inflation rates.

In theory, the merits of make-up strategies depend on how we model ex-

[5]This is a role that, in theory, could also be filled by macroeconomic field experiments, but those are difficult to implement and hard to justify ethically.

pectations. While FIRE agents are assumed to fully understand how monetary policy is conducted – that is, to hold model-consistent expectations – real-world expectations need not align with the way monetary policy is conducted. Learning-to-forecast experiments may shed light on these make-up strategies.

For example, Arifovic and Petersen (2017) show that state-dependent targets, no matter how they are communicated, fail to drive subjects' inflation expectations out of a liquidity trap. The main mechanism is the lack of credibility: subjects need to 'see it to believe it'; in other words, higher inflation must first be realized for them to raise their inflation expectations, which is at odds with the FIRE benchmark.

In a horse-race experiment between make-up strategies, Salle (2021) shows that the negative auto-correlation in the inflation process under AIT is not easy to learn for subjects. In other words, subjects have a hard time averaging inflation across periods and fail to integrate the correct number of lags in their inflation expectations. This entails more diverse and more volatile expectations and resulting inflation than under IT, even outside liquidity-trap episodes.

These doubts on the ability of AIT to reshape inflation expectations towards higher rates echo the results of another recent survey experiment by some of the authors (Coibion et al., 2020). The lab may then complement their findings: if households' inflation expectations fail to adjust upward in the wake of the AIT regime shift announcement in the US, this may be either because they do not immediately believe the announcement, and/or because they fail to see how exactly the new monetary policy rule should be reflected in their expectations.

There are many other pressing CB issues that could be addressed by the combined use of lab and survey work. Whether one approach or the other is better suited depends on the question being asked. For instance, one question of particular interest concerns the implementation, effects and eventual tapering of unconventional monetary policies such as quantitative easing and yield-curve control, the relative merits of which could be evaluated in a stylized lab environment. Another example is the real effects of expected inflation, which has been studied in a simple laboratory economy by Jiang et al. (2021). Yet another question is whether the public understands why CBs are seeking higher inflation, which could potentially be studied by employing an experimental design in surveys.

5 Conclusions

I wish to conclude as I started, by stressing that the ACGG paper constitutes a major advancement in our understanding of expectation formation and

decision making. I hope I have successfully conveyed the idea that the broader research agenda to which this paper belongs – including laboratory experiments – has the potential to guide CBs in the challenges they face within a 'low for long' environment.

References

Jasmina Arifovic and Luba Petersen. Stabilizing expectations at the zero lower bound: Experimental evidence. *Journal of Economic Dynamics and Control*, 82:21–43, 2017.

Tiziana Assenza, Peter Heemeijer, Cars Hommes, and Domenico Massaro. Managing Self-Organization of Expectations through Monetary Policy: a Macro Experiment. *Journal of Monetary Economics*, 117:170–186, 2020.

Romain Baeriswyl, Kene Boun My, and Camille Cornand. Double overreaction in beauty contests with information acquisition: Theory and experiment. *Journal of Monetary Economics*, 118:432–445, 2021.

Edward Chamberlin. An Experimental Imperfect Market. *The Journal of Political Economy*, 56(2):95–108, 1948.

Olivier Coibion and Yuriy Gorodnichenko. What Can Survey Forecasts Tell Us about Information Rigidities? *Journal of Political Economy*, 120(1): 116–159, 2012.

Olivier Coibion, Yuriy Gorodnichenko, and Edward S. Knotek. Average Inflation Targeting and Household Expectations. Working Papers 202026, Federal Reserve Bank of Cleveland, Sep 2020.

Camille Cornand and Paul Hubert. On the external validity of experimental inflation forecasts: A comparison with five categories of field expectations. *Journal of Economic Dynamics and Control*, 110:1037–46, 2020.

John Duffy. Macroeconomics: A Survey of Laboratory Research. In John H. Kagel and Alvin E. Roth, editors, *The Handbook of Experimental Economics, Volume 2*, Introductory Chapters. Princeton University Press, 2016.

Cars Hommes, Domenico Massaro, and Isabelle Salle. Monetary and fiscal policy design at the zero lower bound: Evidence from the lab. *Economic Inquiry*, 57(2):1120–1140, 2019.

Janet H. Jiang, D. Puzzello, and C. Zhang. Inflation, output and welfare in the laboratory. Technical Report forthcoming, Bank of Canada Staff Working Paper Series, 2021.

Oleksiy Kryvtsov and Luba Petersen. Central Bank Communication That Works: Lessons from Lab Experiments. *Journal of Monetary Economics*, 117:760–780, 2020.

Isabelle Salle. What to Target? Insights from Theory and Lab Experiments. Staff working paper, forthcoming, Bank of Canada, 2021.