

Inflation Expectations, Learning and Supermarket Prices

Evidence from Survey Experiments*

Alberto Cavallo

Guillermo Cruces

Ricardo Perez-Truglia[†]

MIT and NBER

CEDLAS-FCE-UNLP and IZA

Microsoft Research

This Draft: May 2016. First Draft: December 2013.

Abstract

Information frictions play a central role in the formation of household inflation expectations, but there is no consensus about their origins. We address this question with novel evidence from survey experiments. We document two main findings. First, individuals in lower-inflation contexts have significantly weaker priors about the inflation rate. This finding suggests that rational inattention may be an important source of information frictions. Second, cognitive limitations also appear to be a source of information frictions: even when information about inflation statistics is made readily available, individuals still place a significant weight on less accurate sources of information, such as their memories of the price changes of the supermarket products they purchase. We discuss the implications of these findings for macroeconomic models and policy-making.

JEL Classification: C93, D83, D84, E31, E58.

Keywords: inflation expectations, survey experiment, rational inattention, supermarket.

*We would like to thank Robert Barro, Raj Chetty and David Laibson for their valuable input at the early stages of the project. We also thank useful comments from Alberto Alesina, Nageeb Ali, Rüdiger Bachmann, Christian Borgs, Sebastian Di Tella, Emmanuel Farhi, Matthew Gentzkow, N. Gregory Mankiw, Markus Mobius, Andrés Neumeyer, Roberto Rigobon, Tanya Rosenblat, Guido Sandleris, Tavneet Suri, Martin Tetaz, Glen Weyl, Fernando Yu, and participants in the seminars at Harvard University, MIT Sloan, Microsoft Research New England, Universidad de San Andres, Universidad Torcuato Di Tella, the Chicago-NYU International Macro Finance Conference, the 2016 AEA Meetings, the 6th Ifo Conference on Macroeconomics and Survey Data, and the NY Federal Reserve Subjective Expectations Conference. Julián Amendolaggin and Nicolás Badaracco did excellent work as research assistants. We would also like to thank Tomás Pessacq and Carolina Yellati for their collaboration in conducting the experiments, and MIT Sloan and CEDLAS-UNLP for their funding. This project was reviewed and approved by the Committee on the Use of Humans as Experimental Subjects at MIT.

[†]Corresponding author: rtruglia@microsoft.com. Microsoft New England Research and Development (NERD) Lab, Office 12073, 1 Memorial Drive, Cambridge MA 02142.

1 Introduction

Expectations about macroeconomic variables play an essential role in economic theory and policy-making. Consumer inflation expectations, in particular, are key to understanding household consumption and investment decisions, and ultimately the impact of monetary policies. Although central banks seek to influence expectations, there is no consensus in the empirical literature on how household inflation expectations are formed or can be affected (See Bernanke, 2007; Bachmann et al., 2015; Coibion and Gorodnichenko, 2015).

Consumer surveys indicate that household inflation expectations tend to be much more heterogeneous than those of professional forecasters (Ranyard et al., 2008; Armantier et al., 2013). Two main explanations for this degree of dispersion have been given in the literature. Some authors attribute it to rational inattention, according to which individuals only partly incorporate information on topics such as inflation because acquiring that information is costly (relative to the potential gains from using that information). This explanation is particularly convincing in contexts of low inflation like the United States, where the potential financial cost of ignoring inflation is negligible for most households. Other authors argue that, in forming inflation expectations, individuals use information derived from their personal experience as consumers, which can be both diverse and inaccurate (Bruine de Bruin et al., 2011; Malmendier and Nagel, 2016; Madeira and Zafar, forthcoming). The existing evidence on information frictions cannot distinguish between different sources of frictions. This distinction can be important, to the extent that different sources can lead to very different policy prescriptions. We present evidence from a series of experiments specifically designed to test the hypotheses of rational inattention and personal consumer experience.

In a series of online and offline surveys, we randomly provided subjects with information related to past inflation and measure the effects of the information provided on the subjects' inflation expectations. We provide information about inflation from different sources, such as inflation statistics and tables with historical prices of specific supermarket products. With the help of a Bayesian learning model, we can estimate how much weight subjects give to a given piece of information – e.g., an inflation statistics – relative to their prior beliefs about inflation.

The first goal of the paper is to provide a sharp test of the rational inattention model. To do so, we conducted survey experiments in both a context of low inflation – the United States, with an average annual inflation rate of 1.8% in the five years prior to our study – and in a context of high inflation – Argentina, where the average annual inflation rate over the same time period was around 22.5%.¹ According to the rational inattention model, individuals in a context of higher inflation should have stronger priors about inflation, because the financial cost of misperceiving inflation is higher. They should thus acquire information of higher quality, and do so more often (Mankiw et al., 2003; Carroll, 2003). Consistent with this hypothesis, we find that individuals in

¹We do not use official inflation statistics for Argentina because they were widely discredited. We use instead alternative indicators compiled by the private sector, which are well known and widely cited in the media.

the lower-inflation context have weaker priors about the inflation rate. For example, when provided with information about inflation statistics or prices of specific supermarket products, individual in the low-inflation context (United States) assigned a weight of just 15% to their prior beliefs, while individuals from the high-inflation context (Argentina) assigned a weight of roughly 50%.

The second goal of the paper is to measure whether cognitive limitations may also be an important source of information frictions. To do so, we compare how individuals incorporate two types of information about inflation: inflation statistics, and historical prices for a handful of randomly selected supermarket products (conceived as a proxy for the type of information that individuals would obtain from their own personal experience as shoppers).² Relative to the average price change from a random set of six products, inflation statistics are extremely more precisely estimated. As a result, when confronted to these two pieces of information, we would expect an economic professional to put all weight on the statistics information and ignore the supermarket prices information. Instead, when subjects were provided with these two types of information simultaneously, they implicitly assigned as much weight to supermarket prices than to inflation statistics. In other words, even when information about inflation statistics is made readily available to them, individuals still place significant weight on less accurate sources of information.³

Our experimental design tries to address what we believe is one of the most common criticism to survey experiments: instead of inducing genuine learning, the information provided in the experiment may elicit spurious reactions. For instance, if an individual is told that the annual inflation rate was 2% and then later on is asked about her inflation expectations, she may report an inflation expectation that is closer to 2% for spurious reasons, such as to please the interviewer (Goffman, 1963), avoid being perceived as ignorant, or because of unconscious numerical anchoring (Tversky and Kahneman, 1974).⁴ Indeed, to illustrate how easy it is to manipulate the subjects' responses, we show that providing explicitly-fictitious information on price changes has an statistically and economically significant effect on inflation expectations.

Our experimental design includes two methods for disentangling how much of the reaction to the information provided corresponds to genuine learning instead of spurious learning. The first method exploits the fact that, if the reaction to the information was spurious, then the experimental effects should not persist months after the information provision. The second method exploits the fact that, if the reaction to the information was spurious, then we should not observe effects on expectations about other nominal variables that are intrinsically related to the inflation rate, such as the the nominal interest rate. Results from these two methods suggest that concerns about spurious learning are justified and must be taken seriously, since half of the reaction to our

²The data was scraped off the websites of some of the largest supermarkets in the United States and Argentina as part of the Billion Prices Project at MIT.

³This result is also consistent with survey evidence presented by Bruine de Bruin et al. (2011), who show that, when asked about the inflation rate, most individuals report that they try to recall the prices of specific products.

⁴This criticism is common to survey experiments in general, not particularly to our application in the area of inflation expectations. See Rosenthal (1966) for a discussion of effects of this type in behavioral research, and Zizzo (2010) for a recent application to experimental economics.

informational treatments is spurious. Nevertheless, our main results remain unchanged after we control for spurious learning.

Another concern is that subjects may have reacted to the information on supermarket prices that we provided because they perceived it as accurate, but they may still not trust their own memories about supermarket prices. Also, using price memories in the formation of inflation expectations is misleading only insofar as those memories are inaccurate. Addressing these remaining questions requires data that would be difficult to collect in an online survey. Thus, we conducted a unique consumer-intercept survey experiment to address these questions. Among other things, we recorded consumers' purchases by scanning participants' supermarket receipts, which we linked to data on the actual historical prices of those same products at the same store. We also asked respondents to recall historical prices for a random selection of the items that they had just purchased, which allowed us to generate exogenous variation in the salience of the subjects' own price memories. The evidence from this experiment suggests that individuals do use their own memories about supermarket prices when forming inflation expectations, despite of the fact that those memories are largely inaccurate and thus induce large errors in expectations.

Our findings provide useful lessons for macroeconomic theory. The idea that monetary policy can have real effects due to information frictions goes back to Phelps (1969) and Lucas (1972). More recently, Mankiw and Reis (2002) show how the New Keynesian Phillips Curve can be the product of sticky information. The policy prescriptions can depend sensibly on how we model the information frictions, but there is no consensus about what the right model may be (Coibion and Gorodnichenko, 2012). Our evidence suggests that, in addition to rational inattention models, the literature should also incorporate cognitive limitations.⁵

Our findings are also related to recent debates about central bank transparency. Some authors argue that information disclosure can enhance welfare (Hellwig, 2005), while others argue that it can reduce welfare (Morris and Shin, 2002). Our findings suggest that, even when the statistics are publicly and readily available, households use less accurate private information. This implies that, in addition to the dissemination of aggregate statistics, central banks may have an additional policy margin in terms of communicating how objective, precise and representative these statistics are. For example, the European Central Bank and the French statistical agency have made notable efforts to create online tools to convey this information and the way it is collected and processed in a user-friendly way.⁶ Central banks interested in affecting expectations could also consider disseminating information about the price changes of specific products, which individuals can find easier to relate to. All these efforts may help central banks increase the speed with which individuals react to monetary policy, and help households make better financial decisions (Armantier et al.,

⁵The literature on memory in psychology and behavioral economics provide useful models for these cognitive limitations: see for example Mullainathan (2002) and Gennaioli and Shleifer (2010).

⁶See <http://www.ecb.europa.eu/ecb/educational/hicp/html/index.en.html> and http://www.insee.fr/en/indicateurs/indic_cons/sip/sip.htm, respectively.

2013).⁷

Our paper belongs to a literature that tries to understand the formation of household inflation expectations. A group of studies measure the role of inflation statistics, exploiting media coverage of statistics (Lamla and Lein, 2008; Badarinza and Buchmann, 2009; Drager, 2011), the publication of official statistics (Carrillo and Emran, 2012), and information-provision experiments (Roos and Schmidt, 2012; Armantier et al., 2016). Other studies have looked at the role of personal experiences. For instance, there is suggestive evidence that individuals use information from their own price memories (Bates and Gabor, 1986; Bruine de Bruin et al., 2011; Coibion and Gorodnichenko, 2015) and that individuals place excessive weight on information about past inflation levels experienced in their lifetime (Malmendier and Nagel, 2016).

The existing evidence in the literature on the formation of inflation expectations cannot distinguish between the different sources of information frictions (Ranyard et al., 2008). First, there is evidence that individuals fail to incorporate all the available information (e.g., Mankiw et al., 2003; Armantier et al., 2016), which some authors interpret as evidence of rational inattention. However, the same result would emerge if individuals irrationally incorporated information from inaccurate sources. Second, there is also evidence that individuals use these inaccurate sources when forming their inflation expectations (e.g., Bruine de Bruin et al., 2011; Malmendier and Nagel, 2016), which some authors interpret as evidence of cognitive limitations. However, this result is also consistent with a model of rational inattention, according to which, if the stakes of misperceiving inflation are low, individuals should rationally use information that is inaccurate as long as it is also costless.⁸ Our contribution to this literature is to design experiments to test each of the two information frictions, rational inattention and irrational learning, by exploiting variations in stakes (i.e., contexts of high vs. low inflation) and sources of information (i.e., inflation statistics vs. supermarket prices).

Methodologically, our paper is related to a recent subset of the literature that employs survey experiments to investigate household inflation expectations. For example, studies by Roos and Schmidt (2012) and Armantier et al. (2016) examine how individuals react to information about U.S. inflation statistics by adjusting their reported inflation perceptions. Bruine de Bruin et al. (2011) show that subjects who are asked to think about products with extreme price changes tend to report higher inflation expectations. We contribute to this literature by extending these methods to answer novel questions about the sources of information frictions. Additionally, we make a number of methodological contributions, such as disentangling genuine from spurious learning and combining survey with administrative data to study how individuals learn about supermarket prices.

⁷The distribution of the bias is relevant as well. If poorer and less educated consumers had larger biases, as observed in many datasets, then correcting their biases may reduce these consumers' relative disadvantage.

⁸For example, Demery and Duck (2007) argue that individuals may optimally decide to use solely information they receive as a byproduct of their economic activity rather than complementing that information with official statistics.

The paper proceeds as follows. Section 2 describes the general experimental design. Section 3 presents evidence from a series of online experiments conducted in the United States and Argentina. Section 4 presents evidence from the consumer intercept survey experiment. The last section concludes.

2 Experimental Design

2.1 Structure of the Survey Experiments

In this section, we describe the experimental framework that will be used as the basis for all the empirical analysis provided in the rest of this paper. This framework builds upon a number of previous experimental studies (e.g., Bruine de Bruin et al., 2011; Roos and Schmidt, 2012; Armantier et al., 2016), but introduces innovations aimed at testing new hypothesis and addressing the concern of spurious learning.

The basic structure of the survey experiments is:

1. Eliciting subjects' inflation perceptions: i.e., the perception of the annual inflation rate over the previous twelve months. This constitutes the individual's prior belief ($\pi_{i,t}^0$ in the model in the following section).
2. Providing the subject with information related to the inflation rate over the previous twelve months, which constitutes the signal ($\pi_{i,t}^T$). In the case of the control group with no information provision, there is no signal. The different pieces of information provided to the subjects is described in the following subsection.
3. Eliciting subjects' expectations about inflation (i.e., the expected annual inflation rate over the following twelve months, $\pi_{i,t+1}$) and other nominal variables (e.g., the nominal interest rate, $i_{i,t+1}$). These expectations may be elicited right after the information provision, or several months later.

The main analysis consists of measuring how the information provided to individuals changes their expectations about the future. When eliciting inflation perceptions and expectations, we always refer to the general price level rather than to the prices of the goods purchased by the respondent.⁹ We did not provide any incentives for respondents to answer accurately (i.e., prizes for guessing the right figures). However, as shown by Armantier et al. (2012), there is a significant correlation between incentivized and non-incentivized beliefs on inflation expectations.

⁹Specifically, for the U.S. online experiment, we asked participants the following two questions, taken directly from the University of Michigan's Survey of Consumers: "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?" with three options: "Go up," "Stay the same" and "Go down." We then asked: "By about what percent do you expect prices to change, on average, during the next 12 months?" with an open numerical answer. For the Argentina online experiment, we opted to repeat the format of the question that had been asked in previous rounds of the opinion poll: "What do you think will be the annual inflation rate for the following 12 months?" (see the Appendix for exact wording in Spanish).

2.2 Treatment Arms

After eliciting past inflation perceptions, subjects were randomly assigned to either a control group (with no information) or one of four treatment arms. This section describes these treatment arms.

The snapshots of the informational treatments and the survey questions are shown in the questionnaire Appendix D. As a summary, Figure 1 provides samples of the information treatments in the U.S. Online experiment (see Figure B.1 for samples from the Argentina Online experiment). Our first treatment arm, shown in Panel (c) of Figure 1, aims to capture how individuals incorporate information from inflation statistics. This *Statistics (1.5%)* treatment consists of a table with the most recent statistics about annual inflation at the time of the survey, and is preceded by an explanation of what they are intended to measure (see the note to Figure 1 for the exact wording). The average of the three statistics indicated an annual average inflation rate of 1.5%, which was also displayed on the table.

Our second treatment arm was designed to capture the degree to which individuals use the information related to their everyday experience when forming inflation expectations, even if that information is not as representative and precise as aggregate inflation statistics. The *Products* treatment arm, illustrated in Panels (a) and (b) of Figure 1, presented respondents with a table containing the prices of six products at the time of the survey and one year earlier, as well as the price change (in percentage points) for each product and the average percentage change for all products presented in the table, also for the period from August 1, 2012 to August 1, 2013. This table was preceded by an explanatory paragraph (see the note to Figure 1 for the exact wording).

The products were selected from six broad types of goods (infant formula, bread, pasta and noodle-related products, cereals, sodas, and shampoos and related products). An algorithm selected the products in the specific tables so that the average price changes would be between -2% to 7% in 1 percentage point increments for a total of ten tables. The algorithm provided tables with products with different average price changes, but it also verified that other characteristics of the tables were roughly constant, leveraging on the availability of price histories for thousands of products and on detailed information on product characteristics. For instance, every table has one product from each of the six categories of goods, and the goods within each category have similar initial prices between tables (the algorithm selects different brands within product categories, since each brand experienced different price changes). This ensured that the initial price level and the representativeness of the products remain broadly comparable across tables. The information provided was entirely truthful, and a note to the table indicated that the products were taken from a large database with information on an existing branch of a large U.S. supermarket chain.¹⁰ There was no indication that the products in the table, or the average of price changes, were representative or that they reflected actual inflation levels. Respondents in this treatment arm were randomly assigned one of the ten tables with different average price changes, which we

¹⁰The data was scraped of the websites of some of the largest supermarkets in the United States and Argentina as part of the Billion Prices Project at MIT. See Cavallo (2013) for details.

indicate in parentheses after the *Products* treatment arm name in the rest of this paper. Panels (a) and (b) in Figure 1 illustrate the -2% and 2% cases respectively.

An additional treatment arm consisted of a combination of the previous two pieces of information: i.e., the respondent was shown the table with inflation statistics and one of the tables with prices for specific products. This is the *Statistics (1.5%)+Products* treatment arm. This was designed to test whether the tables with specific prices induced learning over and above the information conveyed by the inflation statistics.

Finally, we included a fourth treatment arm to gauge the degree of spurious learning, which we call the *Hypothetical* treatment. The respondents were asked to “eyeball” the price change of a product over a period of one year. We phrased the question in terms of the need to assess how comfortable the respondent was with questions about price changes. The table we provided contained only two prices at two points in time (January 1, 2012 and January 1, 2013), without specifying the product. The price of the hypothetical product changed from \$9.99 to \$10.99, a price increase of about 10% (panel (d) of Figure 1). If the number we introduced in the information provision stage, which was unrelated to reality, had any impact on stated inflation expectations, it would comprise evidence of spurious learning.

2.3 Estimating Learning Rates

In the following sections we present some reduced-form evidence on how individuals react to randomly assigned information, by comparing the raw distribution of inflation expectations (e.g., by means of a histogram) across individuals who were assigned to different treatment groups. The main advantage of this model-free approach is its transparency. Additionally, in this section we introduce a simple learning model that can allow us to summarize the reaction to the information in a single parameter that can be easily compared between experimental samples and information treatments.

We denote an individual’s perception of the annual inflation rate over the previous twelve months as $\pi_{i,t}$, and $\pi_{i,t+1}$, in turn, represents the individual’s expected annual inflation rate over the following twelve months. Individuals use information about (perceived) past inflation to form their expectations about future inflation:

$$\pi_{i,t+1} = f(\pi_{i,t}) \tag{1}$$

Note that this is a reduced-form model of expectations: this forecasting rule could represent an agent with rational expectations, an agent with adaptive expectations, or some other model of expectation formation.¹¹ None of the experiments that we conduct intend to distinguish between

¹¹The fact that individuals use information about the past to estimate future inflation may be suggestive of the models of adaptive learning (Sargent, 1993). However, the use of inflation perceptions to assess future inflation may also be consistent with rational expectations: e.g., some rational expectation models predict that inflation expectations follow an AR(1) process (Barr and Campbell, 1997).

these different interpretations, because we want to estimate a model of learning, not a model of expectation formation.

We consider a linear specification for $f()$:

$$\pi_{i,t+1} = \mu + \beta\pi_{i,t} \quad (2)$$

where β is the degree of pass-through from inflation perceptions to inflation expectations. Whether intentionally or not, a simple forward looking model like this seems to be a good strategy from the perspective of forming inflation expectations. For example, Atkeson and Ohanian (2001) report that, since 1984, the one-year-ahead inflation forecast of professionals in the U.S. has been no better than the “naïve” forecast of the inflation rate over the previous year.

Indeed, this linear specification fits the expectations and perceptions data very well (Jonung, 1981). For example, Figure 2 shows a robust linear relationship between perceived past inflation and expected future inflation for our online samples: with a regression coefficient of 0.782 in the United States (panel a) and a regression coefficient of 0.883 in Argentina (panel b).¹² Moreover, a great deal of the variation in inflation expectations can be explained by variation in inflation perceptions: in our U.S. sample, 29% of the variation in inflation expectations is due to variation in inflation perceptions, whereas the equivalent figure for our Argentine sample is 60%. In other words, a significant fraction of the disagreement about future inflation seems to be due to a disagreement about past inflation (see also Blanchflower and MacCoille, 2009). As a result, to understand the biases and dispersion in future inflation expectations, we need to understand the biases and dispersion in perceptions about past inflation.

The experiments we carried out consist of providing information related to past inflation. Let $\pi_{i,t}^0$ denote perceptions prior to the acquisition of new information, and let $\pi_{i,t}^T$ denote the signal from the information provided in the experiment. Any learning process – i.e., how individuals combine their prior knowledge and the new information to form their perceptions – can be represented by the following reduced-form equation:

$$\pi_{i,t} = g\left(\pi_{i,t}^0, \pi_{i,t}^T\right) \quad (3)$$

There are several plausible functional forms for $g()$. A simple and parsimonious alternative is to assume a Bayesian learning model with Gaussian distribution. Under this model, the prior belief is normally distributed with mean $\pi_{i,t}^0$ and standard deviation $\sigma_{i,t}^0$ (indeed, the distribution of reported inflation perceptions and expectation is distributed approximately Normal). The individual is presented with a signal about average inflation, $\pi_{i,t}^T$, which can be interpreted as the price change for one product randomly drawn from the universe of products. The population of price changes for all possible products follows a normal distribution with mean $\pi_{i,t}$ and standard deviation $\sigma_{i,t}^T$ (this functional form is also roughly consistent with the actual distribution of price changes). By

¹²This data is for subjects in the control group, i.e., those who were not provided any information about inflation.

construction, $\pi_{i,t}^{TRUE}$ is the actual inflation level – i.e., the average of price changes for all products. The precision of the signal is given by the inverse of $\sigma_{i,t}^T$, which is assumed to be known. Under these assumptions, the posterior belief is distributed normally with the following mean and variance:

$$\pi_{i,t} = \frac{\left(\frac{1}{\sigma_{i,t}^0}\right)^2}{\left(\frac{1}{\sigma_{i,t}^0}\right)^2 + \left(\frac{1}{\sigma_{i,t}^T}\right)^2} \pi_{i,t}^0 + \frac{\left(\frac{1}{\sigma_{i,t}^T}\right)^2}{\left(\frac{1}{\sigma_{i,t}^0}\right)^2 + \left(\frac{1}{\sigma_{i,t}^T}\right)^2} \pi_{i,t}^T, \quad \sigma_{i,t} = \sqrt{\frac{(\sigma_{i,t}^0 \cdot \sigma_{i,t}^T)^2}{(\sigma_{i,t}^0)^2 + (\sigma_{i,t}^T)^2}} \quad (4)$$

That is, the individual updates her perception based on an average between her prior belief and the realized signal:

$$\pi_{i,t} = (1 - \alpha_{i,t}) \pi_{i,t}^0 + \alpha_{i,t} \pi_{i,t}^T \quad (5)$$

where $\alpha_{i,t}$, the weight assigned to the new information, decreases with the accuracy of the prior belief $1/\sigma_{i,t}^0$ and increases with the accuracy of the signal $1/\sigma_{i,t}^T$. If $\sigma_{i,t}^0$ and $\sigma_{i,t}^T$ are constant across individuals, α is also constant across individuals. Replacing this expression into equation (2), the linearized version of the forward-looking equation, results in the following expression:

$$\pi_{i,t+1} = \gamma_0 + \underbrace{\gamma_1}_{\beta} \pi_{i,t}^0 + \underbrace{\gamma_2}_{\alpha\beta} (\pi_{i,t}^T - \pi_{i,t}^0) + \varepsilon_{i,t+1} \quad (6)$$

Note that the three elements in the above regression equation are all observed in our experimental data: $\pi_{i,t+1}$ is the respondent's stated inflation expectation (post-treatment), $\pi_{i,t}^0$ is the respondent's stated past inflation perception (pre-treatment), and $\pi_{i,t}^T - \pi_{i,t}^0$ is the difference between the signal provided in the informational treatment and the prior belief (defined to be zero for the control group). Thus, we can regress $\pi_{i,t+1}$ on $\pi_{i,t}^0$ and $\pi_{i,t}^T - \pi_{i,t}^0$ to estimate $\hat{\gamma}_1$ and $\hat{\gamma}_2$, and then use those parameters to estimate $\hat{\alpha}$ and $\hat{\beta}$ using the formulas $\hat{\beta} = \hat{\gamma}_1$ and $\hat{\alpha} = \hat{\gamma}_2/\hat{\gamma}_1$. We use the Delta Method to obtain the standard errors of $\hat{\alpha} = \hat{\gamma}_2/\hat{\gamma}_1$.¹³

The parameter β represents the rate of pass-through from perceptions of past inflation to future inflation expectations. The parameter α captures the weight the individual assigns to the information provided in the experiment relative to her prior belief. Intuitively, if the individual started with a prior belief of $\pi_{i,t}^0$ and the informational treatment provides a signal that inflation is $\pi_{i,t}^T$, the posterior belief can be expected to be between $\pi_{i,t}^0$ and $\pi_{i,t}^T$, and the parameter α reflects how much closer $\pi_{i,t}$ is to $\pi_{i,t}^T$ relative to $\pi_{i,t}^0$.

The following example illustrates the intuition behind our empirical model. Let us assume that, among individuals who receive no information from us, the correlation between inflation perceptions and expectations is 0.5: i.e., for each 1% increase in perceived past inflation, an individual believes that future inflation will be 0.5% higher. Now assume that we take a group of individuals who believed that past inflation was 10%, and we randomly provide some of them a

¹³One assumption is that the above OLS regression yields an unbiased estimate for β . Since $\pi_{i,t}^0$ is not randomized, at least in principle β could suffer from omitted variable bias, which in turn could bias the estimation of α .

signal that past inflation was 20%. If – relative to the control group – individuals who received the signal believe that future inflation is going to be 1% higher, that means that the information led them to believe that past inflation was 2% higher (i.e., $1/0.5$). In other words, the signal that past inflation was actually 20% increased their belief about past inflation from 10% to 12%. This indicates that, in forming her posterior belief, the individual assigned a 0.8 weight to the prior belief of 10% and a 0.2 weight to the signal of 20%: i.e., $12\% = 0.8 \times 10\% + 0.2 \times 20\%$.

This model of Bayesian learning makes a number of additional predictions that can be directly tested with the data. For instance, this model predicts that confidence in the posterior belief, $\sigma_{i,t}$, should be higher for individuals that were provided with relevant information. We present results for these tests in the results section and in the Appendix.¹⁴

2.4 Disentangling Genuine from Spurious Learning

A potential issue with our results is that, even if we find that the information provided has an effect on stated inflation expectations, individuals' reactions to this information may be spurious. In this section we present the framework attempting to quantify how much of α responds to genuine learning and how much to spurious learning.

Our first (and preferred) strategy consists of using data on the evolution of expectations obtained through follow-up surveys taken months after the original information provision. Numerical anchoring is, by definition, very short-lived, so we would not expect it to explain effects on beliefs measured months after the information was provided. Regarding interviewer pressure, months after the information provision it is most likely that subjects will not remember the information that was provided to them, so they should not be subject to pressure to agree with the interviewer.

Let $\pi_{i,t+1}^{follow-up}$ denote the inflation expectations elicited in a follow-up survey conducted months after the initial experiment, in which we did not provide any new information or reminded the subject about information provided in the past. Consider this new forward-looking equation: $\pi_{i,t+1}^{follow-up} = \mu_{FU} + \beta_{FU}\pi_{i,t}$, where β_{FU} is the degree of pass-through from inflation perceptions as stated in the original survey to inflation expectations stated in the follow-up survey. The estimate of β_{FU} should be lower than β , because β_{FU} is the product of β (i.e., pass-through from perceptions to expectations) and the pass-through from inflation perceptions in the first survey to inflation perceptions in the second survey (which is expected to be lower than one, because individuals should have incorporated more information in the meantime). In other words, for this estimate we do not need to assume that individuals do not learn new information between the two surveys, because that is already accounted for by the parameter β_{FU} .

If we combine the new forward-looking equation with the learning equation (5), we obtain:

¹⁴Armantier et al. (2016) also provide related tests of Bayesian learning in the context of household perceptions about inflation.

$$\pi_{i,t+1}^{follow-up} = \gamma_0 + \underbrace{\gamma_1}_{\beta_{FU}} \pi_{i,t}^0 + \underbrace{\gamma_2}_{\alpha\beta_{FU}} (\pi_{i,t}^T - \pi_{i,t}^0) + \varepsilon_{i,t+1} \quad (7)$$

In other words, we can use the same estimation procedure with $\pi_{i,t+1}^{follow-up}$ instead of $\pi_{i,t+1}$ as the dependent variable. Intuitively, if in the original survey the information provided by the experimenter truly affected the individual's posterior belief about past inflation, then (after properly accounting for the rate of information renewal) we should see that such effect should have persisted in beliefs elicited at future points in time. Since this new estimation strategy should remove spurious learning (at least to some degree), the ratio between the α coefficient based on $\pi_{i,t+1}^{follow-up}$ and the α coefficient based on $\pi_{i,t+1}$ can provide an estimate of the share of learning that is genuine rather than spurious.

We can provide an intuitive explanation of what our estimate is capturing. Among individuals who did not receive any information from us, suppose that we observe that each extra percentage point in perceived inflation today translates, on average, to about 0.5 additional percentage points of inflation expectations two months from now. If an informational treatment truly convinced a subject today that inflation expectations will be 1 percentage point higher, we should observe an increase in inflation expectations of 0.5 percentage points as measured two months later. If, though, the information induced only a 0.25 increase in inflation expectations two months later, then we would conclude that half of the learning was genuine. If the information did not induce any changes in inflation expectations two months after the treatment, then all learning would be deemed spurious.

The second strategy is based on individuals' perceptions and expectations regarding other economic indicators closely related to inflation. In our experiments, we collected information on perceptions about the expected nominal interest rate over the next 12 months, which – just like inflation expectations – was elicited after the experimental information provision. Let $i_{i,t+1}$ denote the expectation about the nominal annual interest rate. The new forward-looking equation is $i_{i,t+1} = \mu_I + \beta_I \pi_{i,t}$, where β_I is the degree of pass-through from inflation perceptions to interest rate expectations. If we combine the new forward-looking equation with the learning equation (5), we obtain:

$$i_{i,t+1} = \gamma_0 + \underbrace{\gamma_1}_{\beta_I} \pi_{i,t}^0 + \underbrace{\gamma_2}_{\alpha\beta_I} (\pi_{i,t}^T - \pi_{i,t}^0) + \varepsilon_{i,t+1} \quad (8)$$

Again, this corresponds to using $i_{i,t+1}$ instead of $\pi_{i,t+1}$ as dependent variable in our learning regression. By comparing the estimated α coefficients in the two specifications, we have a second way of quantifying genuine vs. spurious learning. The intuition for this test is very similar to that of the first test. Assume that among individuals in the control group, respondents who report expecting a 1 percentage point increase in inflation also report a future nominal interest rate that is 0.3 percentage points higher. If an informational treatment truly convinces a subject that future

inflation will be 1 percentage point higher, it should also convince her that the future nominal interest rate will be 0.3 percentage points higher. If, though, the information induced only a spurious effect on inflation expectations, it would have no impact on interest rate expectations (or any other nominal variables intrinsically related to inflation).

3 Results from Online Experiments in the United States and Argentina

3.1 Subject Pool and Descriptive Statistics

In the United States, we recruited subjects from Amazon’s Mechanical Turk (AMT) online marketplace during the month of September 2013. We followed several guidelines that describe the best practices for recruiting individuals for online surveys and experiments using AMT in order to ensure high quality responses (see, for instance, Crump et al., 2013). In Argentina, a first sample was collected through an online survey of College graduates. The second, larger sample is based on an established public opinion research firm that carries out a quarterly online survey of adults in Argentina. See Appendix B.1 for further details about the samples.

According to the Consumer Price Index (CPI) reported by the Bureau of Labor Statistics (BLS), the annual inflation in the United States for the five years prior to our study (2008-2012) was, on average, 1.8%. In the online survey, the mean for inflation perceptions is 5.07% and (in the control group) the mean for inflation expectations is 5.08%. In Argentina the average rate for 2008-2012 was also stable but around 22.5%. In the larger sample, the mean inflation perception is 27.8% and (in the control group) the mean inflation expectation is 28.4%.

Our U.S. sample is younger and more educated than the U.S. average, while our Argentine sample is more educated than the country average (Appendix B.1 provides a comparison of characteristics with population averages). In any case, as shown in Appendix B.3, the results are similar if we re-weight the observations to make them representative on observables. In turn, Figure 3.a compares the distribution of inflation expectations in our U.S. Online Experiment (for the control group) to the University of Michigan’s Survey of Consumers. Besides originating in different samples, there are several methodological differences between the two survey questions capturing expected inflation. Despite all of these differences, the distribution of inflation expectations in the two samples are very similar. For example, the median expectation is just 1 percentage point higher in our online sample (4%) than in the University of Michigan survey (3%), and the interquartile range is just 1 percentage point wider in our sample (2%-6% compared to 2%-5%).

Figure 3.b provides a similar comparison for the Argentine data. In Argentina there is no nationally representative survey equivalent to the Michigan’s Survey of Consumers or the Federal Reserve Bank of New York’s Survey of Consumer Expectations. Instead, Figure 3.b compares our Argentine sample to the Survey of Consumer Expectations conducted by the University Torcuato

Di Tella. This survey is less comparable to our online survey, among other reasons, because the language of the inflation expectation question is different. In spite of these differences, the distribution of inflation expectation is roughly comparable across these two surveys.

In the United States, the final sample includes 3,945 individuals, with 783 in the *Control* group, 807 in the *Statistics* (1.5%) treatment, 763 in the *Products* treatment (10 tables with average price changes from -2 to 7% in 1 percentage point increments within this treatment), 804 in the *Products+Statistics* (1.5%) combined treatment (same 10 tables as above), and 788 in the *Hypothetical* treatment. In Argentina, the first sample yielded a total of 691 observations, 182 of which were assigned to the control, 161 to *Statistics* (24%)¹⁵ and 348 to the *Products* arm (with average price changes of 19%, 24% and 29%). The second sample yielded 3,653 subjects, with 567 subjects assigned the control group and the rest to the *Products* arm (with average price changes ranging from 16% to 34%, in 1 percentage point increments).

3.2 Rational Inattention Test

In this section we discuss the rational inattention test, which relies on the comparison of learning rates between the United States and Argentina.

While in the U.S. the annual inflation rate in the five years before our study (2008-2012) was stable and, on average, 1.8%, in Argentina the average rate for the same time period was also stable but around 22.5%. As a result, the cost of ignoring inflation in Argentina was substantially higher. For example, individuals must rely on good information on inflation prospects in drawing up contracts because it is illegal to index such contracts (labor, real estate, etc.), or rely on more stable foreign currencies.¹⁶ Opinion polls in Argentina at the time of the survey systematically indicated inflation as one of the population's primary concerns.¹⁷ Inflation statistics were mentioned on offline and online news outlets on a regular basis, frequently making the front page of newspapers. According to the rational inattention model (Sims, 2005; Veldkamp, 2011), individuals in Argentina should be more informed and consequently have stronger prior beliefs about past inflation than their U.S. counterparts.

The reduced-form evidence from the online experiment is summarized in Figure 4 (see Appendix B for more detailed outputs by different treatment arms). All the panels in this figure present the distribution of inflation expectations for two treatment arms, where one of them is always the control group (the histograms accumulate the observations below -5% and above 15% in the

¹⁵The value provided in the *Statistics* treatment arm (and reported in that treatment arm) represents the average inflation estimates of private consultancies, research centers, and provincial public statistical agencies, as compiled and computed by opposition parties in the Argentine Congress since the intervention of the national statistical agency in Argentina in 2012 (Cavallo, 2013). These are the statistics that individuals used on a regular basis (for more details, see Cavallo, Cruces and Perez-Truglia, 2016).

¹⁶See Cavallo, Cruces and Perez-Truglia (2016) for more details on the Argentine macroeconomic and institutional context at the time of our experiments.

¹⁷For our opinion poll sample, 40.7% of those in our control group selected inflation as one of the three main concerns for the country.

extreme bars). Each panel in Figure 4 reports the results from an Epps–Singleton (ES) two-sample test using the empirical characteristic function, a version of the Kolmogorov–Smirnov test of equality of distributions valid for discrete data (Goerg and Kaiser, 2009). All the pairwise differences are statistically significant at the 1% level, indicating that our experimental subjects significantly reacted to the inflation provided by us.

We start with the reduced-form results for the United States. Figure 4.a.i presents the results for the *Statistics* (1.5%) treatment, which consisted of providing the respondent solely with a table of statistics about past inflation. According to the Bayesian learning model, providing a signal about inflation should shift the distribution of inflation expectations (relative to the control group) towards the value of the signal, and to produce a more concentrated distribution of expectations. Consistent with this prediction, this signal shifts the distribution of expectations towards 1.5% and makes the distribution of expectations less dispersed.

Figure 4.a.ii shows the effects of the *Products* (2%-3%) treatments (Figure B.2 shows the results for the rest of the *Products* groups). Consistent with the learning model, the signal that supermarket products increased 2-3% shifted inflation expectations towards this range, and reduced the dispersion of expectations.

We can turn to the reduced-form results for Argentina. Figures 4.b.i and 4.b.ii present the results equivalent to those in Figures 4.a.i and 4.a.ii for the United States. The results in Argentina are qualitatively identical to those from the U.S. experiment: providing a signal about inflation shifted the distribution of inflation expectations towards the value of the signal, and led to a more concentrated distribution of expectations.

Even though the effects of *Statistics* and *Products* information are qualitatively identical between the United States and Argentina, the rational inattention test relies on the quantitative comparison between the two countries. To facilitate such comparison, Table 1 presents the estimates of the learning rates, based on the Bayesian learning model introduced in section 2.3, with results from the U.S. experiment in panel (a) and results from the Argentine experiment in panel (b). The table reports the values of α and β from equation (6). As discussed above, β can be interpreted as the degree of pass-through between perceptions of past inflation and expectations of future inflation, and α as the weight placed by the respondents on the information provided in the experiment, with $(1 - \alpha)$ being the weight placed on respondents' prior belief about past inflation.

Column (1) from Table 1.a reports that, in the United States, the learning rate (α) in the *Statistics* (1.5%) treatment was 0.838 (p-value<0.01), whereas the weight given to its equivalent in the *Products* treatment was 0.689 (p-value<0.01). In other words, U.S. subjects assigned a much greater weight to the information provided by the experiment than to their own prior belief. In turn, column (1) from Table 1.b reports that, in the first Argentine sample, the estimated α is 0.432 (p-value<0.01) for the *Statistics* treatment and 0.458 (p-value<0.01) for the *Products* treatment (column (2) shows a similar learning rate of 0.494 for the *Products* treatment in the

second Argentine sample).¹⁸ Consistent with the prediction of the rational inattention model, U.S. subjects revealed to be less informed about inflation, insofar they assigned between 95% (i.e., $\frac{0.838-0.432}{0.432}$) and 51% (i.e., $\frac{0.689-0.458}{0.458}$) more weight to the information about statistics and supermarket products relative to the Argentine subjects.¹⁹

A potential confounding factor is the general loss of credibility of price data in Argentina after the manipulation of official statistics in the period 2007-2015 (Cavallo 2013; Cavallo, Cruces and Perez-Truglia, 2016). While we used an actual unofficial private sector indicator that was widely reported in the media and that became the *de facto* consensus figure for inflation during this period,²⁰ economic agents in Argentina might have become wary of all economic indicators. However, even though this could explain why Argentines react less to inflation statistics than Americans, this explanation cannot explain why Argentines react less to supermarket prices than Americans. Another confounding factor could be that the difference in learning rates between countries could be explained by differences in characteristics of the subject pools. However, this explanation seems unlikely for at least two reasons. First, the observable characteristics are fairly similar across the two countries such as gender, age, and education composition.²¹ Second, as shown in section 2.4 below, the heterogeneity in learning rates by gender, education and other characteristics are an order of magnitude smaller than the difference in learning rates across the two countries. A last confounding factor could be that differences in learning rates are not due to differences in inflation levels but rather due to differences in the volatility of inflation across the two countries. This explanation seems unlikely for at least two reasons. First, inflation levels were relatively stable in both Argentina and the United States in the five years prior to our study. Second, even if Argentine inflation was deemed as more unstable, that would lead to an under-estimation of our results, insofar individuals should react more to new information in more a volatile context.

3.3 Irrational Learning Test

In this section, we test whether individuals use less accurate information on inflation even when the more accurate information is readily available.

Figure 5.c and 5.d present histograms corresponding to the effects of the treatments *Products (0%)* and *Products (1%)* (panel c) and *Products (2%)* and *Products (3%)* on inflation expectations in the U.S. online experiment sample. These figures illustrate how individuals react to different signals

¹⁸The similar results for our college graduates sample I – all of whom had at least some basic training in economics and most of whom were professional economists or accountants – and our public opinion poll sample II suggest that economic literacy does not drive our findings (see also Burke and Manz, 2011).

¹⁹This evidence is consistent with related evidence from Coibon et al. (2015) on firms’ lack of incentives to collect and process information on macroeconomic conditions (i.e., rational inattention) and its impact on firms’s inflation perceptions.

²⁰This unofficial index was computed by members of the opposition in the Argentine Congress, and was constructed as an average of private sector indicators.

²¹See Table B.1 for more details.

about supermarket prices. Relative to the control group, the treatment groups that received signals that supermarket prices had increased 0% and 1% are more likely to state inflation expectations in this range: the proportion stating 0-1% more than doubles compared to that of the control group (Figure 5.c). The reaction is similar to the signals of prices increases of 2% and 3%: this information increases the likelihood that individuals report inflation expectations close to the 2-3% range (Figure 5.d).

Figure 6.a.i provides further evidence about the effects of the *Products* treatment arm on inflation expectations in the United States. Relative to the control group, each bar represents the effect of each of the ten sub-treatments (with average annual price changes in the tables ranging from -2 to 10% on the horizontal axis) on average inflation expectations. This figure indicates that each percentage point increase in the average price change reported on the table of products increased the average inflation expectations for about 0.5 percentage points. Figure 6.b.i shows that, as in Figure 6.a.i for the United States, average inflation expectations in Argentina responded significantly to the average price changes shown in the table of supermarket products.

However, it is possible that individuals pay attention to those signals only because they do not observe their ideal signals, such as inflation statistics. In the treatment arm *Statistics (1.5%)+Products*, experimental subjects were provided with the table of statistics for past annual inflation averaging 1.5% and, immediately afterward, they were presented with one of the *Products* tables with the price changes of supermarket products. Relative to the average price change from a random set of six products, inflation statistics are extremely more precisely estimated. As a result, when confronted to these two pieces of information, we would expect an economic professional to put all weight on the statistics information and ignore the supermarket prices information. In other words, the null hypothesis of rational learning predicts that the *Statistics (1.5%)+Products (0%)* and *Statistics (1.5%)+Products (3%)* treatments should have the same effects on expectations. Figure 5.e and 5.f show the distribution of expectations under these two treatment arms (see Figures B.2 and B.3 for the rest of treatment arms). We can confidently reject the null hypothesis of rational learning: even though these individuals have the latest inflation statistics readily available to them, they assign a significant weight on much less reliable information on supermarket prices.

Column (2) of Table 1.a shows the estimates of learning rates for the treatment arm *Statistics (1.5%)+Products*. When both statistics and supermarket prices were shown, the α coefficient for the supermarket prices is 0.449 (p-value<0.01), even higher than the α of 0.283 (p-value<0.01) for the statistics (the difference is statistically significant at the 1% level). These results suggest that, whenever the two signals disagree with each other, individuals are more willing to incorporate signals closer to their everyday experience, such as a list of price changes for specific products, than signals derived from statistics.

The behavioral framework for thinking about subjective probabilities provides plausible explanations for this finding. A first interpretation has to do with the under-appreciation of the effects of sample sizes on the likelihood of different proportions of a sample (Kahneman and Tversky, 1972),

also known as non-belief in the law of large numbers. For instance, Benjamin et al. (2013) show that averages estimated with samples sizes of 10 and 1,000 are perceived as equally precisely estimated. This could explain why individuals put as much weight to the supermarket prices (based on less than 10 price changes) than to inflation statistics (based on thousands of price changes). A second interpretation relies on the concept of availability heuristic, according to which individuals give more weight to information that is easier to recall. For instance, individuals incorrectly report that there are more words that start with the letter “k” than words with “k” as the third letter, presumably because it takes a more concentrated effort to think of any words in which “k” is the third letter (Tversky and Kahneman, 1973). According to the availability heuristic, prices of products are more familiar and easier to recall than inflation statistics, therefore explaining why individuals put so much weight on this information. Last, an alternative interpretation may be that, even in a developed country like the United States, some individuals distrust official statistics (see Cavallo, Perez-Truglia and Cruces, 2016), or they may fail to comprehend how representative they are.

This evidence is suggestive that individuals are using their own memories of price changes in forming inflation expectations. For example, Bruine de Bruin et al. (2011) present survey evidence that, when asked about inflation, a majority of individuals report to try to recall prices of specific products. At the end of our surveys we also asked individuals about the information they tried to recall. A 64.4% of subjects from the control group of our U.S. experiment sample reported trying to recall the prices of specific products, twice as much as those trying to recall inflation statistics. In Argentina, even though accurate inflation statistics are widely covered by the media, still 74.9% of respondents reported to try to recall prices of specific products when asked about past inflation.

We would like to test more directly the hypothesis that individuals use their price memories in forming inflation expectations. Furthermore, we would like to assess how misleading the use of price memories may be, which depends on the accuracy of these memories. These additional hypothesis require more detailed information about the consumer experience of the individuals, which is difficult to collect in an online survey. To explore these hypotheses further, we conducted an “offline” survey experiment with supermarket customers, which is presented and discussed in section 4 below.

3.4 Spurious Learning Test

In this section we measure the extent to which the learning induced by our experimental setup is spurious. To illustrate this concern, we included a treatment arm in which, with the excuse of a numerical literacy question, respondents were provided information on the current and past prices of a fictitious product, with an implied price change of about 10%. Under the hypothesis of spurious learning, individuals being asked to eyeball this 10% price change would be more likely to anchor their inflation expectations around 10%. Figure 5.b compares the distribution of

inflation expectations between the control group and the *Hypothetical* treatment group in the U.S. sample. Consistent with the hypothesis of spurious learning, the fictitious signal of a 10% price change increased the density around the 10% range of inflation expectations (the ES test indicates a statistically significant difference in the distribution of inflation expectations between these two groups). For instance, respondents in the *Hypothetical* treatment arm were 4.8 percentage points more likely than those in the control group to state that their inflation expectations for the following year were exactly 10% (p-value<0.01).

Column (1) of Table 1.a reports the implied learning rate from the *Hypothetical* treatment. The corresponding α is 0.232, and statistically significant at the 1% level. Though significant, this rate is economically less significant when compared to the learning rates for the other informational treatments. The effect of this treatment may be attributable to unconscious numerical anchoring. Alternatively, this evidence may reveal that some individuals are so uninformed about inflation that they are even willing to use inflation figures from a hypothetical exercise as a benchmark. In any case, the evidence suggests the presence of some degree of spurious learning.

It is important to note that we are comparing individuals who were shown the financial literacy questions (regardless of what they answered) versus individuals who were not asked that question. Since we randomized who gets to see the financial literacy question, this variation is exogenous. Even though the correct answer to the financial literacy question was 10%, some individuals did not respond correctly (21%, to be precise). However, note that this can only lead to an under-estimation of spurious learning: i.e., if individuals responded some value X% different from 10%, we should observe anchoring of inflation expectations around X% rather than around 10%.²²

The first methodology to weed-out the spurious learning consist of estimating the learning model using the inflation expectations in the follow-up survey. We used data on a subsample of 1,073 subjects who were re-interviewed two months after the original online experiment. This subsample was asked again about their inflation expectations, but they were not subjected to any type of new informational treatment or reminded of previous informational treatments. Column (3) of Table 1.a presents the results of the basic regression with inflation expectations in the original survey as the dependent variable, but only for the subsample of those who later participated in our follow-up survey in the United States. The β and α coefficients are very similar to those presented in column (1) for the full sample.

Column (4) of Table 1.a presents the regression for the same follow-up subsample, but in this case with inflation expectations as reported in the follow-up survey as the dependent variable. The β coefficient is now substantially lower, falling from 0.814 in column (3) to 0.438 in column (4). This result is consistent with our learning model, because the β coefficient is the product of two terms – the pass-through from perceptions to expectations, and the pass-through from

²²We could see whether, relative to the control group, individuals responding X% to the financial literacy question tend to anchor their inflation expectations around X%, for each of the possible values of $X=\{1\%,5\%,10\%,100\%\}$. However, we refrain from making this comparison because the responses to the financial literacy questions were not randomized and thus are most likely endogenous.

inflation perceptions in the original survey to inflation perceptions in the follow-up survey. The α coefficients of 0.360 for the *Statistics* treatment and of 0.336 for the *Products* treatment are both statistically significant (at the 1% and 5% levels respectively). However, they are about half as large as the coefficients of 0.799 and 0.697 from column (3). These estimates suggest that between 45% (i.e., $\frac{0.360}{0.799}$) and 48% (i.e., $\frac{0.336}{0.697}$) of the effect of the information provided can be attributed to genuine, rather than spurious, learning. Notably, the α coefficient for the *Hypothetical* treatment in the follow-up results in column (4) is close to zero and statistically insignificant, in contrast to the small but positive and significant effect in column (3). This evidence suggests that our methodology may have successfully weeded out the spurious learning.

Columns (3) and (4) of Table 1.b reproduce the above results but for the Argentine sample instead of the U.S. sample. A subsample of 1,320 of the subjects from the Argentine opinion poll sample were re-interviewed four months after the original survey.²³ As in the case of the United States, the β coefficient is lower for the follow-up regression, falling from 0.963 in column (3) to 0.754 in column (4).²⁴ The α coefficient of 0.208 is statistically significant, but only half as large as the coefficient in column (3), indicating that about 45.6% (i.e., $\frac{0.208}{0.456}$) of the effect of the information provided can be attributed to genuine, rather than spurious, learning. This reinforces the findings of the U.S. online experiment, which showed a proportion of genuine learning of about 45% in the context of a similar follow-up survey.

The second methodology for weeding-out spurious learning consists of measuring learning rates based on the indirect effect of the information provided on the expected nominal interest rate. In the U.S. sample, we report results from this exercise in column (5) of Table 1, where the dependent variable is an individual's expectation for the nominal interest rates for the following twelve months. The β coefficient indicates that for each additional percentage point in expected inflation, on average, subjects believed that the nominal interest rate would be about 0.3 percentage points higher. This is consistent with Behrend (1977), who presents evidence that individuals have a significant amount of useful understanding of the link between inflation and other economic outcomes such as the nominal exchange rate. The estimated α are 0.314 for the *Statistics* treatment (borderline insignificant at the 10% level) and 0.499 for the *Products* treatment (significant at the 1% level). When these parameters are compared to those presented in column (1), they suggest that between 37.5% (i.e., $\frac{0.314}{0.838}$) and 72.5% (i.e., $\frac{0.499}{0.689}$) of the learning is genuine. The average between these two figures, 55%, is close to the corresponding share of genuine learning inferred from the follow-up survey (46.6%). That is, both of these methodologies provide similar estimates of the degree of spurious learning. The results in column (5) indicate that the *Hypothetical* treatment arm did not have a significant effect on individuals' expected interest rates. This, again, can be

²³There was no significant difference in the probability of participating in the follow-up sample between the treatment and the control groups.

²⁴Comparing the β coefficient between the follow-up and the original samples provides a measure of the persistence of beliefs. This evidence suggests that expectations are more persistent in Argentina than in the United States: i.e., the β for the follow-up survey is only 53.8% of the same coefficient in the original survey in the U.S. (0.438 and 0.814, Table 1) while in Argentina the corresponding proportion is 78.2% (0.754 and 0.963, Table 1).

interpreted as evidence that this methodology effectively weeds out spurious learning.

Column (5) of Table 1.b reproduces the specification from column (5) of Table 1.a, but for Argentina instead of the United States. In this case, the α coefficient of 0.468 is close to the baseline value from column (2) of Table 1.b (0.494). This estimate suggests that 95% (i.e., $\frac{0.468}{0.494}$) of learning is genuine. However, this estimate is less precisely estimated, and as a result we cannot reject the null hypothesis that only 50% of the learning is genuine.

The results for the nominal interest rate also support our findings in a more general way. Our survey questions always refer to inflation expectations in the sense of changes in the average general price level. However, it may be argued that individuals may mistakenly respond as if we asked about their own idiosyncratic inflation – i.e., the price change of their own consumption basket.²⁵ The results described in this paragraph show that this cannot be the case: changes in inflation expectations affect expectations about nominal variables like the interest rate (and the exchange rate in the Argentine case, discussed in Appendix B.3), which should not be affected if the individual was only thinking about her own idiosyncratic experience.

In sum, while there is a significant level of spurious learning, about half of it can be still be considered genuine. More importantly, once we account for spurious learning, the main results still hold: e.g., it is still true that the learning rate in Argentina is substantially lower than that in the United States.

Additionally, to help establish the validity of the estimation of the learning rates, we can test some auxiliary predictions of the Bayesian model. One prediction yielded by this model is that providing relevant information will increase the accuracy of the posterior belief. We can test this hypothesis with our data using the respondents' self-reported confidence in their own inflation expectations in a scale from "very sure" to "very unsure." Higher values of this confidence variable indicate higher confidence, and this variable was standardized to have a standard deviation of 1. In the United States, relative to the control group, the *Statistics (1.5%)* treatment increased confidence by 0.324 (p-value<0.01) the *Products* treatments increased confidence by 0.226 (p-value<0.01), the combined *Products+Statistics (1.5%)* increases confidence by 0.368, and the *Hypothetical* treatment only increases confidence by a non-significant 0.032 (p-value of 0.54). This pattern of results indicates a larger increase in confidence the more factual information we provide to our subjects, and no change when we provide only the (non-factual) *Hypothetical* treatment. Similarly, in Argentina's sample I, relative to the control group, the *Statistics (24%)* treatment increased standardized confidence by 0.197 (p-value of 0.07) and the combined *Products* treatments increased confidence by 0.152 (p-value of 0.09), whereas for sample II the combined *Products* treatments increased confidence by 0.360 (p-value<0.01) compared to that of the control group.

The Bayesian model also predicts that all signals from the same source should be equally

²⁵Indeed, Armantier et al. (2016) find that information about food prices causes consumers to update expectations more for their own-basket inflation rate, but less for their rate of inflation.

informative to respondents, regardless of their value. In terms of the sub-treatments, Figure 6.a.ii compares the impact of each level of price changes for the *Products* treatment arm on the confidence on the posterior belief for the United States. The evidence is consistent with this prediction, although with a slight twist: we cannot reject the hypothesis that all the positive signals (1%-7%) have the same effect on confidence (p-value if 0.77) and we cannot reject the hypothesis that all non-positive signals (-2%, -1% and 0%) have the same effect on confidence (p-value of 0.83), but we can reject the hypothesis that positive and negative signals have the same effect on confidence (p-value<0.01), which is suggestive evidence that individuals might be less prone to incorporate information about price decreases than about price increases. Figure 6.b.ii reproduces the analysis of Figure 6.a.ii but for the Argentine experiment (sample II) instead of the U.S. The results suggest that, consistent with the Bayesian model, all these different signals led to the same gain in confidence about the posterior belief (we cannot reject the null hypothesis that all signals have the same effect on posterior confidence, with a p-value of 0.73).

Another prediction of the Bayesian model is that the strength of the reaction to new information should depend on the strength of the prior. We can test this hypothesis using the subjective measure of confidence in the prior belief. Figure 7.a presents the value of the α coefficient corresponding to the reaction to the *Products* treatment, split by individuals with “High Confidence” and with “Low Confidence” in their stated inflation perceptions (United States corresponds to panel a and Argentina to panel b). The “High Confidence” group corresponds to those who answered “Sure” or “Very Sure” when asked their confidence on their inflation perceptions, and the “Low Confidence” group corresponds to those who replied “Somewhat Sure,” “Unsure” or “Very Unsure.” As predicted, in both countries the learning rates as captured by the α coefficient are higher (i.e., individuals react more to the new information) for those with lower levels of confidence in their prior beliefs about inflation. In the U.S. sample, the learning rate is 0.75 for “Low Confidence” and 0.58 for “High Confidence,” with a statistically significant difference (p-value 0.03). In Argentina, the learning rate is 0.62 for “Low Confidence” and 0.41 for “High Confidence,” also with a statistically significant difference between the two coefficients (p-value<0.01).

As a final robustness check, we can measure the heterogeneity in learning rates with respect to individual demographic characteristics, and compare those results with the findings reported in other papers about the heterogeneity of biases in expected inflation. Figure 7 presents this heterogeneity, by computing the learning rates for different subgroups of the population (panel (a) for United States and panel (b) for Argentina). Overall, the direction of the effects is identical across the two countries, with learning rates being lower for males, educated and older individuals, although the differences tend to be more pronounced in Argentina than in the United States in terms of economic and statistical significance.²⁶ This heterogeneity in learning rates is consistent

²⁶Females have larger learning rates than males: 0.719 for females compared to 0.702 for males in the United States (p-value of the difference of 0.81), and 0.519 compared to 0.432 in Argentina (p-value=0.07). Individuals with a college degree sometimes have lower learning rates than individuals without such degree: 0.715 for college graduates compared to 0.716 for non-graduates for the United States (p-value=0.99), and 0.4 compared to 0.535

with the evidence discussed in existing literature, for example, showing that females, less educated and younger individuals from the United States tend to be less informed about inflation and have more biased inflation expectations (Bruine de Bruin et al., 2011; Malmendier and Nagel, 2016; Armantier et al., 2016).

4 The Supermarket Experiment

4.1 Remaining Hypotheses to be Tested

The evidence from the online experiments indicates that, even when inflation statistics are readily available, individuals pay attention to prices of specific products in forming their inflation perceptions and expectations. This is suggestive evidence that individuals use their price memories to form inflation expectations. In this section, we discuss an “offline” survey with supermarket customers designed to strengthen our understanding of this issue.

A first remaining hypothesis is that, even though suggestive, these findings do not constitute conclusive evidence that individuals use price memories in forming inflation expectations. For example, subjects may have reacted to the price information insofar as they perceived it to be accurate, but they would not trust their own price memories for the same products. A second remaining hypothesis is whether the use of price memories leads to significantly biased inflation expectations. The more inaccurate the price memories, the more misleading its use will be.

Addressing these two remaining hypothesis would require data that is difficult to collect in an online survey, on products purchased by subjects, the actual historical prices of those products, the individual’s memories of those historical prices, and the individual’s inflation perceptions and expectations. Moreover, we would need a source of exogenous variation in the price memories of subjects. We designed and conducted a unique consumer intercept survey at the main exit of some supermarkets in Buenos Aires to meet all of these requirements.

4.2 Subject Pool and Experimental Design

The consumer intercept survey was carried out in four branches of one of the largest supermarket chains in the city of Buenos Aires. The subject pool consisted of supermarket customers who had just finished shopping and in the way out of the supermarket they were invited to participate in a short survey for an academic study. A total of 1,200 subjects were interviewed for about three to five minutes, yielding 1,140 observations with complete information about relevant outcomes. Using hand-held scanners, the interviewers scanned respondents’ receipt from the supermarket purchase, which contained product identifiers that could be matched to our database of scrapped

for Argentina (p-value<0.01). Older subjects (above 35 years old) have lower learning rates than their younger counterparts: 0.586 for the older compared to 0.750 for the younger for the United States (p-value=0.03), and 0.469 compared to 0.535 for Argentina (p-value=0.16).

online data of supermarket prices for the supermarket chain where the study was conducted. After providing purchase receipts for scanning, respondents were asked twelve questions.

Some of these questions were about the current and past prices of supermarket products that the individuals just bought. Specifically, respondents were asked to recall the current price, and the price twelve months earlier, of two specific products they had just purchased, chosen at random by the interviewer from the receipt. The interviewers selected two additional products from the receipt, read each of their prices out loud, and asked the respondents what they thought the prices of these two products had been twelve months earlier. These questions had a double purpose: measuring the accuracy of price memories, and also measuring the effect of price memories on inflation expectations. For the second goal, we followed a design similar in spirit to the online experiments: first, we asked about perceptions of the inflation rate over the past year; we then asked individuals about the prices of a random set of four products; last, we asked about expectations for the inflation rate for the following twelve months. This information-retrieval exercise intended to generate random variation in the salience of the individual’s own price memories about four specific products randomly chosen from the receipt. In the online experiments, we provided a table with specific, pre-selected product prices and price changes. This supermarket experiment also consisted of a list of four products at random, although this time they corresponded to products that the individual had just purchased and thus were relevant for them. Most important, instead of providing the historical prices for these four products, we asked respondents to “fill in the table” by using their own price memories. By chance, some of the products we made salient through this procedure corresponded to products with higher or lower actual price changes, and/or with higher or lower remembered price changes. This design allows us to test whether making salient these products had any effect on subsequent individuals’ inflation expectations. Unlike in the other informational treatments in our study, subjects were not learning new information – we were only making salient some information that they already had.²⁷

4.3 Accuracy of Memories about Current and Past Prices

The goal of this section is to compare the memories about current and past prices to the actual prices. Panel (a) in Figure 8 presents a scatterplot of prices for the products the respondents had just purchased, with the prices the respondents reported paying for (without looking at the receipt) on the vertical axes and the prices they actually paid for them on the horizontal axis. The relationship between the two variables seems to be linear, with most observations clustered around the 45 degree line, indicating that individuals’ memories of the prices of the products they had just purchased were fairly accurate. Panel (b) in Figure 8 presents the results of a more taxing exercise for respondents’ memory: we present a scatterplot of respondents’ reported recollections

²⁷As a benchmark, we also included a second informational treatment that was identical to the one used in the online experiments, consisting of showing the actual price histories for six randomly selected products (results reported in Appendix C).

of the prices of the same goods one year earlier (vertical axis) and of the actual prices one year earlier (horizontal axis), obtained from our database of scrapped prices for the same supermarket chain. The main pattern that emerges indicates that individuals' recalled prices for one year earlier are systematically lower than the actual prices of those products at that time as indicated in our database.²⁸ Moreover, the R^2 of the predictions provided by the individuals about current prices is 0.81 – while not a perfect fit, the relationship is very tight. However, the R^2 drops to just 0.65 when individuals are asked about past prices. A significant part of that drop in predictive power is likely due to the fact that individuals systematically under-estimate past prices. Since individuals have relatively unbiased memories of current prices but tend to underestimate past prices, they often overestimate price changes.

This evidence is consistent with literature in psychology documenting large biases in remembered prices (e.g., Bates and Gabor, 1986; Kemp, 1987; Monroe and Lee, 1999). Unlike the existing literature, however, our evidence shows that these biases are substantial even in a context of high inflation and in a natural environment for consumers such as the products and brands that the consumers actually buy at the supermarket.

Even though price changes are overestimated on average, there may be a correlation between remembered price changes and actual price changes. For instance, individuals might be mistakenly reporting prices for twenty months earlier rather than for twelve months earlier. Panel (d) in Figure 8 presents a comparison of the remembered price changes and the actual price changes observed in our database of supermarket prices. There is a weak correlation between the two: for each percentage point increase in the actual price change, the remembered price change increases by only 0.13 percentage points.²⁹

In Appendix C we show that individuals are similarly inaccurate when asked about the changes in bundles of products instead of specific products. However, individuals may follow the evolution of prices for a different set of products (e.g., a handful of “favorite” goods), and their memories for these products may be more accurate. With this caveat in mind, we show in Appendix A that even with perfectly accurate recollections, if the number of products an individual keeps track of is small, that can generate substantial excess dispersion in inflation expectations, enough to explain the dispersion of inflation expectations observed in the data.

Finally, it should be noted that noisy memories about past prices may also arise mechanically in the absence of behavioral biases. For instance, in a model with information frictions where individuals can only carry a limited amount of information, individuals may use current prices

²⁸This underestimation of past prices may be due in part to the fact that individuals may struggle with the operation of projecting percentage changes into the past. See, for example, the discussion about implicit memory in Monroe and Lee (1999).

²⁹As a benchmark, panel (c) in Figure 8 presents respondents' perceptions of aggregate inflation over the previous twelve months and the implicit average percentage price change of the products for which we requested this information. As expected, the correlation is positive and significant: i.e., individuals who believe inflation was higher also believe that, on average, prices of specific products increased more. For each percentage point increase in perceptions of past inflation, the remembered price change increases by about 0.69 percentage points.

and inflation perceptions to “backcast” past prices (Woodford, 2009). However, this mechanism does not seem to be the whole explanation in our data. For instance, a regression of remembered price changes on the inflation perceptions rate yields an R^2 of 0.057, which suggests that the vast majority of the variation in remembered price changes cannot be explained by “backcasting.”

4.4 Evidence on the Use of Actual and Remembered Price Changes on the Formation of Inflation Expectations

Panel (e) in Figure 8 presents evidence on the effect of remembered prices changes on inflation expectations. It presents a comparison of the distribution of inflation expectations when, conditional on the individual’s inflation perceptions, we made salient products that the individual remembered to have higher and lower price changes.³⁰ The results from this exercise indicate that making salient products with higher remembered price changes generates higher inflation expectations – the distributions differ significantly, as suggested by the ES test, and inflation expectations are 32.77% for the low remembered price changes group compared to 37.38% for the high remembered price changes group (the difference is significant at the 1% level). This finding suggests that individuals use memories of their own experience as consumers when forming their inflation expectations.³¹

As we established above, these memories are highly inaccurate, so this may generate substantial biases in expectations. To show this more directly, Panel (f) in Figure 8 presents a comparison of the distribution of inflation expectations between groups of individuals for which we randomly made salient products whose actual price changes (rather than their price changes as remembered by the respondents) were higher, following the same methodology used in panel (e) for remembered prices. The comparison of the two distributions indicate that making salient products with actual higher price changes did not result in higher inflation expectations: while the p-value of the ES test suggests that the two distributions differ significantly, the differences are economically small. For example, inflation expectations are 33.44% for the low actual price changes group compared to 34.10% for the high actual price changes group, and the difference is not statistically significant at standard levels. In other words, it is the remembered price changes and not the actual price changes that mattered for the formation of our subjects’ inflation expectations.

³⁰Specifically, we computed the remembered price change as the average of the price changes of the four randomly selected products that each respondent was asked to state. We then controlled for each individual’s inflation perceptions by subtracting the variation in the average remembered price change that can be explained by inflation perceptions, and divided those residuals in two extreme groups: the top half (i.e., high price changes) and the bottom half (i.e., low price changes).

³¹In this case, unlike the other informational treatments, we did not randomize the recalled price changes directly, but randomized instead the salience of the recalled price changes for a group of products. As a result, estimating the weight assigned to this information (the α coefficient) with our learning regression would not yield the same interpretation in terms of rate of learning as in the information provision treatments in the online experiments. Table C.1 in Appendix C presents regression for the corresponding rate of learning, although these results should be interpreted with this caveat in mind.

All in all, the use of price memories as inputs for the formation of inflation expectations tends to induce large errors in beliefs and may cause the significant dispersion observed in expectations. This evidence is consistent with the fact that, even though their price memories are actually strongly biased, subjects are largely unaware of these biases: when asked how confident they were, only 9.81% of subjects reported to be either “unsure” or “very unsure” about their answers to the questions about prices of specific products.³²

5 Conclusions

We presented evidence from a series of survey experiments in which we randomly assigned respondents to treatments that provided different information related to inflation, such as inflation statistics or price changes for specific products. We used that exogenous variation to estimate the rate of learning from different sources of information. We document two main findings. First, consistent with the rational inattention model, individuals in lower-inflation contexts have significantly weaker priors about the inflation rate. Second, we found that rational inattention is not the only significant source of information frictions: even when information about inflation statistics is made readily available to them, individuals still place significant weight on less accurate sources of information, such as their own memories on prices of supermarket products.

Our findings have a number of implications for macroeconomic models and policy-making. How households form inflation expectations is an important consideration for central banks insofar as, by anchoring expectations, the policies of monetary authorities attempt to influence decisions that households make about consumption and investment. It is, then, important to incorporate realistic informational frictions in models of households expectations and monetary policy (e.g., Coibion and Gorodnichenko, 2015). From a more practical perspective, our findings imply that central banks could have a greater influence on inflation expectations by disseminating information on individual product prices and communicating how objective, accurate and representative inflation statistics are.

Our findings also contribute to the discussion on the potential usefulness of survey data on inflation expectations. Some researchers attribute the biases in household inflation expectations to the inherent limitations of self-reported data (Manski, 2004), which would imply that survey data on household expectations is not useful.³³ Other authors argue that the failure to incorporate public information is a natural outcome of rational inattention (Mankiw et al., 2003). This would imply that survey data on expectations has limited value, because inaccurate expectations merely

³²In comparison, only 9.72% responded to be “unsure” or “very unsure” about their responses to the question about the inflation rate over the past 12 months.

³³Of course, the limitations with subjective reports must explain at least part of the dispersion in expectations. For example, Armantier et al. (2012) show that even though individuals’ inflation expectations are correlated to their actual behavior in a financially incentivized investment experiment where future inflation affects payoffs, there are substantial discrepancies correlated to numeric and financial literacy.

reveal that the respondents do not care about inflation. Our evidence suggests that individuals report biased beliefs on inflation partly because they use private sources of information (e.g., price memories), even when inflation statistics are readily available. This implies that some of the observed heterogeneity in reported inflation expectations reflects actual heterogeneity in deep beliefs rather than measurement error or rational inattention.³⁴

³⁴Consistent with this interpretation, our survey data reveals that even individuals with biased inflation expectations report significant confidence about their stated expectations. For individuals in the control group in the U.S., the average levels of confidence about perceptions of past inflation of 1%, 2%, and 3% (i.e., closest to the average of statistics, 1.5%) are 2.6 for past inflation and 2.69 for inflation expectations (on a scale of 1 to 5). The figures for confidence are 2.95 and 2.85 respectively for those whose stated perceptions of past inflation were -4% or lower or 7% or higher.

References

- [1] Armantier, O., Bruine de Bruin, W., Potter, G., Topa, G., van der Klaauw, W. and Zafar, B. (2013). “Measuring Inflation Expectations,” *Annual Review of Economics*, Vol. 5, pp. 273-301.
- [2] Armantier, O., Bruine de Bruin, W., Topa, G., der Klaauw, V., Wilbert, H., and Zafar, B. (2012). “Inflation expectations and behavior: Do survey respondents act on their beliefs?,” Federal Reserve Bank of New York Staff Report No. 509.
- [3] Armantier, O., Nelson, S., Topa, G., van der Klaauw, W. and Zafar, B. (2016). “The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economics and Statistics*, forthcoming.
- [4] Atkeson, A., and Ohanian, L.E. (2001). “Are Phillips Curves Useful for Forecasting Inflation?,” FRB Minneapolis Quarterly Review (Winter) pp. 2-11.
- [5] Bachmann, R., Berg, T. and Sims, E. (2015). “Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence,” *American Economic Journal: Economic Policy*, Vol. 7 (1), pp. 1-35.
- [6] Badarinza, C. and Buchmann, M. (2009). “Inflation Perceptions and Expectations in the Euro Area: The Role of News,” ECB Working Paper 1088.
- [7] Barr, D.G. and Campbell, J.Y. (1997). “Inflation, real interest rates, and the bond market: A study of UK nominal and index-linked government bond prices,” *Journal of Monetary Economics*, Vol. 39, pp. 361-383.
- [8] Bates, J. M. and Gabor, A. (1986). “Price perception in creeping inflation: Report on an enquiry,” *Journal of Economic Psychology*, Vol. 7, pp. 291–314.
- [9] Behrend, H. (1977). “Research into inflation and conceptions of earnings,” *Journal of Occupational Psychology*, Vol. 50, pp. 169–176.
- [10] Benjamin, D.; Moore, D. and Rabin, M. (2013). “Misconceptions of Chance: Evidence from an Integrated Experiment,” mimeo.
- [11] Bernanke, B. (2007). “Inflation Expectations and Inflation Forecasting,” Speech at the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, Massachusetts, July 10, 2007.
- [12] Blanchflower, D. and MacCoille, C. (2009). “The formation of inflation expectations: an empirical analysis for the UK,” Working Paper.
- [13] Brachinger, H., (2008). “A new index of perceived inflation: Assumptions, method, and application to Germany,” *Journal of Economic Psychology*, vol. 29(4).
- [14] Bruine de Bruin, W., van der Klaauw, W. and Topa, G. (2011). “Expectations of inflation: The biasing effect of thoughts about specific prices,” *Journal of Economic Psychology*, Vol. 32 (5).

- [15] Burke, M. and Manz, M. (2011). “Economic literacy and inflation expectations: evidence from a laboratory experiment,” Public Policy Discussion Paper No. 11-8.
- [16] Carrillo, P.E., and Shahe Emran, M. (2012). “Public information and inflation expectations: Microeconomic evidence from a natural experiment,” *Review of Economics and Statistics*, Vol. 94 (4), pp. 860-877.
- [17] Carroll, C. (2003). “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, 118(1).
- [18] Cavallo, A., Cruces, G. and Perez-Truglia, R. (2016). “Learning from Potentially Biased Statistics,” *Brookings Papers on Economic Activity*, forthcoming.
- [19] Cavallo, A. (2013). “Online and Official Price Indexes: Measuring Argentina’s Inflation,” *Journal of Monetary Economics*. Vol. 60 (1).
- [20] Coibion, O. and Gorodnichenko, Y. (2015). “Is The Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation,” *American Economic Journal: Macroeconomics*, Vol. 7 (1), pp. 197-232.
- [21] Coibion, O. and Gorodnichenko, Y. (2012). “What can survey forecasts tell us about informational rigidities?,” *Journal of Political Economy*, Vol. 120, pp. 116-159.
- [22] Coibion, O., Gorodnichenko, Y. and Kumar, S. (2015). “How Do Firms Form Their Expectations? New Survey Evidence.” NBER Working Paper No. 21092.
- [23] Crump, M.J.C., McDonnell, J.V. and Gureckis, T.M. (2013). “Evaluating Amazon’s Mechanical Turk as a Tool for Experimental Behavioral Research,” *PLoS ONE*, Vol. 8 (3).
- [24] D’Acunto, F.; Hoang, D. and Weber, M. (2016), “Unconventional Fiscal Policy, Inflation Expectations, and Consumption Expenditure,” *Mimeo*.
- [25] Demery, D. and Duck, N. (2007). “The theory of rational expectations and the interpretation of macroeconomic data,” *Journal of Macroeconomics*, Vol. 29 (1), pp. 1-18.
- [26] Drager, L. (2011). “Inflation Perceptions and Expectations in Sweden - Are Media Reports the ‘Missing Link’?,” *KOF Swiss Economic Institute Working Paper No. 273*.
- [27] Eurobarometer (2008). “Europeans’ knowledge of economic indicators,” *Special Eurobarometer 323, Wave 67.2 – TNS Opinion & Social*, European Commission, Brussels.
- [28] Fuster, A., Laibson, David and Mendel, Brock (2010). “Natural Expectations and Macroeconomic Fluctuations,” *Journal of Economic Perspectives*, Vol. 24 (4). pp. 67-84.
- [29] Gennaioli, N. and Shleifer, A. (2010). “What Comes to Mind,” *Quarterly Journal of Economics* 125, no. 4: 1399-1433.

- [30] Goerg, S. and Kaiser, J. (2009). “Nonparametric testing of distributions—the Epps–Singleton two-sample test using the empirical characteristic function,” *Stata Journal*, Vol. 9(3), pp. 454-465.
- [31] Goffman, E. (1963). “Stigma: Notes on the Management of Spoiled Identity,” New Jersey: Prentice-Hall.
- [32] Hellwig, C. (2005). “Heterogeneous Information and the Benefits of Public Information Disclosures,” Working Paper.
- [33] Jonung, L. (1981). “Perceived and expected rates of inflation in Sweden,” *The American Economic Review*, Vol. 71 (5), pp. 961-968.
- [34] Lamla, M.J. and Lein, S.M. (2008). “The Role of Media for Consumers’ Inflation Expectation Formation,” KOF Swiss Economic Institute, Working Paper No. 201.
- [35] Kahneman, D. and Tversky, A. (1972), “Subjective probability: A judgment of representativeness,” *Cognitive Psychology*, Vol. 3 (3), pp. 430–454.
- [36] Kemp, S. (1987). “Estimation of past prices,” *Journal of Economic Psychology*, Vol. 8, pp. 181–189.
- [37] Lucas, R. E., Jr. (1972). “Expectations and the Neutrality of Money,” *Journal of Economic Theory* 4 (2): 103–124.
- [38] Madeira, Carlos and Zafar, Basit (forthcoming). “Heterogeneous Inflation Expectations, Learning, and Market Outcomes,” *Journal of Money, Credit, and Banking*.
- [39] Manski, C. (2004). “Measuring expectations,” *Econometrica* 72 (5), pp. 1329–1376.
- [40] Mankiw, N.G. and Reis, R. (2002). “Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, Vol. 117 (4), pp. 1295-1328.
- [41] Mankiw, N.G., Reis, R. and Wolfers, J. (2003). “Disagreement About Inflation Expectations.” In *NBER Macroeconomics Annual 2003*, ed. by M. Gertler, and K. Rogoff.
- [42] Malmendier, U. and Nagel, S. (2016). “Learning from Inflation Experiences,” *The Quarterly Journal of Economics* (2016) 131 (1): 53-87.
- [43] McGranahan, L., and Paulson, A. (2006). “Constructing the Chicago Fed Income Based Economic Index– Consumer Price Index: Inflation Experiences by Demographic Group: 1983-2005,” Federal Reserve Bank of Chicago Working Paper.
- [44] Monroe, K.B. and Lee, A.Y. (1999). “Remembering versus knowing: Issues in buyers’ processing of price information,” *Journal of the Academy of Marketing Science*, Vol. 27, pp. 207–225.
- [45] Morris, S. and Shin, H.S. (2002). “Social Value of Public Information,” *The American Economic Review*, Vol. 92 (5), pp. 1521-1534.
- [46] Mullainathan, S. (2002), “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, Vol. 117 (3), pp. 735-774.

- [47] Phelps, E.S. (1969). "The New Microeconomics in Inflation and Employment Theory," *American Economic Review: Papers and Proceedings*, 59.
- [48] Ranyard, R., Missier, F.D., Bonini, N., Duxbury, D. and Summers, B. (2008). "Perceptions and expectations of price changes and inflation: A review and conceptual framework," *Journal of Economic Psychology*, Vol. 29(4), pp. 378-400.
- [49] Rosenthal, R. (1966). *Experimenter effects in behavioral research*. New York : Appleton-Century-Crofts.
- [50] Roos, M.W.M. and Schmidt, U. (2012). "The Importance of Time-Series Extrapolation for Macroeconomic Expectations," *German Economic Review*, Vol. 13(2), pp. 196–210.
- [51] Sargent, T.J. (1993). "Bounded Rationality in Macroeconomics," Oxford: Oxford University Press.
- [52] Sims, C. (2005). "Rational inattention: a research agenda," Discussion Paper Series 1, Deutsche Bundesbank.
- [53] Tversky, A. and Kahneman, D. (1973), "Availability: A heuristic for judging frequency and probability," *Cognitive Psychology*, Vol. 5 (2), pp. 207–232.
- [54] Tversky, A., and Kahneman, D. (1974). "Judgement under uncertainty: Heuristics and biases," *Science*, Vol. 185, pp. 1124–1130.
- [55] Veldkamp, L. (2011). "Information Choice in Macroeconomics and Finance," New Jersey: Princeton University Press.
- [56] Woodford, M. (2009). "Information-constrained state-dependent pricing," *Journal of Monetary Economics*, Vol. 56, p. 100-124.
- [57] Zizzo, D. J. (2010). "Experimenter demand effects in economic experiments," *Experimental Economics*, Volume 13, Issue 1, pp 75-98.

Figure 1: Example of *Products* (various levels), *Statistics* (1.5%) and *Hypothetical* (10%) Treatments, U.S. Online Experiment

a. *Products* (-2%)

Product	Price on August 1, 2012	Price on August 1, 2013	Price change in %
Infant Formula (Enfamil Gentlease)	\$18 ⁶⁹	\$18 ⁶⁹	0.0%
Bread (Anzio & Sons Sub Rolls)	\$3 ⁵⁹	\$3 ⁵⁹	0.0%
Pasta Sauce (Barilla Marinara)	\$2 ⁷⁹	\$2 ⁸⁰	0.4%
Cereal (Cheerios Honey Nut)	\$5 ²⁹	\$4 ⁹⁹	-5.7%
Soda (Schweppes Ginger Ale)	\$1 ⁷⁹	\$1 ⁶⁷	-6.7%
Body Wash (Dial Spring Water)	\$6 ⁰⁹	\$6 ⁰⁹	0.0%
Average change:			-2.0%

b. *Products* (2%)

Product	Price on August 1, 2012	Price on August 1, 2013	Price change in %
Infant Formula (Similac with Iron)	\$7 ²⁹	\$7 ⁵⁹	4.1%
Bread (Pepperidge Farm Sliders)	\$3 ⁰⁰	\$2 ⁹⁹	-0.3%
Noodles (No Yolks)	\$2 ⁷⁹	\$2 ⁷⁹	0.0%
Cereal (Natures Path Envirokidz)	\$4 ⁹⁹	\$5 ³⁹	8.0%
Soda (Dr Pepper)	\$1 ⁷⁹	\$1 ⁷⁹	0.0%
Body Wash (Dial Spring Water)	\$6 ⁰⁹	\$6 ⁰⁹	0.0%
Average change:			2.0%

c. *Statistics* (1.5%)

Official Statistic	Average Annual Change in Prices
Consumer Price Index ¹	2.0%
Personal Consumption Expenditures Price Index ²	1.1%
Gross Domestic Product Deflator ³	1.5%
Average of the three statistics:	1.5%

Sources: 1 Bureau of Labor Statistics, 2 and 3: Bureau of Economic Analysis.

d. *Hypothetical* (10%)

Please consider the following prices of a hypothetical product at two different moments.


Price on January 1st 2012:	\$9.99
Price on January 1st 2013:	\$10.99

What is the approximate price change for this product over this period? Please do not use a calculator, pen, or pencil to calculate the exact figure. We want your best guess from eye-balling these prices.

- ☐ About 1%
- ☐ About 5%
- ☐ About 10%
- ☐ About 100%


Notes: The *Statistics* treatment was preceded by the following text: “Before answering, please look at the table below. The table shows indicators used by different government agencies to measure the annual inflation rate - that is, how much prices have changed on average over the last 12 months, from August 1 2012 to August 1 2013.” The *Products* treatments were preceded by the following text: “Before answering, please look at the table below. The table shows the price of each listed product on August 1st, 2012 and on August 1st, 2013 (that is, one year later). These prices were taken from the same branch of a large supermarket chain. The six products that appear in this table were randomly selected from a database containing hundreds of products.” The *Hypothetical* treatment was preceded by the following text: “In this survey we ask you questions about how “prices in general” evolve over time. The following question is meant to assess how comfortable you are with the way these questions are phrased. Please consider the following prices of a hypothetical product at two different moments” and, immediately afterward, included the following question: “What is the approximate price change of this product over this period? Please do not use a calculator, pen or pencil to calculate the exact figure. We want your best guess from eye-balling these prices.” See the questionnaire appendix for more details. Figure B.1 in the Appendix presents examples of the *Products* treatment for the Argentina Online Experiment.

Figure 2: Past Inflation Perceptions vs. Future Inflation Expectations, U.S. and Argentina Online Experiments



Notes: The total number of observations are 783 for the U.S. and 567 for Argentina's sample II. These observations correspond to the control group only in both cases. The figures are binned scatterplots. The darker markers represent the average inflation expectations for quantiles of inflation perceptions (12 quantiles for the U.S. and 24 for Argentina). The solid line represents the 45 degree line.


Figure 3: Comparison of Inflation Expectations between U.S. and Argentina Online Experiment Samples and Third-Party Samples




Notes: Both figures plot the distribution of inflation expectations for following 12 months for each country according to two different sources. Panel a) presents the distribution for the United States from the control group of our U.S. Online Experiment sample (N=697 – September 2013 observations only) and from the University of Michigan's Survey of Consumers (N=468 – September 2013 wave),. Panel b) presents the distribution of inflation expectations for the control group of our Argentina Online Experiment sample II (opinion poll, N=567) and from the Universidad Torcuato Di Tella's Encuesta de Percepciones de Inflación (N=1,878), both for the month of April 2013 .

Figure 4: Reduced Form Evidence: Rational Inattention Test, U.S. and Argentina Online Experiments

a. United States



b. Argentina



Notes: For panel (a), we use the U.S. Online Experiment sample, with 783 observations from the *Control* group, 807 in the *Statistics (1.5%)* treatment and 156 observations from the *Products (2%)* and *Products (3%)* groups. For panel (b), we use observations from the Argentina Online Experiment sample I, with 182 observations from the *Control* group, 161 observations in the *Statistics (24%)* group and 135 observations in the *Products (24%)* group. ES is the Epps–Singleton characteristic function test of equality of two distributions. The histograms are censored at -5% and 15% (inclusive) in panel (a) and at 5% and 55% in panel (b), but these bins represent the cumulative observations below the minimum and above the maximum for each country.

Figure 5: Reduced Form Evidence: Irrational Learning and Spurious Learning Tests, U.S. Online Experiment

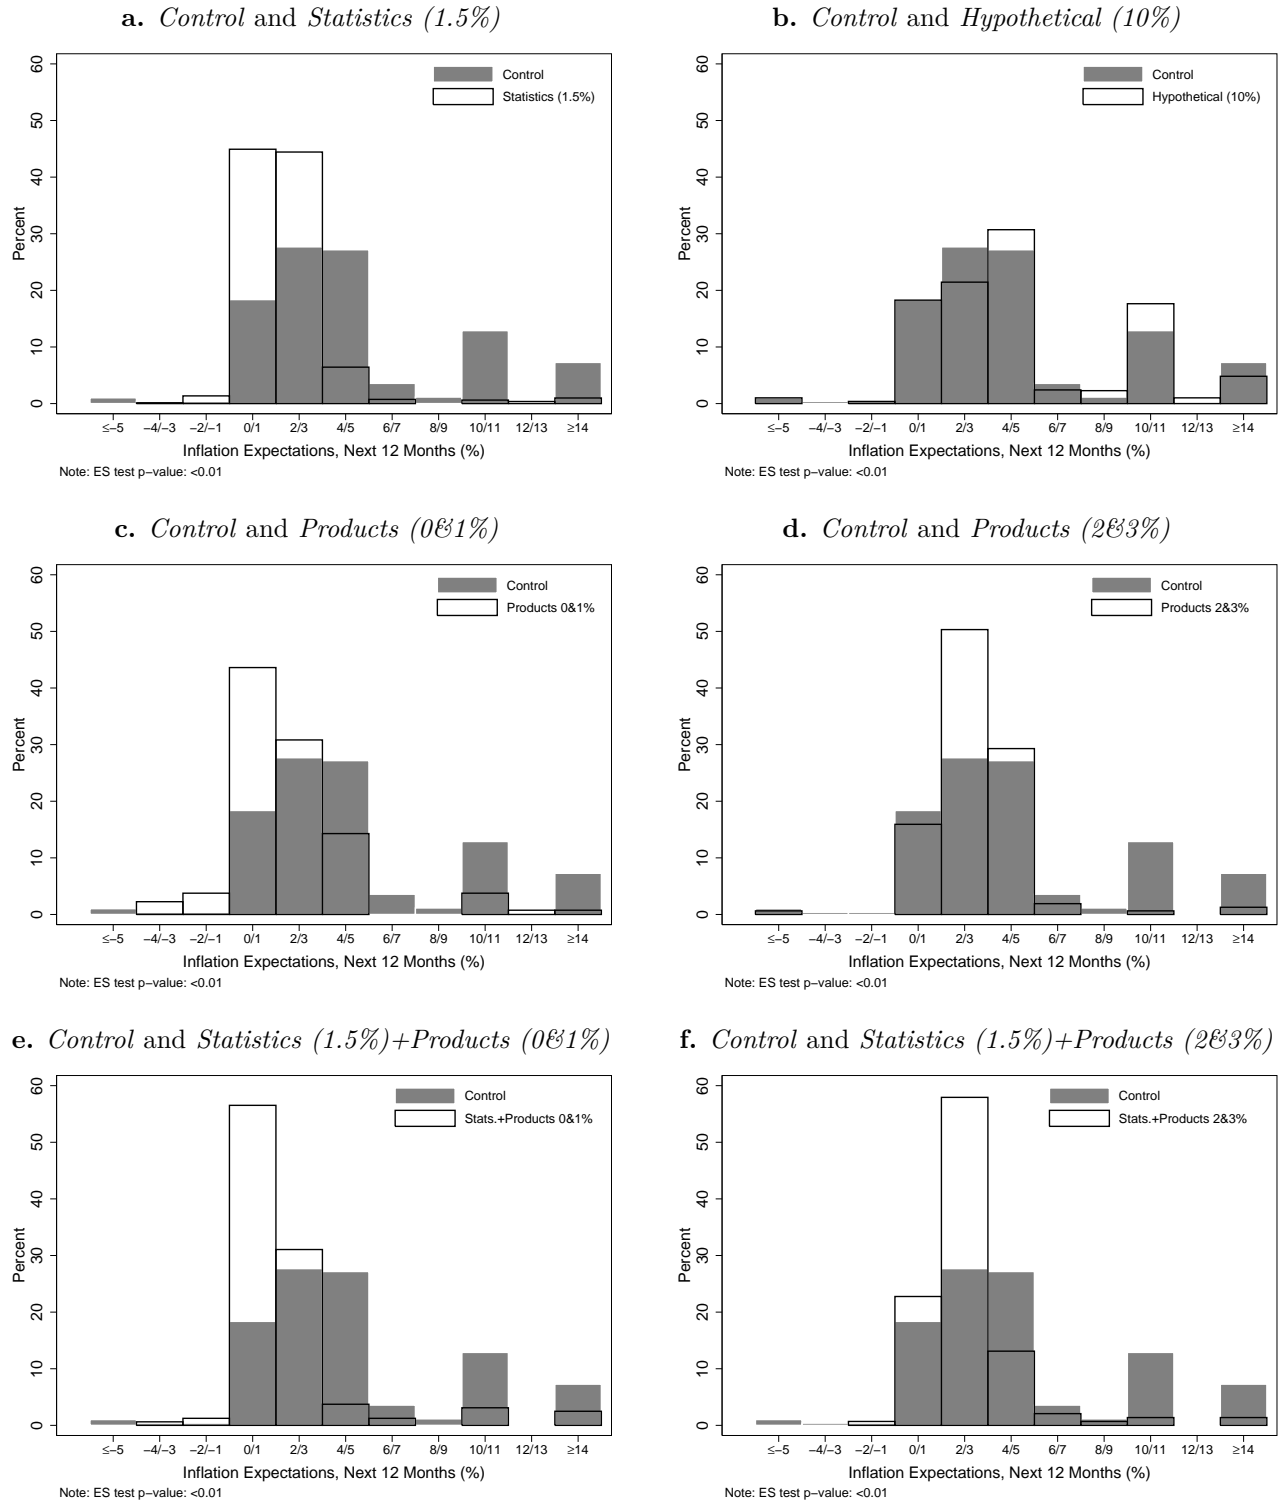




Figure 6: Effects of *Products* Treatments, U.S. and Argentina Online Experiments

a. United States

a.i. Effects on Inflation Expectations




a.ii. Effects on Confidence in Inflation Expectations
(Higher is More Confident)




b. Argentina

b.i. Effects on Inflation Expectations




b.ii. Effects on Confidence In Inflation Expectations
(Higher is More Confident)



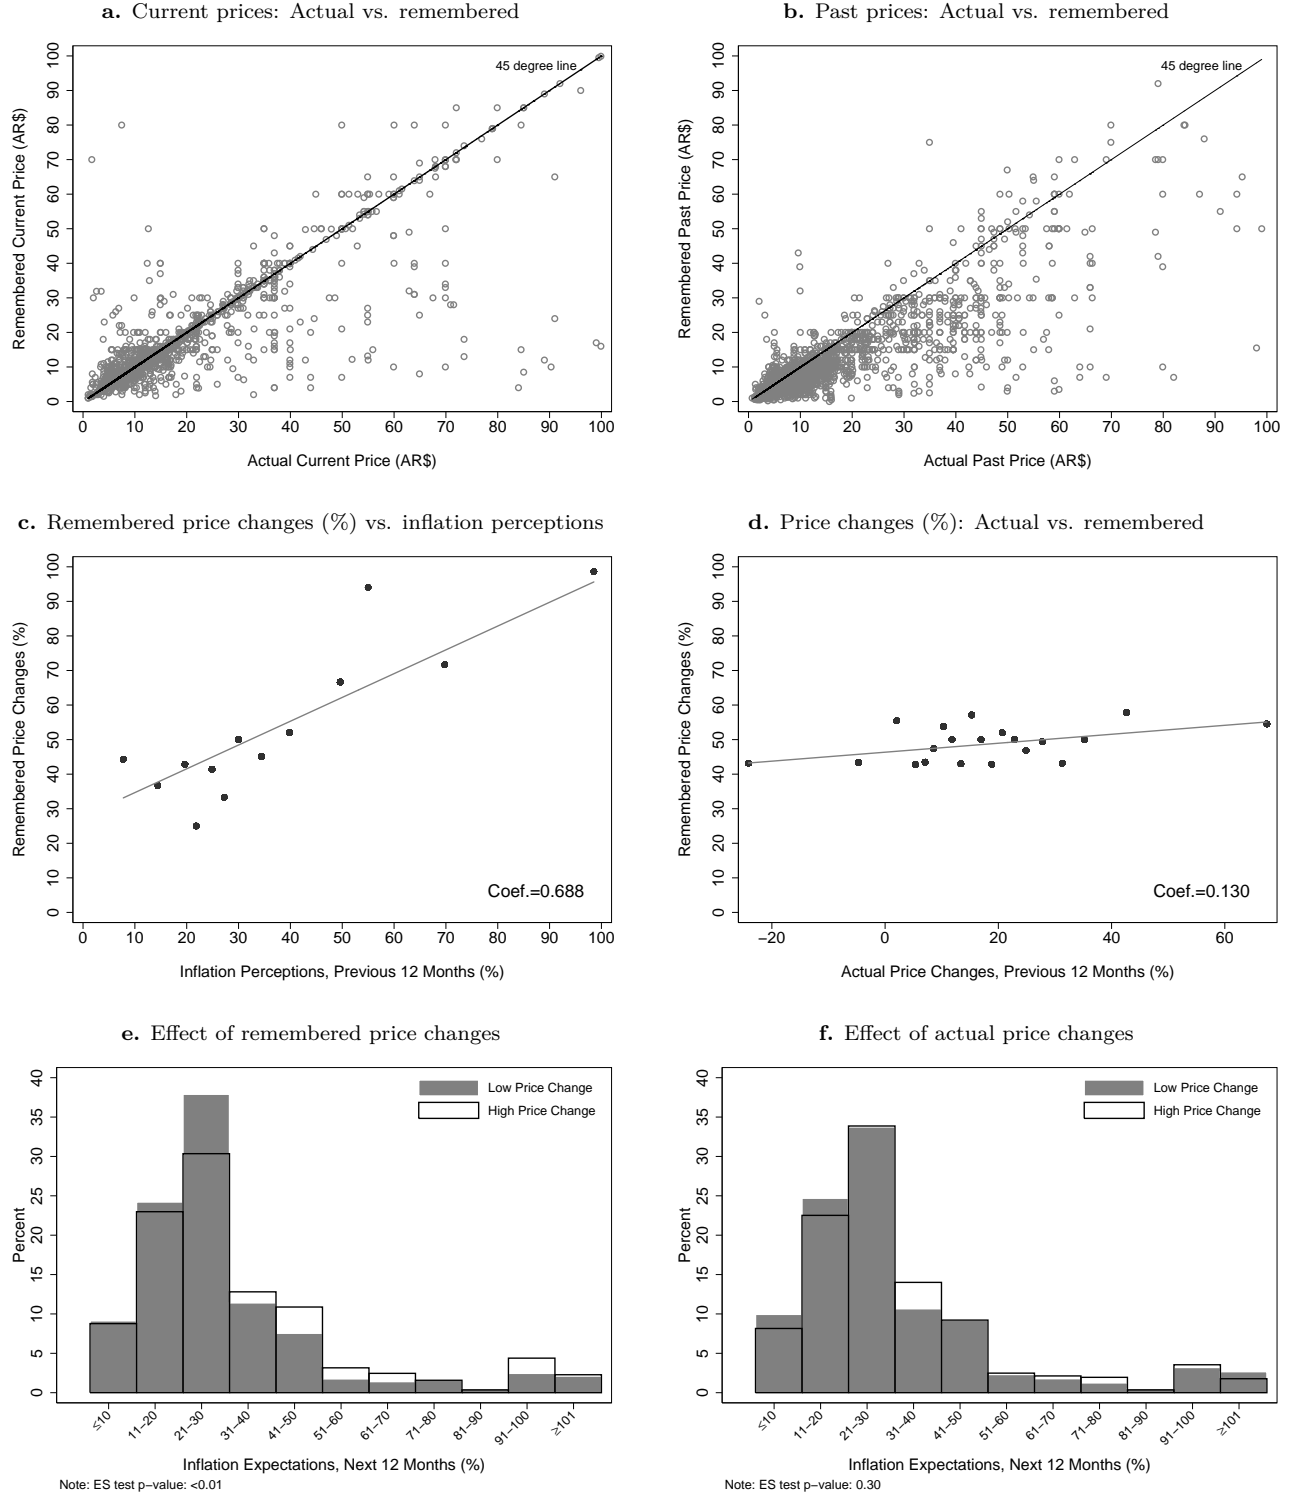
Notes: The results in Panel (a) were obtained from the U.S. Online Experiment sample, with a total of 1,546 observations (783 from the control group and 763 in 10 variations of *Products* treatment). The results in Panel (b) correspond to the Argentina Online Experiment sample II, with a total of 3,686 observations (567 in the control group and 146-181 in each of the 19 *Products* treatment groups). Each bar represents the point estimate of the effect of the specific sub-treatment (average product price changes in the table presented) compared to the control group, with confidence intervals computed with robust standard errors. The confidence variable from Panels a.ii and b.ii corresponds to answer to a question about respondents' self-confidence on their inflation expectations, with higher values denoting higher confidence, and standardized to have a standard deviation of 1.

Figure 7: Heterogeneity in Learning Rates, U.S. and Argentina Online Experiments



Notes: The total number of observations in Panel (a) (U.S. Online Experiment) is 1,552 (789 in the control group and 763 in the 10 variations of the *Products* treatment). The total number of observations in Panel (b) (Argentina Online Sample II) is 3,653 (567 in the control group and 3,086 in the 19 variations of the *Products* treatment). For both panels, the “High Confidence” group corresponds to those who answered “sure” or “very Sure” when asked their confidence on their inflation perceptions, and the “Low Confidence” group corresponds to those who replied “somewhat sure,” “unsure” or “very unsure.” The bars represent the coefficient α obtained from the regression described by equation (6) for each of the pairs of mutually exclusive groups. The confidence intervals were computed with robust standard errors.

Figure 8: Results from the Supermarket Experiment



Notes: The data in this Figures corresponds to the Argentina Supermarket Experiment. The total number of observations is 1,140 for Panels (a), (b), (c) and (d), and 1140 for Panel (e) and 1127 for Panel (f). Panels (c) and (d) represent binned scatterplots, where the number of observations are almost identical across bins. The annual price changes in panels (c) and (d) are implicit; they are obtained from the current and past prices in pesos (AR\$) reported by the respondents. In panel (e), the *High/Low Price Change* corresponds to individuals above/below the median of remembered price changes, after controlling for inflation perceptions. In panel (f), the *High/Low Price Change* corresponds to individuals above/below the median actual price changes, after controlling for inflation perceptions. ES is the Epps–Singleton characteristic function test of equality of two distributions.

Table 1: Estimates of Learning Rates, Online Experiments

a. United States					
	(1)	(2)	(3)	(4)	(5)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}^{follow-up}$	$i_{i,t+1}$
β	0.757*** (0.033)	0.817*** (0.058)	0.814*** (0.046)	0.438*** (0.055)	0.291*** (0.040)
α -Products	0.689*** (0.036)	0.449*** (0.050)	0.697*** (0.045)	0.336** (0.150)	0.499*** (0.135)
α -Statistics	0.838*** (0.034)	0.283*** (0.063)	0.799*** (0.058)	0.360*** (0.138)	0.314 (0.212)
α -Hypothetical	0.232*** (0.027)		0.215*** (0.046)	-0.021 (0.092)	0.131 (0.112)
Observations	3,141	1,587	1,073	1,073	3,141
Simultaneous treatments	No	Yes	No	No	No
b. Argentina					
	(1)	(2)	(3)	(4)	(5)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}^{follow-up}$	$i_{i,t+1}$
β	1.138*** (0.118)	0.902*** (0.042)	0.963*** (0.041)	0.754*** (0.086)	0.155*** (0.035)
α -Products	0.458*** (0.062)	0.494*** (0.027)	0.456*** (0.037)	0.208** (0.094)	0.468*** (0.133)
α -Statistics	0.432*** (0.098)				
Observations	691	3,653	1,320	1,320	3,373
Sample (experts, online)	I	II	II	II	II

Notes: * significant at the 10% level, ** at the 5% level, *** at the 1% level. Heteroskedasticity-robust standard errors in parenthesis. The source for the data in Panel (a) is the U.S. Online Experiment sample. The source for the data in Panel (b) is the Argentina Online samples I (college graduates) and II (opinion poll). The α and β coefficients are obtained from the regression given by equation 6, section 2.3: $\pi_{i,t+1} = \gamma_0 + \gamma_1 \pi_{i,t}^0 + \gamma_2 (\pi_{i,t}^T - \pi_{i,t}^0)$, where $\pi_{i,t}^0$ is the respondent's stated past inflation perception, $\pi_{i,t}^T$ is the mean inflation provided in the treatment, and $\pi_{i,t+1}$ is the post-treatment inflation expectation ($\pi_{i,t+1}$). We estimate $\hat{\alpha}$ and $\hat{\beta}$ by running this linear regression and setting $\hat{\gamma}_1 = \hat{\beta}$ and $\hat{\alpha} = \hat{\gamma}_1 / \hat{\gamma}_2$ (standard errors of this ratio computed with the Delta Method). The parameter β represents the rate of pass-through from perceptions of past inflation to future inflation expectations. The parameter α captures the weight the individual assigns to the information provided in the experiment relative to her prior belief. In Panel (a), the results presented in column (2) represent the case of the *Products+Statistics (1.5%)* combined treatment, in which treated individuals received two pieces of information simultaneously. In both Panels, the dependent variable in columns (1) to (3) is inflation expectations (for the following 12 months) at the time of the original survey, with the sample restricted in column (3) to a subset of respondents who were re-interviewed after the original survey (two months in the U.S. online experiment and four months in Argentina). The dependent variable in column (4) is inflation expectations (for the following 12 months) at the time of that follow-up interview. The dependent variable in column (5) is the expected interest rate (for the following 12 months) in the original survey. For the number of observations in each treatment group, please refer to Section 3.1.