

Inflation Expectations, Learning and Supermarket Prices

Evidence from Survey Experiments*

Alberto Cavallo

MIT and NBER

Guillermo Cruces

CEDLAS-FCE-UNLP and IZA

Ricardo Perez-Truglia[†]

Microsoft Research

This version: May 2015.

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 Borge, Sebastian Di Tella, Emmanuel Farhi, Matthew Gentzkow, N. Gregory Mankiw, Markus Mobius, Andrés Neumeyer, Roberto Rigobon, Tanya Rosenblat, Tavneet Suri, Martin Tetaz, Glen Weyl, Fernando Yu, and participants in the seminars at Universidad de San Andres, Universidad Torcuato Di Tella, Microsoft Research, MIT Sloan, Chicago/NYU International Macro Finance Conference and Harvard University. Julián Amendolagine 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., 2012; 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, 2013; 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 disentangle the roles 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.¹ 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%.² 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

¹The 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.

²We do not use official inflation statistics for Argentina because they are 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, individuals 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.³ The latter set of data was conceived as a proxy for the type of information that individuals would obtain from their own personal experience as shoppers. A rational individual would be expected to pay much more attention to inflation statistics, because their precision and representativeness is an order of magnitude larger than the precision and representativeness of the average price change computed over a handful of supermarket products. 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.⁴ 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 (i.e., the price changes of a few familiar products, such as bread and milk).⁵ In other words, subjects incorporated in a similar measure information about and information based on price changes for thousands of products.

While this evidence suggests that inefficient use of the information available may be a significant source of information frictions, there are still some important caveats. First, subjects may have reacted to the information on specific products because they perceived it as accurate, but they may still not trust their own memories about supermarket prices. Second, using price memories in the formation of inflation expectations is irrational only insofar as those memories are inaccurate.⁶ To address these remaining questions, we conducted a consumer-intercept survey experiment with some unique features at several branches of a supermarket chain in Argentina. We recorded consumers'

³The 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.

⁴An alternative explanation might be that respondents interpreted our survey question as addressing the prices of the specific goods they purchased rather than their country's overall inflation rate. If that were the case, individual product prices could be considered more informative. However, we find that information about product prices also has a significant impact on expectations of other variables, such as the nominal exchange rate and the nominal interest rate, which would only make sense insofar as individuals interpreted the question as referring to the average price level. Following the University of Michigan's Survey of Consumers methodology, the phrasing of our questions referred to changes in "prices in general." See the online Questionnaire Appendix for more details and variations across different surveys and languages.

⁵This 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. In our own samples, 64.4% of the subjects in the U.S. reported trying to recall the prices of specific products when answering questions about inflation expectations, twice as many as those who report trying to recall inflation statistics. Even in Argentina, where credible (non-official) inflation statistics are regularly covered in the media, 74.9% of respondents reported trying to recall prices of specific products when asked about past inflation.

⁶This could happen, for example, if the individual was trying to correct the mismatch between her own consumption bundle and the consumption bundle used in the computation of inflation statistics. However, the evidence suggests that the differences in inflation rates arising from heterogeneity in consumption patterns are small (McGranahan and Paulson, 2006).

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 also suggests that individuals use their own memories about price changes of specific products when forming inflation expectations (the products they just purchased, in this case), and that those memories are inaccurate and thus induce large errors in expectations.

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.⁷ 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).⁸ Our experimental design includes two methods for measuring how much of the reaction to the information provided corresponds to genuine learning, and how much can be attributed to a 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 nominal interest rate and the nominal exchange 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 informational treatments is spurious. Nevertheless, our main results remain unchanged after we control for spurious learning.

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. Models of information frictions can explain a number of facts such as the frequency of price changes and the dispersion of inflation expectations. 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 accounting for costly information and limited attention, such as in the rational inattention model, is a step in the right direction. However, our evidence indicates that these models should also incorporate cognitive limitations on how individuals process the available information, such as failing to weighting information according to its precision (for example, using highly inaccurate price memories even when inflation statistics are available).⁹

⁷This criticism is common to survey experiments in general, not particularly to our application in the area of inflation expectations.

⁸See Rosenthal (1966) for a discussion of effects of this type in behavioral research, and Zizzo (2010) for a recent application to experimental economics.

⁹Coibion and Gorodnichenko (2015) use an expectations-augmented Phillips curve, incorporating household in-

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 official 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.¹⁰ 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 (Armentier et al., 2013).¹¹

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; Armentier et al., 2014). 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, 2013).

The existing evidence in the literature on the formation of inflation expectations cannot disentangle 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; Armentier et al., 2014), 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, 2013), 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, rational individuals should use information that is inaccurate as long as it is also costless. Our

inflation expectations, to account for the lack of a strong disinflation experience during the Great Recession. They argue that household expectations rose significantly during this period because individuals assigned a disproportionate weight to information about changes in gas prices, which increased during this period. Our findings support that assumption. The literature on memory in psychology and behavioral economics may be useful for developing models of information frictions (see for example Mullainathan, 2002).

¹⁰See <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.

¹¹The distribution of the bias is relevant as well. If poorer and less educated consumers have a higher level of bias, correcting it may reduce those consumers' relative disadvantage.

contribution to this literature is to provide an experimental setting that can disentangle the two sources of information frictions, rational inattention and irrational learning. We can accomplish this by exploiting variations in stakes (i.e., contexts of high vs. low inflation) and by providing different 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. (2014) 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.

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., 2014), 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 seconds after the information provision, or months after.

The main analysis consists of measuring how the information provided to individuals change 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.¹² 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 responses on inflation expectations.

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 subsection describes these treatment arms.

Figure 1 provides some samples of the pieces of information provided to the subjects in each treatment arm in the U.S. Online experiment (for snapshots of the informational treatments and the survey questions, see the questionnaire in Appendix E.3). Our first treatment arm aims to capture how individuals incorporate information from official inflation statistics. The *Statistics (1.5%)* treatment arm consisted of providing a randomly selected group of participants a table with the most recent official statistics about annual inflation at the time of the survey, including the source of the information (the price changes referred to the period from August 1, 2012 to August 1, 2013 – all the questions and the information in the survey have a twelve-month reference period). Panel (c) of Figure 1 illustrates this treatment. The table included the annual inflation implied by the Bureau of Labor Statistics’ Consumer Price Index, and the Personal Consumption Expenditures and Gross Domestic Product deflators as computed by the Bureau of Economic Analysis.¹³ 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 presented respondents with a table containing the prices of six products at the

¹²Specifically, 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).

¹³The table 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.”

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. 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.¹⁴ 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.¹⁵ Respondents in this treatment arm were randomly assigned one of the ten tables with different average price changes, which we 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.¹⁶

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 official inflation statistics.

Finally, we included a fourth treatment arm to gauge the relevance of the potential anchoring effects of the information provided (Tversky and Kahneman, 1974), 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

¹⁴The 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.

¹⁵This treatment arm covers only changes in prices of supermarket goods, which are a subset of the price changes included in the official statistics in the previous treatment arm (the Consumer Price Index, GDP deflators, etc.). The latter take into account also durable goods, services, rent and gas, among many others expenditure items. The effect of the information provided should be stronger if we included a broader set of goods in the information about changes in prices of specific products. Moreover, Cavallo (2013) shows that, in practice, supermarket prices follow closely the evolution of the CPI.

¹⁶The tables 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.”

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).¹⁷ Finding a significant degree of “learning” from this information would be suggestive of a spurious effect.

2.3 Estimating Learning Rates

2.3.1 Baseline Model

In the following sections we present some reduced-form evidence on how individuals react to randomly assigned information. The main advantage of this model-free approach is its transparency. Additionally, in this subsection 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.

The goal of the learning model is to infer how much “weight” individuals assign to a particular type of information (e.g., inflation statistics) from the joint distribution of $\{\pi_{i,t}^0, \pi_{i,t}^T, \pi_{i,t+1}\}$, even if more than one signal is provided simultaneously. Recall that $\pi_{i,t}$ and $\pi_{i,t+1}$ denote perceptions about past inflation (e.g., inflation rate over the past twelve months) and inflation expectations (e.g., expected inflation rate over the next twelve months), respectively. Individuals use information about (perceived) past inflation to form their expectations about future inflation (Jonung, 1981):

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

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 (Sargent, 1993), or some other model of expectation formation.¹⁸ None of the experiments that we conduct intend to distinguish between these different interpretations, because we want to estimate a model of learning, not a model of expectation formation.

We consider a linear specification for $f(\cdot)$: i.e., $\pi_{i,t+1} = \mu + \beta\pi_{i,t}$, 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

¹⁷The table 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.” Immediately afterward, we asked 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,” with “About 1%,” “About 5%,” “About 10%” and “About 100%” as the possible answers. See the questionnaire appendix for more details.

¹⁸The 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).

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 strong linear relationship between perceptions of past inflation and expectations of future inflation fits the 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 in the United States (panel (a)) and Argentina (panel (b)).¹⁹ 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%.²⁰ 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 can be represented by the following reduced-form equation:

$$\pi_{i,t} = g\left(\pi_{i,t}^0, \pi_{i,t}^T\right) \quad (2)$$

In our setup, we have information on these three elements $\pi_{i,t}^0, \pi_{i,t}^T, \pi_{i,t+1}$: $\pi_{i,t}^0$ is the respondent’s stated past inflation perception (pre-treatment), $\pi_{i,t}^T$ is the mean inflation (or inflation-related information) provided in one of the treatments, and $\pi_{i,t+1}$ is the respondent’s stated inflation expectation (post-treatment).

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 $\sigma_{i,t}^0$. This functional form is in fact consistent with the distribution observed in our survey data. The individual is presented with a signal about average inflation, $\pi_{i,t}^T$, which represents the price change for one product randomly drawn from the universe of products.²¹ 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 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:

¹⁹This data is for subjects in the control group, i.e., those who were not provided any information about inflation.

²⁰These proportions could be higher if we took into account the measurement error in the reporting of these variables.

²¹The results would be equivalent if the price change corresponded to an average over multiple randomly-drawn products.

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

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

$$\pi_{i,t} = (1 - \alpha_{i,t})\pi_{i,t}^0 + \alpha_{i,t}\pi_{i,t}^T \quad (3)$$

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 in the forward-looking equation (1) results in the following expression:

$$\pi_{i,t+1} = \gamma_0 + \underbrace{\gamma_1}_{\beta} \pi_{i,t}^0 + \underbrace{\gamma_2}_{\alpha\beta} (\pi_{i,t}^T - \pi_{i,t}^0) \quad (4)$$

Since $\pi_{i,t+1}$, $\pi_{i,t}^0$ and $\pi_{i,t}^T - \pi_{i,t}^0$ are all observed in our experimental data, we can estimate $\hat{\alpha}$ and $\hat{\beta}$ by simply running the above linear regression.²² 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 past and future inflation 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 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. We present results for these tests in the results section and in the Appendix.

²²One 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 . In unreported results (available upon request), we conducted an auxiliary experiment and found strong evidence that this is not a cause for concern.

This model predicts that confidence in the posterior belief, $\sigma_{i,t}$, should be higher for individuals that were provided with relevant information. The model also predicts that, for a given level of confidence in the information signal ($\sigma_{i,t}^T$), the effect of providing a signal on $\sigma_{i,t}^T$ should be independent of the particular value of the signal that was drawn ($\pi_{i,t}^T$). α should be lower for individuals with lower reported confidence in their prior belief on past inflation, $\sigma_{i,T}^0$. This model also predicts that an individual's adjustment to the new information is a linear function of the distance between the new information and her prior belief. We can test whether this prediction is accurate by estimating the basic model including an additional quadratic term, $\pi_{i,t+1} = \gamma_1 \pi_{i,t}^0 + \gamma_2 (\pi_{i,t}^T - \pi_{i,t}^0) + \gamma_3 (\pi_{i,t}^T - \pi_{i,t}^0)^2$, and testing whether $\hat{\gamma}_3 = 0$. Similarly, we can test the possibility that individuals react differently to price increases than to price decreases (Brachinger, 2008) by estimating the model $\pi_{i,t+1} = \gamma_1 \pi_{i,t}^0 + \gamma_+ \cdot 1 \{ \pi_{i,t}^T > \pi_{i,t}^0 \} \cdot (\pi_{i,t}^T - \pi_{i,t}^0) + \gamma_- \cdot 1 \{ \pi_{i,t}^T < \pi_{i,t}^0 \} \cdot (\pi_{i,t}^T - \pi_{i,t}^0)$ and then testing whether $\hat{\gamma}_- = \hat{\gamma}_+$. All these additional robustness checks are .²³

2.3.2 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. As mentioned above, when a subject is told that the annual inflation rate was 2% and is then asked about her inflation expectations, she may report an inflation expectation that is closer to 2% for spurious reasons: e.g., to show agreement with the interviewer due to a desirability bias (Goffman, 1963), to avoid being perceived as ignorant, or due to numerical anchoring (Tversky and Kahneman, 1974). These spurious effects are a major concern for our experiments and for information provision experiments in general.²⁴ To assess whether the main results are contaminated by spurious learning, our framework attempts 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. We conducted follow-up interviews with the same subjects several months after the initial experiments, in which we did not provide any new information or reminded the subject about information provided in the past. We simply elicited their inflation expectations at the time of the follow-up ($\pi_{i,t+1}^{follow-up}$).

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

²³Armantier et al. (2014) also provide related tests of Bayesian learning in the context of household perceptions about inflation.

²⁴See Rosenthal (1966) for a discussion of the effects of factors of this sort on behavioral research, and Zizzo (2010) for a recent application to experimental economics.

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 rate of 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 (3), we obtain:

$$\pi_{i,t+1}^{follow-up} = \gamma_0 + \underbrace{\gamma_1}_{\beta_{FU}} \pi_{i,t}^0 + \underbrace{\gamma_2}_{\alpha\beta_{FU}} (\pi_{i,t}^T - \pi_{i,t}^0) \quad (5)$$

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.²⁵ 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.²⁶ 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.

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. The test is based on the following intuition: among individuals in the control group, respondents who report expecting a 1 percentage point increase in inflation also report a future interest rate that is 0.3 percentage points higher. This correlation suggests that individuals partially understand the Fisher equation (Behrend, 1977). 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).

²⁵As before, an implicit assumption is that β – which is identified solely with non-experimental variation – is not subject to omitted-variable bias.

²⁶We can illustrate the intuition with an example. Let us assume that, among individuals who receive no information from us, the correlation between perceived inflation in the first survey and inflation expectations in the second survey is 0.5. Now assume that we take a group of individuals who (in the first survey) believed that past inflation was 10%, and we randomly provide some of them a signal that past inflation was 20%. If – relative to the control group – individuals who received the signal believe (in the follow-up survey) that future inflation is going to be 1 percentage points higher, that means that the information led them to believe that past inflation (as of the first survey) was 2 percentage points higher (i.e., $1/0.5$). That is, 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., $12\% = 0.8 \times 10\% + 0.2 \times 20\%$.

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 (3), we obtain:

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

Again, this corresponds to using $i_{i,t+1}$ instead of $\pi_{i,t+1}$ as dependent variable in our learning regression.²⁷ By comparing the estimated α coefficients in the two specifications, we have a second way of quantifying genuine rather and spurious learning.

3 Results from Online Experiments in the United States and Argentina

3.1 Evidence from the United States

3.1.1 Subject Pool and Descriptive Statistics

We conducted the U.S. online experiment during the month of September 2013. 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%. The subject pool for the U.S. online experiment was recruited from Amazon’s Mechanical Turk (AMT) online marketplace. 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). The final sample includes 3,945 individuals. The subjects in our sample are younger and more educated than the average U.S. citizen (the Online Appendix provides a description of the sample and a comparison with the U.S. population), but the results are similar if we re-weight the observations to make them representative on observables.

The main variables on which our analysis is based are perceptions of past inflation and expectations of future inflation. The mean for inflation perceptions is 5.07% with a median of 5% and a standard deviation of 4.02%, and the mean for inflation expectations is 5.08% with a median of 4% and a standard deviation of 5.8% (all values for the control group). Figure 2.a depicts the relationship between the two variables by means of a binned scatterplot. There is a strong positive association between the two, with a regression coefficient of 0.782 (p-value<0.01).

²⁷And, again, implicitly assumes that β is not subject to omitted-variable bias.

3.1.2 Reduced-Form Effects of the Informational Treatments on the Distribution of Inflation Expectations

The basic results of our information provision U.S. online experiment are summarized in Figure 3 (see Appendix C 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 extreme bars).

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. For instance, our informational treatment with a table depicting products with average price changes of 2% is expected to shift the mean of inflation expectations closer to 2%, and also to compress this distribution.

Panel (a) presents the results for the *Statistics* (1.5%) treatment, which consisted of providing the respondent solely with a table of official statistics about past inflation. As expected, this signal shifts the distribution of expectations towards 1.5% and makes the distribution of expectations less dispersed. Each panel in Figure 3 also presents 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). The comparisons indicate that in all cases the distributions of inflation expectations between all treatment groups and the control group are significantly different (all p-values below 1%). This indicates that our experimental subjects reacted substantially and incorporated into their inflation expectations the information on inflation statistics that we provided .

Panels (c) and (d) in Figure 3 present two examples from the *Products* treatments, in which we provided respondents with tables with the price changes, and the average of these changes, for a series of products.²⁸ Panel (c) indicates that, as expected, the signal that supermarket products increased 0-1% shifted inflation expectations towards this range, and reduced the dispersion of expectations. The distributions in panel (d) indicate that the signal that prices increased 2-3% had the same effect, although the distribution did not shift to the left as much as with the 0-1% signal. Figure 3 provides further evidence about the effects of the *Products* treatment arm. Panel (a) shows the effect of all the levels of the *Products* treatments on average inflation expectations. Each bar represents the point estimate for each of the ten sub-treatments (with average annual price changes in the tables ranging from -2 to 10% on the horizontal axis) compared to the control group. The evidence in this Figure confirms that the treatments with specific products had a systematic impact on average reported expectations. The average price changes that appear on the tables have an increasing and roughly linear impact on inflation expectations. Each percentage point increase in the average price change reported on a table as part of our treatments yielded an increase in inflation expectations of about 0.5 percentage points. These results indicate that individuals incorporate the

²⁸See panels (a) and (b) in Figure 1 for two examples of the actual information provided.

prices of specific products when forming their inflation expectations if this information is available to them.

In the treatment arm *Statistics (1.5%)+Products*, experimental subjects were provided with the table of official 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. If individuals only cared about statistics, then the inflation expectations of subjects who already receive nationally representative, aggregate official statistics on inflation should not be affected by information on the price changes of a few arbitrarily selected products. For instance, the *Statistics (1.5%)+Products (0%)* and *Statistics (1.5%)+Products (3%)* treatments should have the same effects on expectations. However, panels (e) and (f) in Figure 3 indicate that this is not the case: individuals changed their inflation expectations on the basis of information about the price changes of specific products even when aggregate representative statistics were made readily available to them.

Finally, we also included a treatment in which respondents were provided information about price changes of about 10% for fictitious products. The results from this *Hypothetical* treatment are presented in panel (b) of Figure 3. The ES test indicates a statistically significant difference between the distribution of inflation expectations for this treatment group and for the control group. This can be attributed to an small increase in density around the 10-11% range. Since this information, even though non-factual, acted as an anchor for expectations, this evidence is suggestive of the existence of non-negligible spurious reaction to the information provided.

3.1.3 Inferring Learning Rates from the Effects of the Informational Treatments

This section presents our quantification of the effects of our experiment’s informational treatments in the context of the Bayesian learning model introduced in section 2.3. The main estimates from the learning model for the U.S. online experiment are presented in Table 1. The table reports the values of α and β from equation 4.²⁹ 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.

The first pattern that emerges from Table 1 is that, consistent with panel (a) in Figure 2, there is a high correlation between inflation perceptions and inflation expectations, reflected in a relatively high value for β which varies from 0.757 to 0.817, all highly significant (for the control group only, the coefficient of perceptions in a regression with expectations as the dependent variable is 0.782). The second notable result from Table 1 is the high level of α for the factual informational treatments in columns (1) and (2). Results from the main regression in column (1) indicate that the weight given to the information in the *Statistics (1.5%)* treatment was 0.838, whereas the weight given to

²⁹We estimated this model for the *Control* group and the *Statistics*, *Products* and *Hypothetical* treatment groups (column 1), and separately for the *Control* group and the combined *Products+Statistics (1.5%)* treatment (column 2), since in that case the two pieces of information were provided simultaneously.

its equivalent in the *Products* treatment was 0.689 (the difference between the two is statistically significant at the 1% level). In the case where information about statistics and products were provided simultaneously, reported in column (2), the combined α is 0.732, which also falls in the same range.

An estimated α of about 0.7-0.85 means that, in forming their posterior beliefs about inflation expectations, individuals in our sample assigned a much greater weight to the information provided by the experiment than to their own prior belief. This is consistent with the rational inattention model (Sims, 2005; Veldkamp, 2011), which predicts that in a low-inflation country most individuals will be uninformed about inflation because the cost of misperception is low.³⁰ It is costly to acquire, update and understand inflation statistics and, therefore, individuals will only pay that cost if and when they really need to.³¹ For example, learning about inflation consumes attention, which is a limited resource that can be better used on financial information for which the stakes are higher, such as information on taxes and benefits, on how to best finance a large purchase, on the best alternatives for credit cards or mortgages, and so on.³²

A second notable result from Table 1 is that both the information from the specific product tables and the information from the official statistics treatments had significant and substantial effects on reported inflation expectations in the *Statistics (1.5%)+Products* treatment, as captured by the respective α coefficients reported in column (2). When both statistics and supermarket prices were shown, the α coefficient for the supermarket prices is 0.449, even higher than the α of 0.283 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. There are several plausible explanations for this result. Individuals may distrust official statistics, or they may fail to comprehend how representative the figures in them are. Again, this may not be surprising in a country like the United States, where the stakes for misperception of the actual inflation rate are relatively low. The same result would be surprising, though, in a country with a high level of inflation where the inflation rate is a major concern for every household. We explore this hypothesis in more depth in the online experiment conducted in Argentina (section 3.2 below).

We can also test some auxiliary hypothesis that help us establish the validity and the robustness of the Bayesian learning model used to estimate the learning rates. One prediction yielded by this model is that providing relevant information will increase the accuracy of the later belief, $\sigma_{i,t}$. We

³⁰Coibon et al. (2015) present related evidence 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.

³¹Significantly, the cost of acquiring information about inflation exceeds a simple visit to the Bureau of Labor Statistics website or other sources to check the most recent estimate of the Consumer Price Index or other measures. While that might be a simple enough task for those with some training in economics, it is not for those without that training; the cost of acquiring information about inflation includes, among other things, learning how inflation is measured and who measures it.

³²Demery 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.

can test this with our data using the respondents’ confidence in their own inflation expectations, which is self-reported in a question we included immediately after the elicitation of expectations. This variable was standardized to have a standard deviation of one. As expected, the confidence is significantly higher when individuals received factual information (*Products*, *Statistics (1.5%)* and *Products+Statistics* treatments), but not higher when the information was not factual (*Hypothetical*).³³ Moreover, the learning model also predicts that all signals from the same source, regardless of its value, should be equally informative to respondents. Figure 4, panel (b), compares the impact of each treatment level for the *Products* treatment arm on the standardized confidence variable. The different signals seem to have similar effects on respondents’ confidence in their stated expectations, although with a slight asymmetry. Indeed, we can reject at standard levels the equality of all ten coefficients (p-value 0.0475). This is suggestive evidence that individuals might be less prone to incorporate information about price decreases than about price increases.

Another test of the Bayesian model described in section 2.3 consist in testing for non-linearities or asymmetries in the reaction to the information provided (e.g., if individuals learn more from signals that are closer to their prior belief). Columns (1) and (2) in Table 2 present some robustness tests of the learning results for the *Statistics (1.5%)* treatment arm, and columns (3) and (4) present similar results for the *Products* treatment. The coefficients in columns (1) and (3) present a specification with a quadratic term (as discussed at the end of section 2.3). The corresponding estimates for this coefficient are virtually zero (0.007 and -0.003, respectively), and the linear terms for α and β are very similar to those presented in Table 1. This evidence also suggests that the Bayesian model fits the data very well. Columns (2) and (4) present the results yielded by a specification that allows differential learning for positive and negative differences between the signal and the prior belief, with a coefficient α of 0.632 (*Statistics*) and 0.606 (*Products*) for those with $\pi_{i,t}^T - \pi_{i,t}^0 \geq 0$, and of 0.859 and 0.736 for those with $\pi_{i,t}^T - \pi_{i,t}^0 < 0$. The difference between the two pairs of coefficients is statistically significant for the *Statistics* treatment (p-value of 0.08) but not for the *Products* treatment (p-value of 0.22). Thus, there is some weak evidence of a mild asymmetry, indicating that individuals seem more prone to revise their expectations downwards rather than upwards. A mechanical interpretation is that individuals with $\pi_{i,t}^T - \pi_{i,t}^0 < 0$ are those who have high perceptions of past inflation ($\pi_{i,t}^0$), and they tend to be less informed and less confident about their own prior beliefs. Appendix C.3 presents further tests of the Bayesian learning model; the results are also strongly supportive of this simple model.

3.1.4 Disentangling Genuine from Spurious Learning

While the robustness and validation checks indicate that the data is consistent with the Bayesian learning model, a pressing concern is whether or not the learning induced by our experimental setup

³³The difference in standardized confidence between the control and the *Products* treatments (pooled) is 0.226 (p-value<0.001); between the control and the *Statistics (1.5%)* treatment it is 0.324 (p-value<0.001); and between the control and the *Statistics (1.5%)+Products* it is 0.368 (p-value<0.001). The difference between the control and the *Hypothetical* treatment is a not significant and very close to zero (0.032, p-value of 0.540).

is spurious. As previously discussed, respondents may have reacted to the information provided by changing their reported inflation expectations, not their true inflation expectations, to acquiesce with the statements or information presented in the survey or for other reasons unrelated to genuine learning.

The results of our *Hypothetical* treatment arm yields a first test along these lines. This treatment arm was designed to gauge the relevance of potential anchoring effects pursuant to the provision of information (Tversky and Kahneman, 1974). These results are presented in column (1), Table 1. The coefficient α for the *Hypothetical* treatment 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.

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.³⁴ Column (3) in Table 1 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 (for all individuals except those in the combined *Statistics+Products* treatment group). The β and α coefficients are very similar to those presented in column (1) for the full sample (0.814 compared to 0.757 for β , and equally similar for the three α coefficients corresponding to the different treatment arms). Column (4) 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 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). Although they are about half as large as those in column (3), the results still indicate that 45% to 48.2% 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 is consistent with the interpretation of short-term anchoring effects, according to which the spurious effect induced by the *Hypothetical* treatment arm should disappear over time.

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. We report results from this exercise in Column (5) of Table 1, where the dependent variable is

³⁴Multiple tests suggest that selective attrition is not a concern (results not reported).

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.³⁵ The estimated α coefficient for both the *Statistics* and *Products* treatment are close to the corresponding coefficients estimated with the follow-up survey: 0.314 for the *Statistics* treatment (borderline insignificant at the 10% level) and 0.499 for the *Products* treatment (significant at the 1% level). Although the point estimates are different between the *Statistics* and *Products* treatments, we cannot reject the null hypothesis that they are equal at conventional levels. When these parameters are compared to those presented in column (1), they suggest that between 37.5% (0.314/0.838) and 72.5% (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 very similar estimates of the degree of spurious learning. The results in column (5) for the *Hypothetical* treatment arm, on the other hand, indicate that this treatment, which provided a non-factual signal, did not have a significant effect on individuals’ expected interest rates. This, again, can be interpreted as evidence that this non-factual treatment did not induce genuine learning on participants.³⁶

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. 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 below), which should not be affected if the individual was only thinking about her own idiosyncratic experience.

3.2 Evidence from Argentina

3.2.1 Subject Pool and Descriptive Statistics

In this section, we replicate the main results yielded by the U.S. online experiment with a series of samples from Argentina. The comparison of results from similar experiments in the two countries is interesting because they were at the opposite ends of the spectrum in terms of inflation experiences at the time of our study. 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

³⁵This is also 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.

³⁶We present similar results for additional tests based on alternative outcome variables in the Appendix.

rely on more stable foreign currencies.³⁷ Opinion polls in Argentina at the time of the survey systematically indicated inflation as one of the population’s primary concerns.³⁸ 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, then, individuals in Argentina should be more informed and, therefore, have stronger prior belief about past inflation than their U.S. counterparts.

The results of the Argentina online experiment are drawn from two different sets of respondents. The first group consists of a sample of college graduates (see Appendix D for details about the samples). This sample, which yielded a total of 691 observations, was assigned to a control group, or to the *Statistics* (24%)³⁹ or the *Products* treatment arms, the latter of which was divided into three sub-treatments where respondents were provided with tables showing average price changes of 19%, 24%, and 29%. The second, larger sample is based on an established public opinion research firm that carries out a quarterly online survey of adults in Argentina.⁴⁰ This sample, which yielded 3,653 responses, is also not representative of the Argentine population: while it is roughly similar in terms of age and gender composition, our sample is substantially more educated (and, therefore, richer) than the country average.⁴¹ In any case, the results are similar if we re-weight the observations to make them representative on observables. For this sample, we concentrated our efforts on a detailed version of the *Products* treatment.⁴²

The main variables on which our analysis is based are perceptions of past inflation and expectations of future inflation. For the large (opinion poll) sample, the mean inflation perception is 27.8% with a median of 25% and a standard deviation of 13.75%; the mean inflation expectation is 28.4% with a median of 25% and a standard deviation of 15.7% (all values for the control group). Panel (b) in Figure 2 provides a binned scatterplot showing the relationship between inflation perceptions and expectations. As in the U.S. sample, there is a strong, linear, and positive association between the two, with a regression coefficient of 0.883 (p-value<0.01).

³⁷See Cavallo, Cruces and Perez-Truglia (2014) for more details on the Argentine macroeconomic and institutional context at the time of our experiments.

³⁸For our opinion poll (general population) sample, 40.7% of those in our control group selected inflation as one of the three main concerns for the country.

³⁹The 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, 2014).

⁴⁰The survey has contained the same set of basic questions since 2011.

⁴¹See the Appendix for comparative descriptive statistics of our samples and the Argentine population.

⁴²The total of 3,653 respondents were randomly assigned to a control group (N=567) or to the *Products* treatment (N=3,086); respondents in the latter group were then randomly assigned to one of 19 *Products* sub-treatments with average price changes in the tables of products provided ranging from 16% to 34% in one percentage point increments.

3.2.2 Reduced-Form Effects of the Informational Treatments on the Distribution of Inflation Expectations

Figure 5 presents the results for the online experiment in Argentina. The first two panels present the results for the sample of college graduates. Panel (a) presents the distribution of inflation expectations for the control group and for the *Statistics* (24%) treatment, and panel (b) presents the distribution of the same variable for the control group and for the *Products* (24%) treatment. As in the case of the U.S. experiment, relative to the control group, providing any signal about inflation always shifted the distribution of inflation expectations towards the value of the signal, and led to a more concentrated distribution of expectations.⁴³ The ES tests suggest that these differences are always statistically significant at the 1% level.

A summary of the basic results of the *Products* experiment in Argentina is presented in panels (c) and (d) of Figure 5, which, for the sake of comparison, displays the distribution of inflation expectations for a subset of the treatment groups and for the control group of the opinion poll sample. The inflation expectations of the respondents in the *Products* (18%-19%) treatments, in which average price changes were substantially lower than ongoing inflation (the annual inflation rate at the time of the survey was 24.4%), dropped substantially, with the distribution being shifted to the left of the control group's. Conversely, inflation expectations of the respondents in the *Products* (31%-32%) treatments increased substantially, with distribution to the right of the control group's. These differences are all statistically significant (p-value of 1% or lower).⁴⁴ Another summary of the effect of the *Products* treatments is presented in panel (a) of Figure 6. Each bar represents the point estimate of the effect of the *Products* treatment for each of the ten sub-treatments compared to the control group, with average annual price changes in the tables ranging from 16 to 34% on the horizontal axis (for the opinion poll sample). The evidence in that Figure suggests that the effect of the treatment in which tables with price changes for specific products were presented was roughly linear with respect to the average price change presented in each table: each one percentage point increase in the information provided on products' average price changes yielded an increase in inflation expectations of about half a percentage point, on average.

3.2.3 Inferring Learning Rates from the Effects of the Informational Treatments

Table 3 presents the estimates of the learning rates for the Argentina online experiments. Column (1) presents the results of the learning model for the *Statistics* and *Products* treatments based on the sample of college graduates. The results in the table also indicate a very high pass-through from perceptions of past inflation to expectations of future inflation of 1.138 (significant at the 1% level). The estimated α is 0.432 for *Statistics* and 0.458 for *Products*. Notably, the coefficient α in the

⁴³For both treatment groups, the distribution of inflation expectations seems to have shifted to the left (the means are 2.2 and 1.5 percentage points lower, respectively, than the mean of 28.4% for the control group). Most importantly, the dispersion of inflation expectations was reduced (standard deviations of 6.5 for *Statistics* (24%) and 4.8 for *Products* (24%) versus 10.3 for the control group).

⁴⁴See Appendix Figure D.2 for a more detailed analysis by treatment level.

estimate for the *Products* treatment in the larger opinion poll sample (column (2)), 0.494, is very close to that of the smaller college graduates sample.⁴⁵ This evidence implies that, in forming their posterior expectations about future inflation, individuals in the two samples placed a roughly equal weight on their prior belief and on the information provided in the experiment. While substantial, the weight individuals in Argentina assigned to the information about prices changes for specific products is substantially less than the value of 0.689-0.838 (*Statistics* and *Products* treatments, respectively) that what we found for our U.S. sample.

The fact that learning rates were significantly lower in Argentina is consistent with the prediction of the rational inattention model, where individuals in a context of higher inflation would tend to be more informed because the cost of inflation misperception is higher (Mankiw et al., 2003; Carroll, 2003). An alternative explanation could be that price data are generally less credible in Argentina after the manipulation of official statistics in recent years (see Cavallo 2013). Even though we used a non-official private sector statistic in our experiments, it is possible that this situation made Argentines distrust all inflation statistics, although the fact that the lower learning applies also to information about supermarket products is not consistent with this argument. Another explanation for the difference in learning rates between countries might be found in the subject pools. However, the characteristics of individuals in our samples for the two countries are quite similar in terms of sex, age, and education levels.⁴⁶ Moreover, there is low heterogeneity in learning rates by demographic characteristics (see results on heterogeneous effects in the Appendix) when compared to the difference in the α coefficient between countries, making this an implausible explanation. Finally, this difference may be due to the underlying volatility of inflation: individuals should react more to new information in more a volatile context. However, inflation levels were relatively stable in both Argentina and the United States in the five years prior to our study. While Argentina’s general macroeconomic conditions can be deemed more unpredictable, this would imply a prediction of the opposite sign – i.e., that learning rates should be higher than in the U.S.⁴⁷

In the opinion poll sample for Argentina we replicated some of the tests of the rational model that we conducted on the U.S. online experiment data (more details and results are provided in Appendix D). As expected, we find that both the information on statistics and supermarket products increased the confidence in the posterior belief. Panel (b) in Figure 6 compares the impact of each treatment level on the standardized level of self-reported confidence about the

⁴⁵One concern with our experimental results is that they may reflect a lack of basic literacy in economics. For example, Burke and Manz (2011) show that in a laboratory experiment more economically literate individuals tend to choose more relevant information and make better use of that information. The similar results for our college graduates—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 samples suggest that economic literacy does not drive our findings.

⁴⁶See Appendix Table B1 for more details.

⁴⁷A further result can be obtained by comparing the β coefficient between the follow-up and the original samples, which provides a measure of how “persistent” beliefs are over time. 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) compared to about 78.2% (0.754 and 0.963, Table 3) for Argentina. This finding is also consistent with the prediction of the rational inattention model, since individuals in the U.S. are on average less informed about inflation and this implies that beliefs will be more volatile over time.

answer to the question regarding inflation expectations. The results suggest that, consistent with the Bayesian model, all these different signals led to the same gain in confidence about the posterior belief. Another test, which entails the alternative specification with a quadratic term, is provided in column (3) of Table 3. The results indicate that the linear terms for α and β are very similar to those presented in column (2), while the coefficient for the quadratic term is not statistically significant (it is virtually equal to zero). Column (4) in Table 3 presents the results of an alternative specification that contemplates differential learning for upward and downward corrections of the prior beliefs. The estimated coefficient α is 0.484 for those with $\pi_{i,t}^T - \pi_{i,t}^0 \geq 0$ and of 0.497 for those with $\pi_{i,t}^T - \pi_{i,t}^0 < 0$, and their difference is not statistically significant. This evidence suggest that learning was perfectly symmetric, as predicted by the Bayesian model. This result contrasts with the evidence in the U.S. sample, where we found some weakly statistically significant evidence of a mild asymmetry.

3.2.4 Disentangling Genuine from Spurious Learning

Our first test of spurious learning is based on the effects of our treatments in the medium term. Table 3 presents the results of the learning model based on a subsample of individuals in our opinion poll sample who were re-interviewed four months after the original survey.⁴⁸ This subsample of 1,320 individuals was asked again about their perceptions of past inflation and their expectations for future inflation, but they were not subjected to any type of informational treatment or reminded about the treatment in the original survey.⁴⁹ Column (5) in Table 3 presents the results of the basic regression where inflation expectations in the original survey is the dependent variable, but the parameters are estimated with the subsample of the *Products* treatment group that later participated in the follow-up survey. The α and β coefficients are very similar to those presented in column (2) for the full sample (0.963 compared to 0.902 and 0.456 compared to 0.494, respectively). Column (6) presents the regression for the same follow-up subsample, but with inflation expectations as reported in the follow-up survey as the dependent variable. The α coefficient of 0.208 is statistically significant. While it is only half as large as the coefficient in column (5), it indicates that about 45.6% 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.

In column (7) of Table 3, we present the second test of genuine learning, specifically the results of a learning equation where individuals' expectation of the nominal interest rate is the dependent variable. Notably, the α coefficient of 0.468 is very close to the value for the inflation expectations learning equation (column (1) of the same table). This estimate suggests that the vast majority

⁴⁸This is longer than the period after which we carried out our follow-up interview in our U.S. online experiment. The Argentina follow-up had to be timed with the public opinion firm's quarterly survey.

⁴⁹There was no significant difference in the probability of participating in the follow-up sample between the treatment and the control groups. As an additional robustness check, we estimated the learning regression with an attrition indicator as the dependent variable and neither α nor β was statistically significant (results not reported).

of learning is genuine rather than spurious. We carried out a similar exercise with the log of the expected nominal exchange rate of the Argentine Peso with respect to the U.S. Dollar on the free currency market. This is a key macroeconomic variable in Argentina: due to a history of high inflation, a substantial fraction of savings are held in U.S. dollars, so most individuals are aware of the market value of this exchange rate⁵⁰ and have interest in its future evolution. The α coefficient from this estimation, presented in column (8), is 0.435, that is, very close to the figure for the nominal interest rate (column (8), 0.468) and for inflation expectations (column (2), 0.494). We must note that, unlike the methodology based on follow-up surveys, this methodology suggests a lower degree of spurious learning. Given that this methodology may not remove salience effects, we prefer the estimate obtained from the follow-up survey.⁵¹

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.

4 The Supermarket Experiment

4.1 Remaining Hypotheses to be Tested

The tables for the *Products* treatments in the U.S. and Argentina online experiments indicate 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. 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. Following Bruine de Bruin et al. (2011), at the end of the surveys we asked individuals in the control group about the information they recalled when asked about inflation expectations. A 64.4% of subjects from our U.S. experiment sample report 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.

⁵⁰The perception of the exchange rate at the time of the original survey was AR\$ 8.17 per U.S. Dollar in the case of the control group, a figure very close to its actual value, with a standard deviation of only 0.66.

⁵¹The provision of information (for instance, inflation statistics) can have two effects in the short run. The first effect is learning: those who did not know the current figures incorporate this information. A second effect is salience: even for those already aware of these figures, the provision of this information makes it more salient, which may lead individuals to assign more weight to it. In Argentina at the time of our experiment most individuals followed information about the evolution of prices very closely, and thus the salience effect may have been larger. The nature of this effect implies that it will be short-lived, and it was thus likely to have disappeared by the time of our follow-up survey. This exercise could potentially have removed both spurious learning and the salience effect (which is not truly learning) at the same.

A first remaining question 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 question is that inflation expectations will be misleading only to the extent that those memories are inaccurate. For example, there could be a very rational reason for incorporating price memories in the formation of inflation expectations: if price memories were reasonably accurate, such information could be used to correct for biases in inflation statistics arising from differences between individual consumption baskets and the average basket used to compute inflation statistics.⁵² Even though this conjecture would not explain why the information about supermarket prices affected expectations about the nominal interest rate, it serves to illustrate that, as long as it is not too inaccurate, there could be ways of rationalizing the use of this information.⁵³ Even though a literature in psychology suggests that price memories are subject to large biases (Bates and Gabor (1986); Kemp (1987); Monroe and Lee (1999)), the evidence is subject to a number of caveats. For example, the evidence applies to a context of low inflation, which is subject to the rational inattention criticism. Additionally, the evidence is not collected in a natural environment, like a supermarket, and the products and brands whose prices are being remembered are not products that the subjects buy regularly.

Addressing these remaining questions would require data 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 several 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, having just made a purchase, were invited to participate in a short survey for an academic study. About half of the individuals approached agreed to participate in the survey, and interviewers reported that most of those who agreed to take part showed great interest in the exercise. A total of 1,200 subjects were interviewed for about three to five minutes. Using handheld scanners, the interviewers scanned respondents' receipt from the supermarket purchase, which contained product identifiers that could be matched to our database of scrapped online data of supermarket prices for the chain where the study was conducted. After providing purchase receipts for scanning, respondents were

⁵²Indeed, results from our U.S. online experiment indicate that a majority (72.2%) of respondents in the control group do not think that inflation statistics are representative.

⁵³Additionally, in practice differences in inflation rates due to heterogeneous consumption patterns are small (McGranahan and Paulson, 2006).

asked twelve questions. Following our basic experimental design, we asked about perceptions of the inflation rate over the past year. We then implemented some randomly assigned informational treatments, and finally we asked about expectations for the inflation rate for the following twelve months.

The first informational treatment was aimed at generating 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. The first treatment in the 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. Additionally, instead of providing the historical prices for these four products, we asked respondents to “fill in the table” by using their own price memories. Specifically, respondents were asked to recall the current price, and the price twelve months earlier, of two specific products they had just purchased. 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. This design is consistent with models of beliefs and expectations formation – for instance, Gennaioli and Shleifer (2010) propose the idea of “local thinking, in which an agent combines data received from the external world with information retrieved from memory to evaluate a hypothesis.” 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.⁵⁴

The second informational treatment was identical to the one used in the online experiments, and consisted of showing the individual the actual price histories for six randomly selected products. We randomly assigned one of three tables with average price changes of 19%, 24%, and 29%.

4.3 Accuracy of Memories about Current and Past Prices

The goal of this subsection is to compare the memories about current and past prices to the actual prices. Panel (a) in Figure 7 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 7 presents the results of a more taxing exercise for respondents’ memory: we present a scatterplot of respondents’ reported recollections of the prices of the same goods one year earlier (vertical axis) and of the actual prices one year earlier (horizontal axis),

⁵⁴Unlike 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.

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.⁵⁵ Individuals seem to underestimate the past prices of the products they had purchased, and this effect remains irrespectively of whether we use the prices of products chosen by the respondents or of products randomly selected by the interviewers.⁵⁶ 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 underestimate past prices.

Interestingly, individuals seem greatly unaware of how bad their price memories really are. We asked respondents how confident they were about their answer to questions about prices and inflation. Only 9.81% of subjects reported to be unsure (i.e., either “unsure” or “very unsure”) about their answers to the questions about prices of specific products. This high level of confidence is very similar to the level of confidence on the inflation rate over the past 12 months, about which only 9.72% responded to be “unsure” or “very unsure.”

Since individuals have relatively unbiased and accurate memories of current prices but tend to underestimate past prices, they often overestimate price changes. 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 (c) in Figure 7 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.

Panel (d) in Figure 7, in turn, presents a comparison of the remembered price changes and the actual price changes observed in our database of supermarket prices. There is not a statistically significant 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. In other words, there is a very small association between memories and reality: individuals' memories of price changes for specific products appear to be orthogonal to actual price changes.

Although individuals seem to have a poor memory about price changes for individual products, they may have a better recollection of the price of bundles of products, for instance, the price of the basket of products they had just purchased. To test this hypothesis, immediately after asking about

⁵⁵Bates and Gabor (1986) and Kemp (1987) also find that individuals' implicit price changes overestimate the actual price changes.

⁵⁶This 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).

perceived inflation, the interviewer read out loud the total amount of the purchase as reported on the receipt and asked the respondent how much they thought they would have spent twelve months earlier for exactly the same bundle of products. We compared the individual’s estimate of the change in the total purchase amount and the actual total cost according to our price database. We find similar results than those with individual products (reported in the Appendix), which indicates that respondents do not seem to fare any better when asked about total purchase amounts 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 observed heterogeneity in the data.

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

The supermarket experiment also included an informational treatment with tables of products with three levels of average price changes. Panels (a) and (b) in Figure 8 present the distributions of inflation expectations in pairwise comparisons between the *Products* treatments. While there is no statistically significant difference between the distributions of the 19% and the 24% treatments (the ES test does not reject the null of equality of distributions – p-value of 0.24), the *Products (19%)* and *Products (29%)* treatments are statistically different: average inflation expectations are clearly higher when the subjects were shown tables with the highest average price changes. This evidence merely confirms the findings from the online experiments that individuals incorporate objective information about prices of specific products.

Panel (c) 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.⁵⁷ The results from this exercise indicate that making salient products with higher remembered price changes generates higher inflation expectations. This finding is suggestive that individuals use memories of their own experience as consumers in when forming their inflation expectations.⁵⁸ As we established above, these memories are highly

⁵⁷Specifically, 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. Finally, we divided those residuals in two extreme groups: the top third (i.e., high) and the bottom third (i.e. low) of the distribution.

⁵⁸In 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.

inaccurate, so this may generate substantial biases in expectations. To show this more directly, Panel (d) 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. The comparison of the two distributions (and the result of the ES test) indicate that making salient products with actual higher price changes did not result in higher inflation expectations. 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. This is due to the fact that the price changes that our subjects remembered were nearly orthogonal to the real price changes experienced by the same products.⁵⁹

All in all, far from correcting a representativeness bias, 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, and they report to be very confident about them.

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.

Appendix E presents regression for the corresponding rate of learning, although these results should be interpreted with this caveat in mind.

⁵⁹We obtain similar results if, instead of using price changes for individual products, we use the changes in the total amount of the purchase on the receipt, which we scanned in the context of the survey (see Appendix E for more details on this additional result).

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.⁶⁰ 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 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.⁶¹

⁶⁰Of 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.

⁶¹Consistent 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 official 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., Nelson, S., Topa, G., van der Klaauw, W. and Zafar, B. (2014). “The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economics and Statistics*, forthcoming.
- [3] 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.
- [4] Atkeson, A., and Ohanian, L.E. (2001). “Are Phillips Curves Useful for Forecasting Inflation?,” FRB Minneapolis Quarterly Review (Winter) pp. 2-11.
- [5] Brachinger, H., (2008). “A new index of perceived inflation: Assumptions, method, and application to Germany,” *Journal of Economic Psychology*, vol. 29(4).
- [6] Bachmann, R., Berg, T. and Sims, E. (2012). “Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence,” NBER Working Paper No. 17958.
- [7] Badarinza, C. and Buchmann, M. (2009). “Inflation Perceptions and Expectations in the Euro Area: The Role of News,” ECB Working Paper 1088.
- [8] 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.
- [9] 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.
- [10] Behrend, H. (1977). “Research into inflation and conceptions of earnings,” *Journal of Occupational Psychology*, Vol. 50, pp. 169–176.
- [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. Available at: <http://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm>, last accessed on December 2012.
- [12] Blanchflower, D. and MacCoille, C. (2009). “The formation of inflation expectations: an empirical analysis for the UK,” Working Paper.
- [13] 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).
- [14] Burke, M. and Manz, M. (2011). “Economic literacy and inflation expectations: evidence from a laboratory experiment,” Public Policy Discussion Paper No. 11-8.

- [15] 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.
- [16] Carroll, C. (2003). “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, 118(1).
- [17] Cavallo, A., Cruces, G. and Perez-Truglia, R. (2014). “Learning from Biased Statistics Field Experimental Evidence,” Working Paper.
- [18] Cavallo, A. (2013). “Online and Official Price Indexes: Measuring Argentina’s Inflation,” *Journal of Monetary Economics*. Vol. 60 (1).
- [19] 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.
- [20] Coibion, O. and Gorodnichenko, Y. (2012). “What can survey forecasts tell us about informational rigidities?” *Journal of Political Economy*, Vol. 120, pp. 116-159.
- [21] Coibion, O., Gorodnichenko, Y. and Kumar, S. (2015). “How Do Firms Form Their Expectations? New Survey Evidence.” NBER Working Paper No. 21092.
- [22] 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).
- [23] 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.
- [24] Drager, L. (2011). “Inflation Perceptions and Expectations in Sweden - Are Media Reports the ‘Missing Link’?” KOF Swiss Economic Institute Working Paper No. 273.
- [25] Eurobarometer (2008). “Europeans’ knowledge of economic indicators,” Special Eurobarometer 323, Wave 67.2 – TNS Opinion & Social, European Commission, Brussels.
- [26] Fuster, A., Laibson, David and Mendel, Brock (2010). “Natural Expectations and Macroeconomic Fluctuations,” *Journal of Economic Perspectives*, Vol. 24 (4). pp. 67-84.
- [27] Gennaioli, N. and Shleifer, A. (2010). “What Comes to Mind,” *Quarterly Journal of Economics* 125, no. 4: 1399-1433.
- [28] 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.
- [29] Goffman, E. (1963). “Stigma: Notes on the Management of Spoiled Identity,” New Jersey: Prentice-Hall.
- [30] Hellwig, C. (2005). “Heterogeneous Information and the Benefits of Public Information Disclosures,” Working Paper.

- [31] Jonung, L. (1981). "Perceived and expected rates of inflation in Sweden," *The American Economic Review*, Vol. 71 (5), pp. 961-968.
- [32] 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.
- [33] Kemp, S. (1987). "Estimation of past prices," *Journal of Economic Psychology*, Vol. 8, pp. 181–189.
- [34] Lucas, R. E., Jr. (1972). "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4 (2): 103–124.
- [35] Madeira, Carlos and Zafar, Basit (forthcoming). "Heterogeneous Inflation Expectations, Learning, and Market Outcomes," *Journal of Money, Credit, and Banking*.
- [36] Manski, C. (2004). "Measuring expectations," *Econometrica* 72 (5), pp. 1329–1376.
- [37] 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.
- [38] 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.
- [39] Malmendier, U. and Nagel, S. (2013). "Learning from Inflation Experiences," Working Paper, Berkeley.
- [40] 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.
- [41] 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.
- [42] Morris, S. and Shin, H.S. (2002). "Social Value of Public Information," *The American Economic Review*, Vol. 92 (5), pp. 1521-1534.
- [43] Mullainathan, S. (2002), "A Memory-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, Vol. 117 (3), pp. 735-774.
- [44] Phelps, E.S. (1969). "The New Microeconomics in Inflation and Employment Theory," *American Economic Review: Papers and Proceedings*, 59.
- [45] 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.
- [46] Rosenthal, R. (1966). *Experimenter effects in behavioral research*. New York : Appleton-Century-Crofts.
- [47] 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.

- [48] Sargent, T.J. (1993). "Bounded Rationality in Macroeconomics," Oxford: Oxford University Press.
- [49] Sims, C. (2005). "Rational inattention: a research agenda," Discussion Paper Series 1, Deutsche Bundesbank.
- [50] Tversky, A., and Kahneman, D. (1974). "Judgement under uncertainty: Heuristics and biases," Science, Vol. 185, pp. 1124–1130.
- [51] Van Praag, B., and A. Ferrer-i Carbonell (2007). "Happiness Quantified: A Satisfaction Calculus Approach." Oxford: Oxford University Press.
- [52] Veldkamp, L. (2011). "Information Choice in Macroeconomics and Finance," New Jersey: Princeton University Press.
- [53] 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.

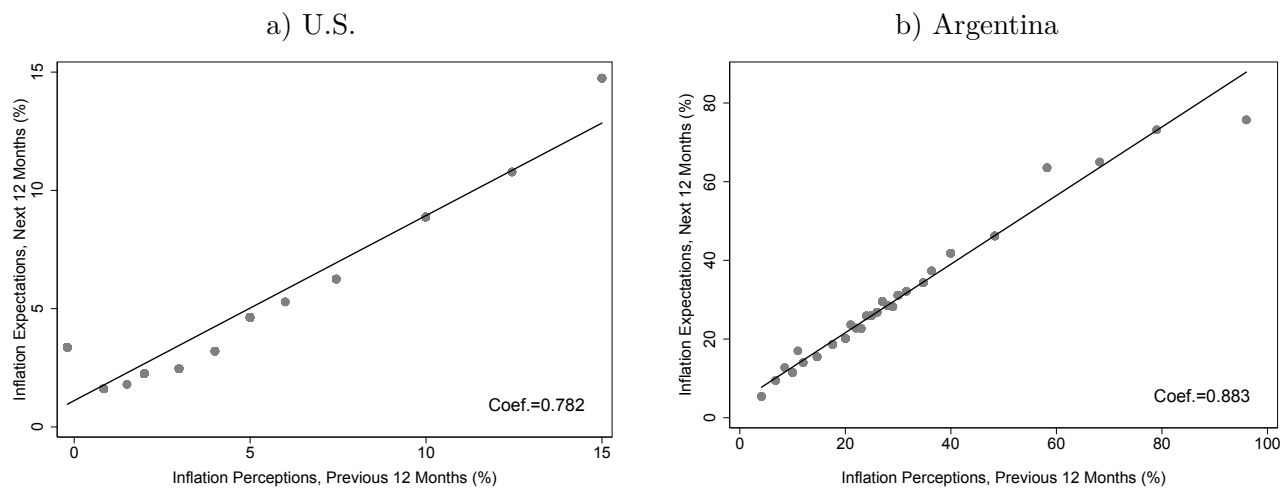
Price on January 1st 2012:	\$9.99
Price on January 1st 2013:	\$10.99

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 eyeballing these prices.

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

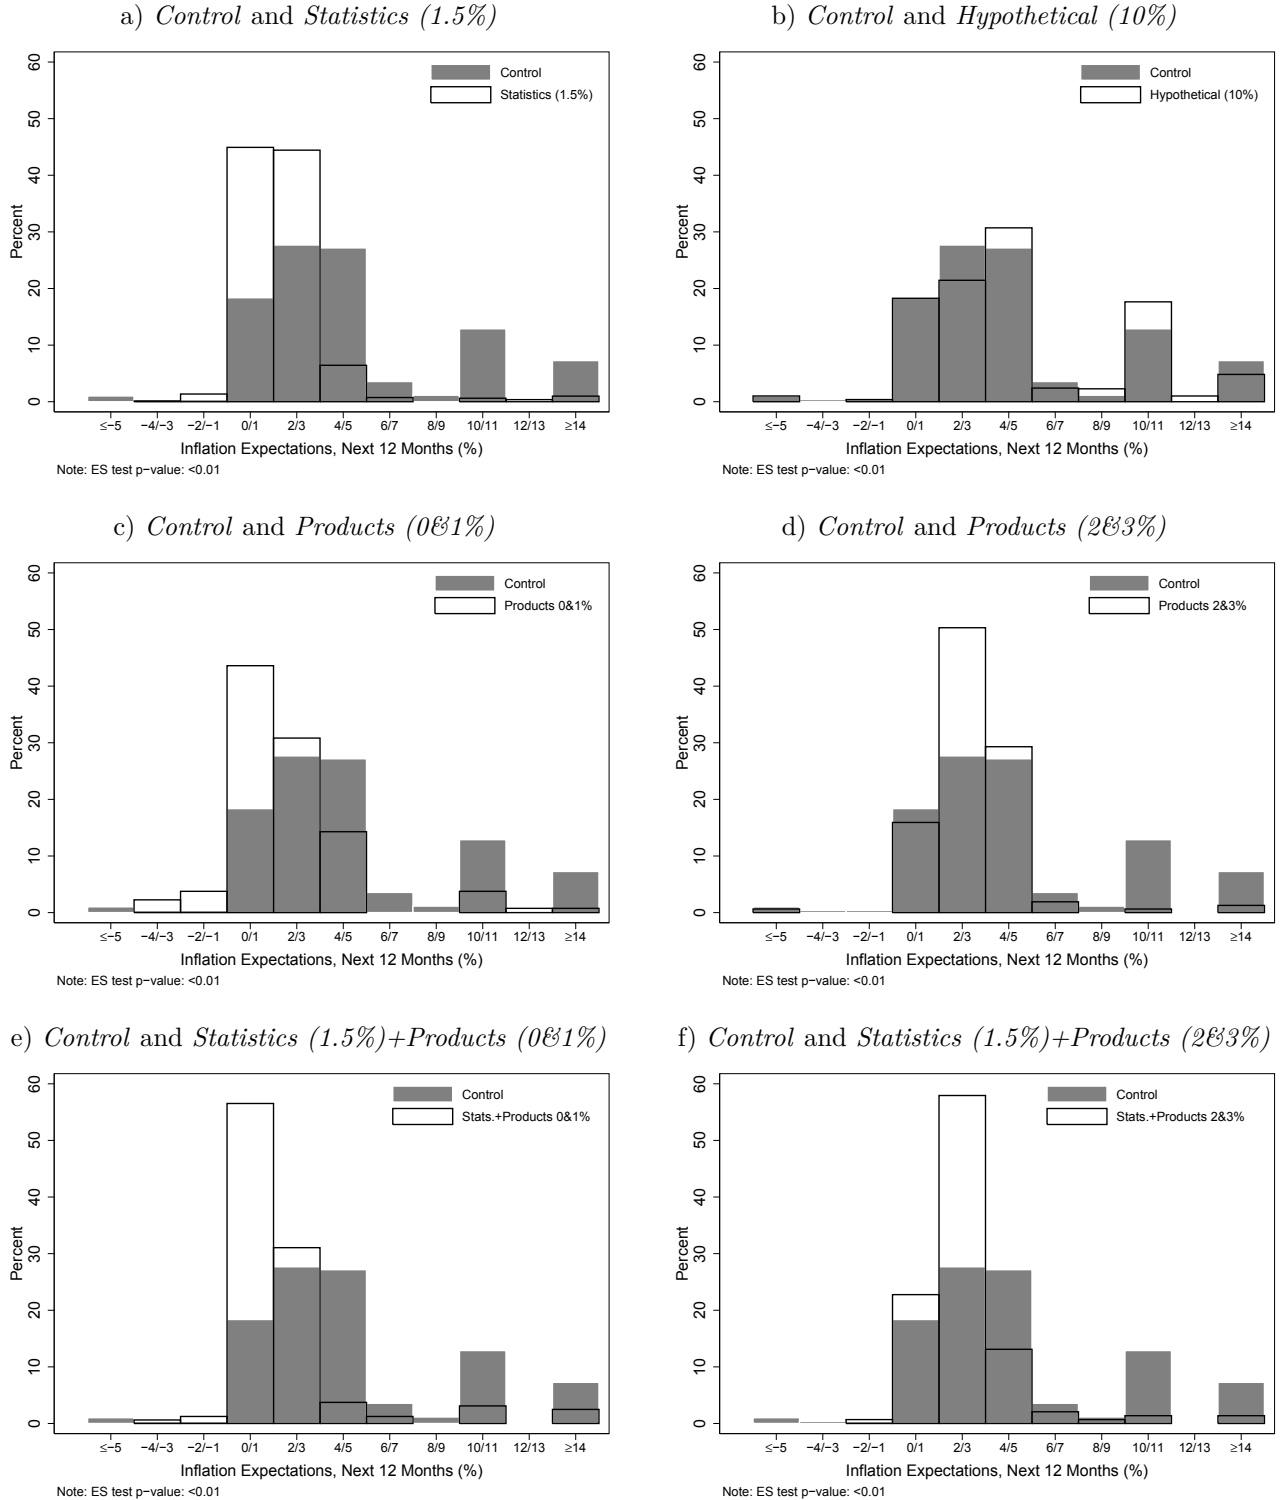
Notes: The *Products* treatment arm consisted of 10 tables similar to those presented here in panels (a) and (b). The average price changes in these tables ranged from -2% to 7% in 1 percentage point increments. The prices were obtained from scrapped online supermarket prices from one of the largest supermarket chains in the United States.

Figure 2: Past Inflation Perceptions and Future Inflation Expectations, Individuals in the *Control* group, U.S. and Argentina Online Experiments



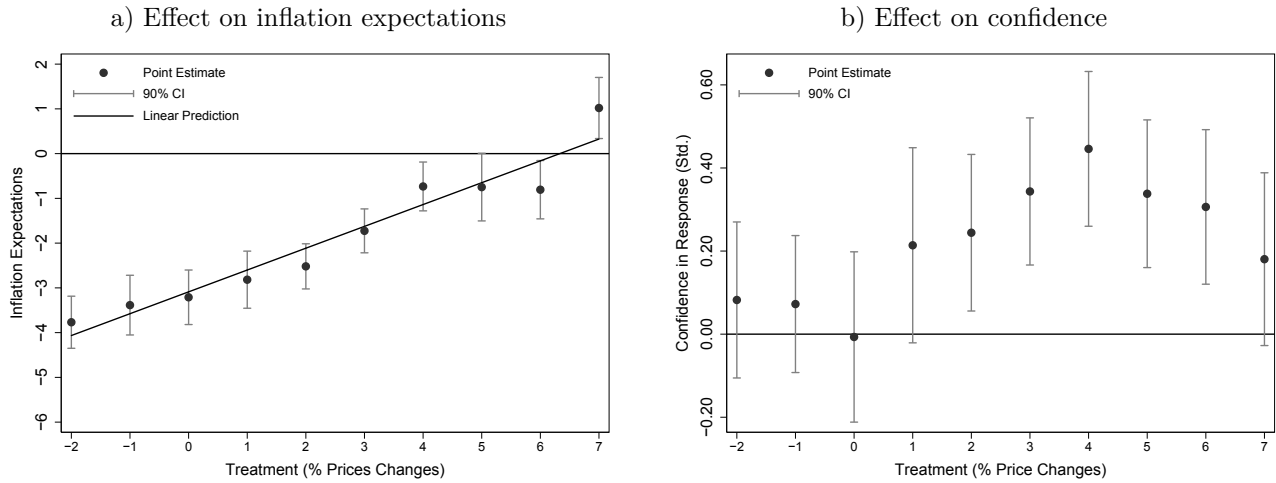
Notes: The total number of observations are 783 for the U.S. and 567 for Argentina (control group only). The darker markers represent the average inflation expectations for quantiles of inflation perceptions.

Figure 3: Inflation Expectations by Informational Treatments, U.S. Online Experiment



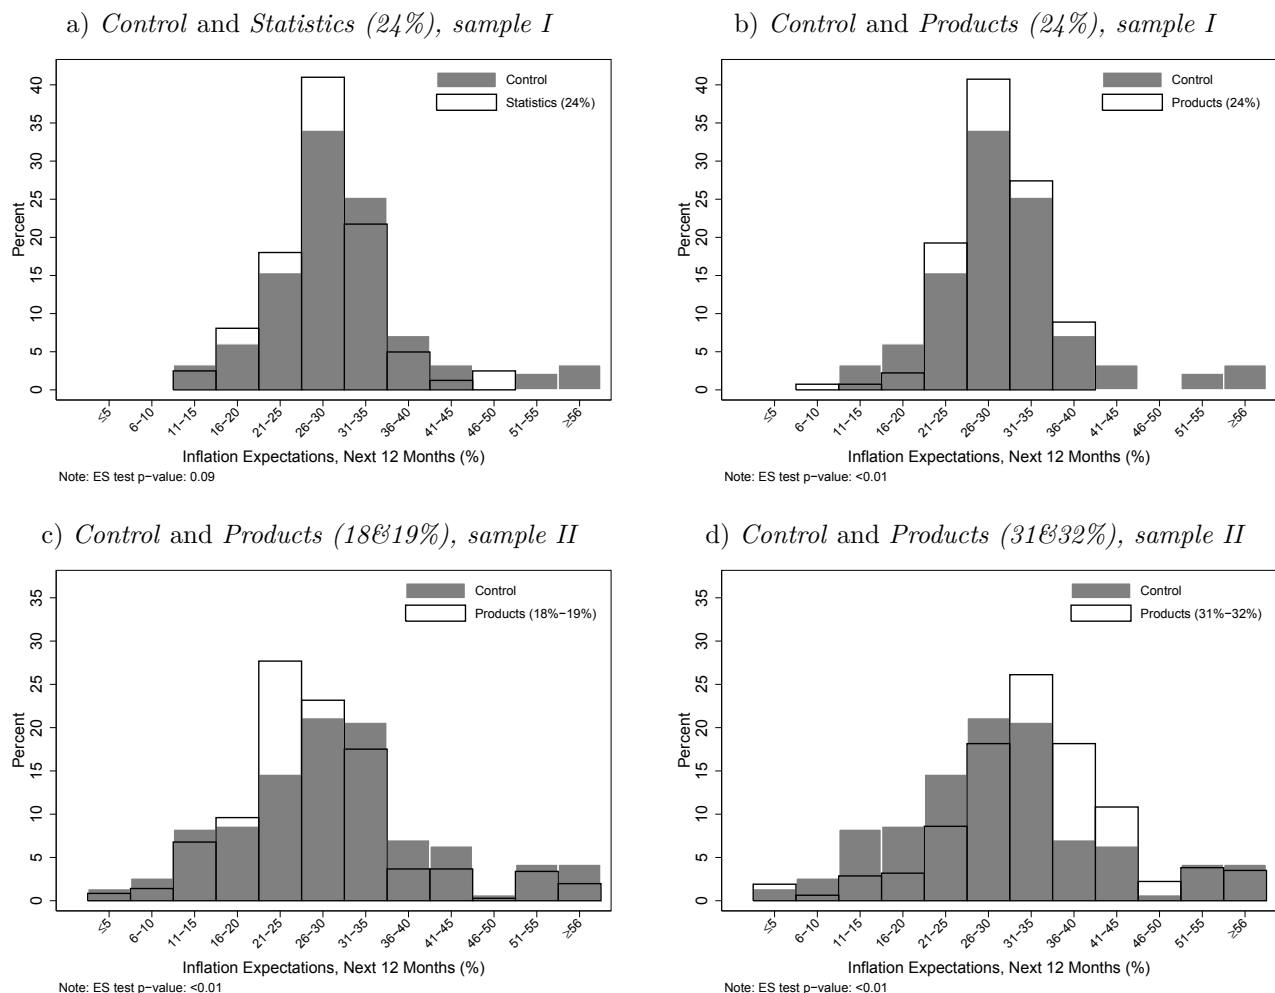
Notes: The total number of observations is 3,945, with 783 in the *Control* group, 807 in the *Statistics (1.5%)* treatment, 763 in the *Products* treatment (10 tables with average price changes from -2 to 7% in 1 percentage point increments within this treatment), 804 in the *Products+Statistics (1.5%)* combined treatment (same 10 tables as above), and 788 in the *Hypothetical* treatment. Panels (c) and (e) pool observations from the 0% and 1% average product price change tables, and panels (d) and (f) pool those from the 2% and 3% tables (see example in the previous Figure). ES is the Epps–Singleton characteristic function test of equality of two distributions. The histograms are censored at -5% and 15% (inclusive) for inflation expectations, but these bins represent the cumulative observations below -5% and above 15% respectively.

Figure 4: Treatment Effects on Inflation Expectations and Confidence about Own Expectations by Levels of *Products* Treatment, U.S. Online Experiment



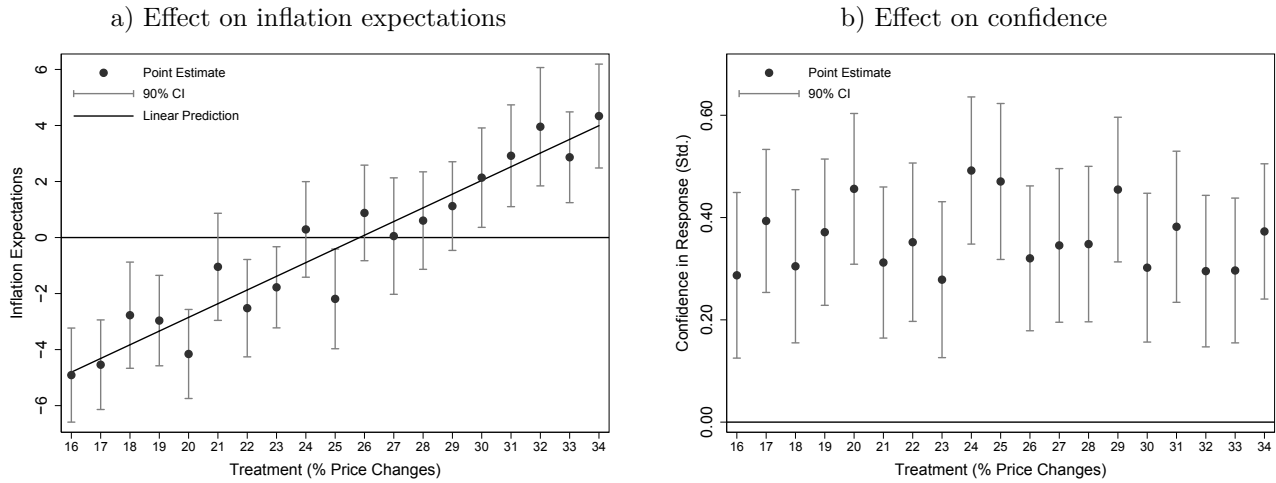
Notes: The total number of observations is 1,552 (783 in the control group and 763 in the 10 variations of *Products* treatment). Each bar represents the point estimate of the effect of the specific sub-treatment (average product price changes in the table presented) compared to the control group. Robust standard errors reported. The confidence variable from panel b) is based on a categorical question that was converted into a numerical scale using the Probit-OLS method (Ferrer-i-Carbonell and van Praag, 2008), and then standardized to have a standard deviation of one. For example, if a fraction q reports the lowest category (not sure at all), the highest confidence among the lowest category must be $\Phi^{-1}(q)$, where Φ is the cumulative distribution of a standard normal. Thus, the POLS method assigns the lowest category a score of $E[z|z < q]$, where z is distributed standard normal.

Figure 5: Inflation Expectations by Informational Treatments, Argentina Online Experiment



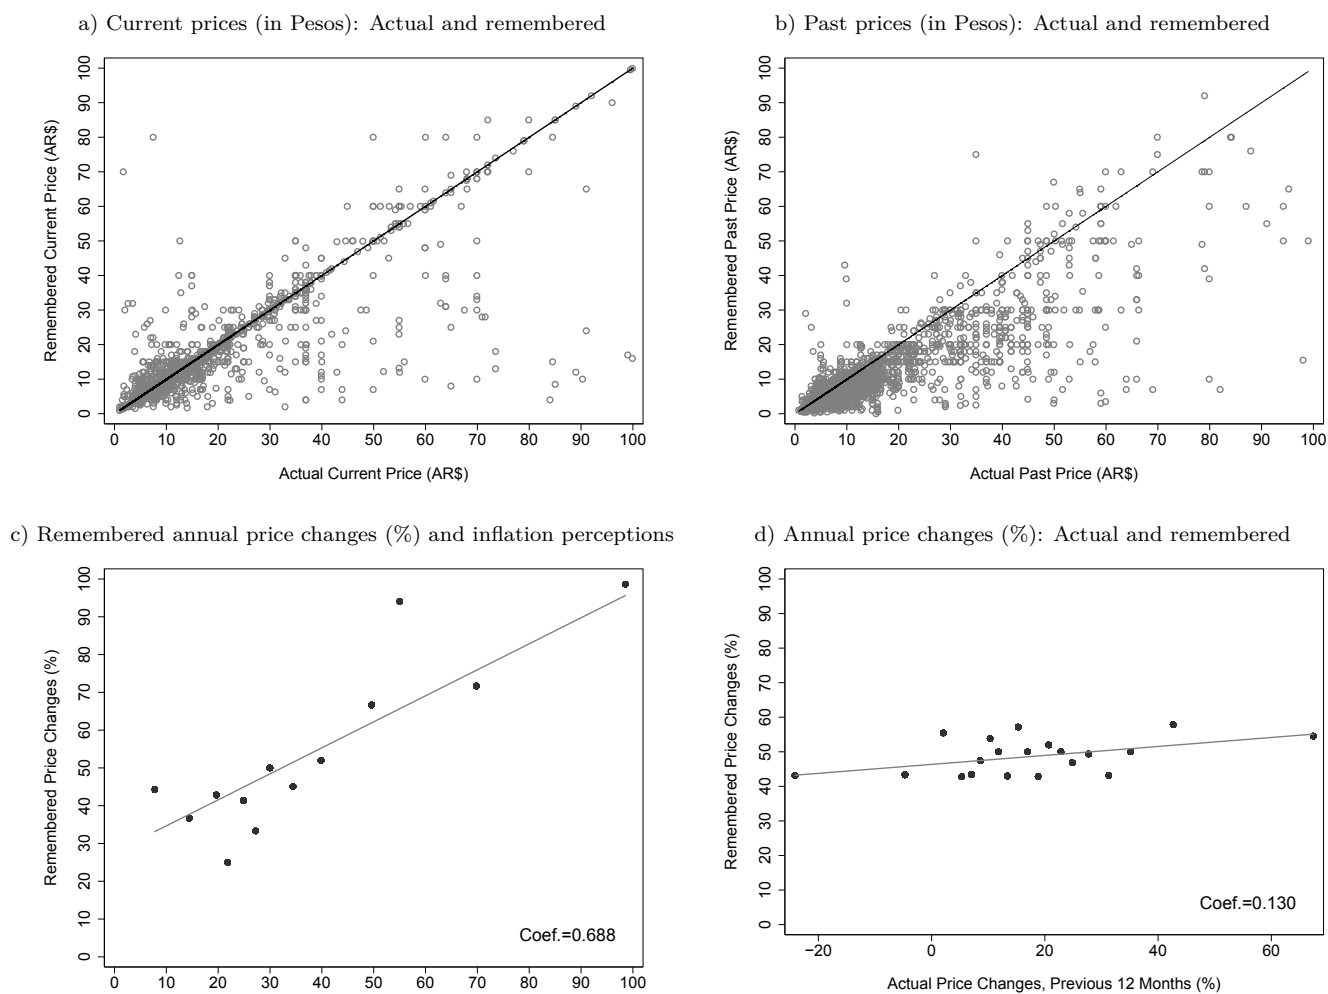
Notes: Panels (a) and (b) present results for the college graduates online experiment sample (sample I). The total number of observations is 641, with 174 in the *Control* group, 127 in the *Products (24%)* group, and 146 in the *Statistics (24%)* group. Panels (c) and (d) present results for the opinion poll online experiment sample (sample II). The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. Panel (c) pools observations from the 18% and 19% average product price change tables, and panel (d) pools those from the 31% and 32% tables. ES is the Epps–Singleton characteristic function test of equality of two distributions. The histograms are censored at 5% and 56% (inclusive), but these bins represent the cumulative observations below 5% and above 56% respectively.

Figure 6: Treatment Effects on Inflation Expectations and Confidence about Own Expectations by Levels of *Products* Treatment, Argentina Online Experiment, Opinion Poll Sample



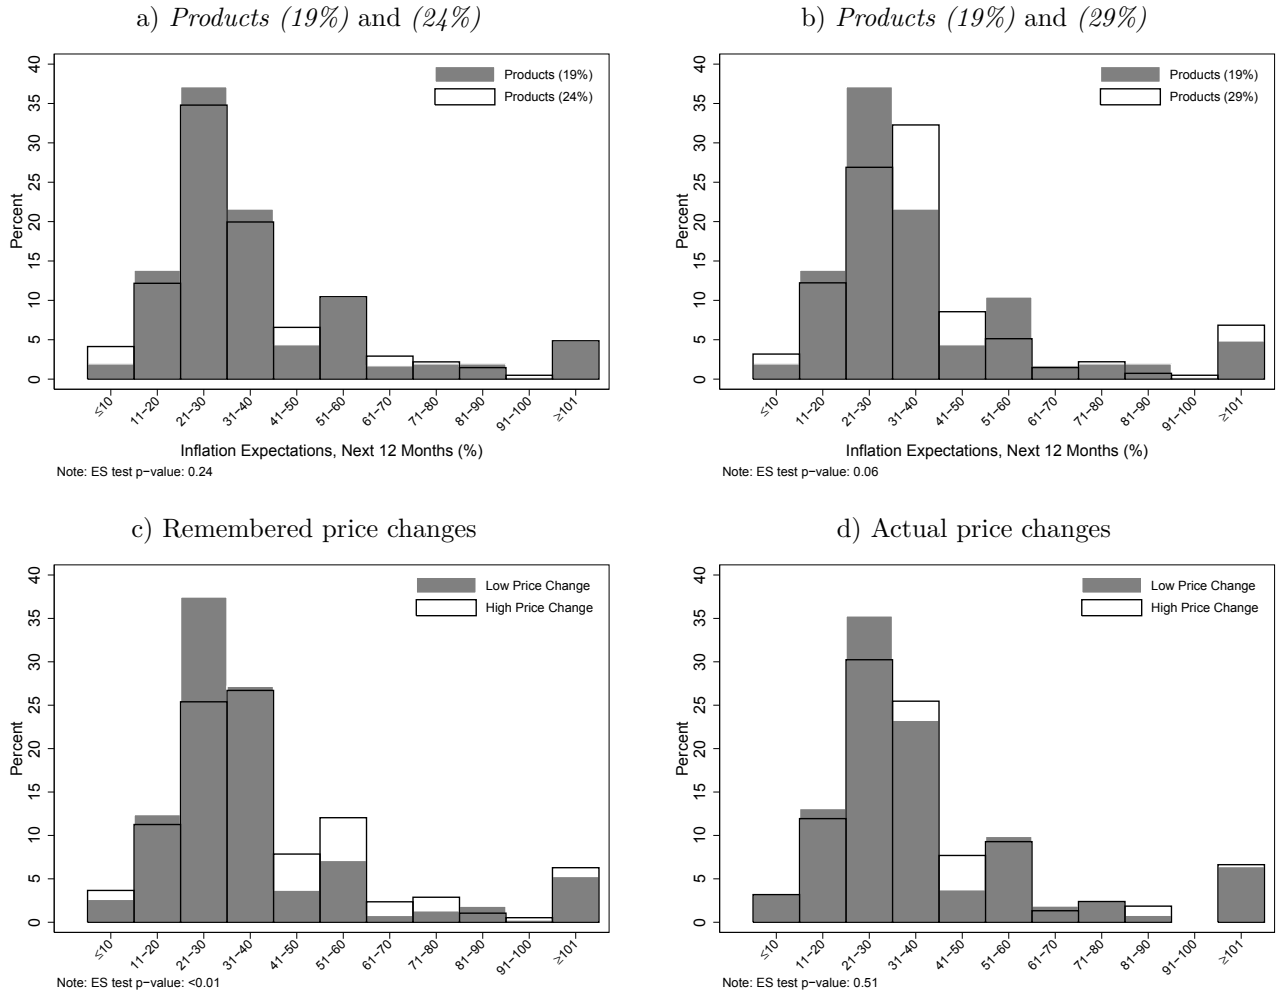
Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 *Products* treatment groups. Each bar represents the point estimate of the effect of the specific sub-treatment (average price changes for each product in the table presented) compared to the control group. Robust standard errors reported. The confidence variable from panel b) is based on a categorical question that was converted into a numerical scale using the Probit-OLS method (Ferrer-i-Carbonell and van Praag, 2008), and then standardized to have standard deviation of one. For example, if a fraction q reports the lowest category (not sure at all), that means that the highest confidence among the lowest category must be $\Phi^{-1}(q)$, where Φ is the cumulative distribution of a standard normal. Thus, the POLS method assigns the lowest category an score of $E[z|z < q]$, where z is distributed standard normal.

Figure 7: Remembered and Actual Past Prices, Implicit Price Changes and Inflation Expectations, Supermarket Experiment, Argentina



Notes: The total number of observations is 1,140. 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.

Figure 8: Inflation Expectations by *Product* Treatment Levels and by Remembered and Actual Price Changes, Supermarket Experiment, Argentina



Notes: The total number of observations is 1,232 for panels (a) and (b) (412 in the *Products (19%)* group, 411 in the *Products (24%)* group and 409 in the *Products (29%)* group). The number of observations in panels (c) and (d) are 379 and 381. In panel (c), the Low Price Change corresponds to individuals in the bottom third of remembered price changes (and, correspondingly, the High Price Change corresponds to individuals in the top third of remembered price changes). In panel (d), the Low Price Change corresponds to individuals in the lowest third of actual price changes (and, correspondingly, the High Price Change corresponds to individuals in the top third of actual price changes). ES is the Epps–Singleton characteristic function test of equality of two distributions.

Table 1: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), U.S. Online Experiment

	(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)
<i>Statistics</i>					
α	0.838*** (0.034)	0.283*** (0.063)	0.799*** (0.058)	0.360*** (0.138)	0.314 (0.212)
<i>Products</i>					
α	0.689*** (0.036)	0.449*** (0.050)	0.697*** (0.045)	0.336** (0.150)	0.499*** (0.135)
<i>Hypothetical</i>					
α	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

Notes: The α and β coefficients are obtained from the regression given by Equation 4, section 2.3. 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. The dependent variable in columns (1) to (3) is inflation expectations (for the following 12 months) at the time of the original survey (September 2013), with the sample restricted in column (3) to a subset of respondents who were re-interviewed two months after the original survey (November 2013). 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. The total number of observations in columns (1) and (5) is the sum of the 783 in the *Control* group and the observations in each treatment group (807 in the *Statistics (1.5%)* treatment, 763 in the *Products* treatments, and 788 in the *Hypothetical (10%)* treatment), with the same groups with less observations in the follow-up surveys for columns (3) and (4). The total number of observations in column (2) is the sum of the 783 controls and 804 in the *Products+Statistics (1.5%)* combined treatment. The p-value of the difference between the α coefficients for *Statistics* and *Products* in column (1) is 0.0015; the p-value of the difference between the two α coefficients in column (2) (*Statistics+Products*) is 0.0038; and the p-values of the differences between the sum of the α coefficients in column (2) (*Statistics+Products*) and the α coefficients *Statistics* and *Products* in column (1) are 0.0077 and 0.8209 respectively. Robust standard errors. *significant at the 10% level, ** at the 5% level, *** at the 1% level.

Table 2: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), Robustness Checks, *Statistics (1.5%)* and *Products* Treatments, U.S. Online Experiment

Treatment:	<i>Statistics</i>		<i>Products</i>	
	(1)	(2)	(3)	(4)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$
β	0.827*** (0.057)	0.822*** (0.059)	0.778*** (0.051)	0.775*** (0.051)
α	0.918*** (0.049)		0.690*** (0.042)	
α^2	0.007 (0.007)		-0.003 (0.005)	
α_+		0.632*** (0.108)		0.606*** (0.078)
α_-		0.859*** (0.037)		0.736*** (0.046)
Observations	1,590	1,590	1,546	1,546

Notes: The total number of observations in each column is the sum of the 783 in the *Control* group and the observations in each treatment group (807 in the *Statistics (1.5%)* treatment – columns (1) and (2) – and 763 in the *Products* treatments – columns (3) and (4)). The α and β coefficients are obtained from the regression given by Equation 4, section 2.3. α^2 represents the squared learning weight parameter. α_+ and α_- represent the learning weight parameters differentiated for those with positive and negative differences (respectively) between the reported value of the difference between the informational signal provided and the own reported value of past inflation perception, $(\pi_{i,t}^T - \pi_{i,t}^0)$. The p-values for the differences between the α_+ and α_- parameters are 0.0754 for column (2) (*Statistics*) and 0.1985 for column (4) (*Products*). Robust standard errors. *significant at the 10% level, ** at the 5% level, *** at the 1% level.

Table 3: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), Argentina Online Experiment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}^{follow-up}$	$i_{i,t+1}$	$\log(e_{i,t+1})$
β	1.138*** (0.118)	0.902*** (0.042)	0.909*** (0.043)	0.902*** (0.042)	0.963*** (0.041)	0.754*** (0.086)	0.155*** (0.035)	0.328*** (0.088)
<i>Statistics</i>								
α	0.432*** (0.098)							
<i>Products</i>								
α	0.458*** (0.062)	0.494*** (0.027)	0.472*** (0.025)		0.456*** (0.037)	0.208** (0.094)	0.468*** (0.133)	0.435** (0.173)
α^2			-0.001 (0.001)					
α_+				0.484*** (0.040)				
α_-				0.497*** (0.037)				
Observations	691	3,653	3,653	3,653	1,320	1,320	3,373	1,660
Sample	I	II	II	II	II	II	II	II

Notes: The dependent variable in columns (1) to (5) is inflation expectations (for the following 12 months) at the time of the original survey (June 2013 for sample I and March 2013 for sample II), with the sample restricted to a subset of respondents who were re-interviewed four months after the original sample II survey (August 2013). The dependent variable in column (6) is inflation expectations (for the following 12 months) at the time of the follow-up interview. The dependent variable in column (7) is the expected interest rate (for the following 12 months) in the original survey. The dependent variable in column (8) is the log of the expected nominal exchange rate of the Argentine Peso with respect to the U.S. Dollar (for the following 12 months) in the original sample II survey. Sample I is a sample of college graduates and sample II is a general population sample from the WP Public Opinion Survey (see section 3.2 for details). The total number of observations in column (1) is the sum of 182 observations in the *Control* group, 161 in the *Statistics 24%* treatment and 348 in the *Products (19%, 24% and 29%)* for the college graduates sample. The total number of observations for columns (2)-(4) is 3,653, with 568 in the control group and 146-181 in each of the 19 *Products* treatment groups for the WP Public Opinion Survey. The 1,320 observations in columns (5) and (6) represent the subsample of the WP Public Opinion Survey respondents who were re-interviewed four months after the original survey (March and August 2013 respectively). The 3,373 observations in column (7) represent the respondents of the WP Public Opinion Survey who provided a valid answer to the expected interest rate question. The 1,660 observations in column (8) represent the half of respondents of the WP Public Opinion Survey who were randomly assigned to be asked about the nominal exchange rate and provided a valid answer to this question. The α and β coefficients are obtained from the regression given by Equation 4, section 2.3. α^2 represents the squared learning weight parameter. α_+ and α_- represent the learning weight parameters differentiated for those with positive and negative differences (respectively) between the reported value of the difference between the informational signal provided and the own reported value of past inflation perception, $(\pi_{i,t}^T - \pi_{i,t}^0)$. Robust standard errors. *significant at the 10% level, ** at the 5% level, *** at the 1% level.

APPENDICES: NOT FOR PUBLICATION


A Using the Findings to Understand the Excess Dispersion in Inflation Expectations

Figure A.1 presents the distribution of inflation expectations for 2013 at the end of 2012 obtained from household surveys and professional. As previously documented in the literature on inflation expectations, the general population’s inflation expectations are substantially more dispersed than those of professional forecasters. In the U.S. the median household expectation is higher than that of the forecasters, but the difference is lower (and with the opposite sign) in the Argentine data. A related question is whether the mechanisms that we identify – the use of price memories in forming inflation expectations – could explain a small or a large share of excess dispersion in inflation expectations. The answer seems to be a lot, based on the evidence that individuals assign a significant weight to the price changes of individual products jointly with the finding of a nearly-orthogonal relationship between remembered price changes and actual price changes.

As a final empirical exercise, we illustrate how – due to the substantial dispersion in the distribution of price changes, both in low- and high-inflation contexts – even small limitations in the ability to recall prices can generate substantial dispersion in perceptions about inflation. Denote $p_{j,t}^a$ the actual price of product $j = 1, \dots, J$, with corresponding price changes for j given by $1 + \pi_{j,t}^a = \frac{p_{j,t}^a}{p_{j,t-1}^a}$. One way of modeling memory limitations is to assume individuals have perfect memory about price changes, but they can only recall prices for a limited number of products – a subset J^* . To estimate the aggregate inflation rate, individuals simply compute the average of price changes for their own basket of J^* products. Using our data on actual price changes for supermarket products, we can simulate how these perceptions vary for different values of J^* .⁶² Figure A.2 shows the distribution of annual price changes for $J^* = 5$ and $J^* = 20$, as well as the distribution of individual inflation expectations for the same time period for the U.S. (panel (a)) and Argentina (panel (b)). This Figure illustrates that even if individuals exhibited a remarkable memory and were able to perfectly recall the current and past prices of 20 products (i.e., 40 individual prices) and correctly compute all changes and their averages, the inflation perceptions resulting from these limited samples would still be substantially dispersed. This evidence complements our finding about the noisiness of individuals’ memories about specific prices. Taken together, these two pieces of evidence reinforce the case for a link between memory limitations and the heterogeneity of inflation expectations.

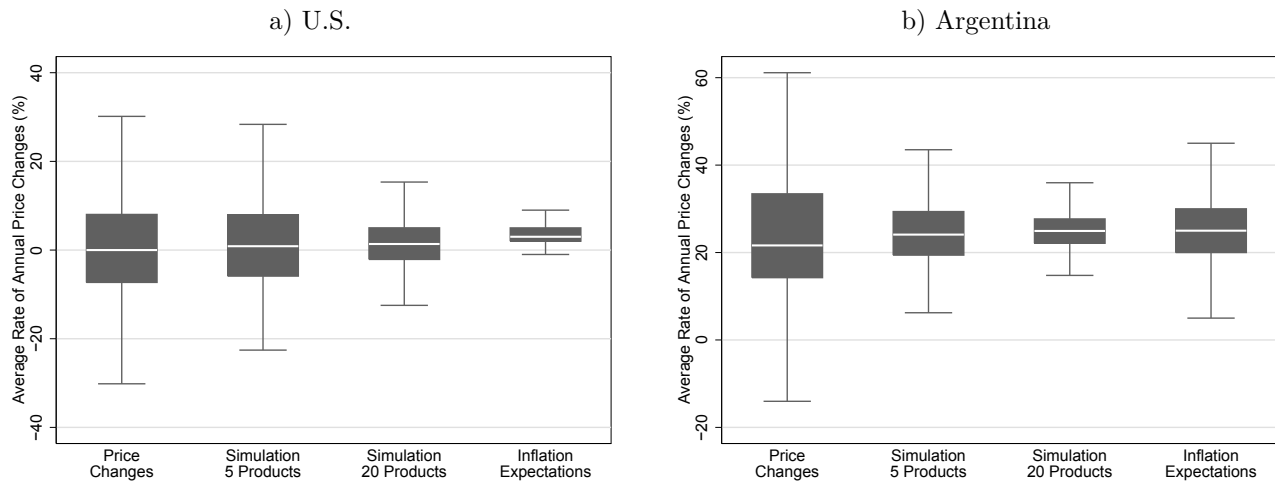
⁶²The dataset consists of 10,518 products for the U.S. and 9,276 products for Argentina, with prices observed on January 1 2012 and January 1 2013.

Figure A.1: Inflation Expectations for 2013, Household Surveys and Surveys of Professional Forecasters, U.S. and Argentina



Notes: Expected inflation for the period January 1-December 31 2013, reported in December 2012. Sources: University of Michigan’s Survey of Consumers, December 2012 (household survey, U.S., N=502), Federal Reserve Bank of Philadelphia’s Survey of Professional Forecasters, fourth quarter of 2012 (professional forecasters, U.S., N=48), WP Public Opinion Survey (household survey, Argentina, N=777) and LatinFocus Consensus Forecast, January 2013 (professional forecasters, Argentina, N=16).

Figure A.2: Price Changes from Supermarket Price Data (Total and Simulated Randomly Selected Baskets) and Inflation Expectations, U.S. and Argentina



Notes: The price changes refer to the period January 1 2012 to January 1 2013 for both countries. The first box in each panel represents the actual distribution of price changes for the products in each database (N=10,518 and N=9,276 for the U.S. and Argentina, respectively). The following two boxes represent the distributions of 1,000 simulations of average price changes for baskets of 5 and 20 randomly selected products. Inflation expectations correspond to December 2012 (University of Michigan’s Survey of Consumers for the U.S. and WP Public Opinion Survey for Argentina).

B Descriptive Statistics and Representativeness of the Subject Pools

Table B.1: Descriptive Statistics, U.S. and Argentina Samples

	Female	Age	College Degree	Observations
U.S. Online Experiment	52.6%	31.4	52.7%	3,945
U.S. Average, 18+ (ACS, 2011)	51.4%	46.5	33.4%	-
Argentina Online Experiment, Sample I	40.7%	35.0	100%	691
Argentina Online Experiment, Sample II	58.8%	42.7	54.5%	4,101
Argentina Supermarket Experiment	58.6%	47.1	41.9%	1,250
Argentina Average, 18+ (EAHU, 2012)	52.6%	43.6	26.9%	-

Notes: ACS stands for American Community Survey (U.S. Census Bureau), and EAHU stands for Encuesta Anual de Hogares Urbanos (INDEC).

C U.S. Online Experiment: Further Details and Results

C.1 Subject Pool and Descriptive Statistics

The subject pool for the U.S. online experiment was recruited from Amazon’s Mechanical Turk (AMT) online marketplace. We followed several references that describe the best practices for recruiting individuals for online surveys and experiments using AMT, and adopted some of these recommendations to ensure high quality responses.⁶³

Potential recruits were offered to participate in a short online “public opinion survey” – we avoided conditioning the subjects by using this vague description and by refraining from using words such as “economic expectations”, inflation and others. We collected data during the month of September 2013. Participants were paid \$0.50 for their participation, which is about average for this type of studies in AMT (the average duration of the questionnaire in our sample was about three minutes). We restricted the sample of participants to U.S. residents only,⁶⁴ and we included attention checks to ensure participants read the instructions and the questions thoroughly.⁶⁵ The descriptive statistics in Table B.1 indicate that, as it is common with this type of studies, subjects in our sample are younger and more educated than the average of the U.S.

We excluded from the final sample a number of participants who reported extreme values for past inflation perceptions. In the University of Michigan’s Survey of Consumers of 2012, about 98% of respondents provided an estimate for the future annual inflation rate between -5 and 15%. We restrict the sample to include inflation perceptions in that range (about 90% of the observations in our sample), which corresponds to 10 percentage points above and below the median perception in our sample (5%). It should be noted that the question about inflation perceptions precedes the informational experiment, and thus these perceptions are orthogonal to the treatments. In any case, all the results presented in the paper are robust to the inclusion of these extreme observations. See Appendix E.3 for the screen captures of the full questionnaire and for all the specific product tables.

⁶³See for instance:

Berinsky, A. J., Huber, G. A., and Lenz, G. S. (2012), “Evaluating online labor markets for experimental research: Amazon. com’s Mechanical Turk,” *Political Analysis*, 20(3), 351-368.

Crump, M.J.C., McDonnell, J.V., Gureckis, T.M. (2013), “Evaluating Amazon’s Mechanical Turk as a Tool for Experimental Behavioral Research,” *PLoS ONE* 8(3).

Paolacci, G., Chandler, J. and Ipeirotis, P. (2010), “Running experiments on Amazon Mechanical Turk,” *Judgment and Decision Making*, vol. 5, no. 5.

Rand, D. G. (2012), “The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments,” *Journal of Theoretical Biology*, 299, 172-179.

⁶⁴While Amazon checks the identity of AMT workers by requiring IDs, social security numbers, and U.S.-based bank accounts for payment, we still discarded a small number (about 2%) of IP addresses originating from outside of the U.S.

⁶⁵All of these controls were done before the experimental treatments to ensure that there is no relationship between the individuals dropped from the sample and the treatments.

C.2 Further Results

Figure 3 in the body of the paper presented the distribution of inflation expectations for selected levels for the *Products* and the *Statistics (1.5%)+Products* treatments for our U.S. online experiment. Figures C.1 (*Products*) and C.2 (*Statistics (1.5%)+Products*) present the distribution of results for all levels of these treatments from -2% average price changes to 7% average price changes in the tables presented, grouped in two one percentage point sets. The results discussed in the body of the paper are even more apparent by inspection of these two detailed figures: lower levels of specific products average price changes shifted the distribution of inflation expectations to the left, and higher levels shifted them to the right.

In the body of the paper, panel (a) in Figure 4 depicted the effect of the *Product* treatments on the average of inflation expectations, and panel (b) in the same Figure compares the impact of each treatment level for the *Products* treatment arm on the standardized confidence variable. Figure C.3 reproduces these results for different levels of the *Statistics (1.5%)+Products* treatment. Each bar in panel (a) represents the point estimate of the effect of the *Statistics (1.5%)+Products* treatment for each of the ten sub-treatments compared to the control group, with average annual price changes in the tables ranging from -2 to 10% on the horizontal axis. The evidence in panel (a) of Figure C.3 confirms that the impact of the treatments with specific products modified average reported expectations in a systematic manner. The *Products* and the *Statistics (1.5%)+Products* treatments had similar effects on the distribution of inflation expectations (panel (a)) and on the respondents' confidence on their stated expectations (panel (b)).

The learning model predicts that any heterogeneity in confidence in own prior beliefs will result in heterogeneity in learning rates. Figure C.4 presents the value of α for the *Products* treatment for different subgroups of the population (the results are qualitatively similar for the *Statistics (1.5%)* and *Statistics (1.5%)+Products* arms). Learning rates are higher for individuals with lower levels of confidence in their own reported inflation perceptions, as predicted by the learning model. On average, learning rates are also higher for those less educated, for females and for those under 30 years old, although none of the pairwise differences are statistically different from zero. This lack of heterogeneity in learning by individual characteristics may simply reflect the fact that most individuals are equally uninformed about inflation levels, which results in no significant variations in confidence about the prior belief.

C.3 Additional Test of Spurious Learning

A key assumption for the test between spurious and genuine learning is that the observational correlation between $\pi_{i,t+1}$ and the outcome variable ($i_{i,t+1}$) reflects a causal effect running from the first to the latter. For other outcomes, denoted $y_{i,t+1}$, the observational correlation with $\pi_{i,t+1}$ may suffer from substantial omitted variable bias. For example, a negative correlation between inflation expectations and expected growth rate could be due to individuals believing that inflation is bad

for growth, while a positive correlation could imply that individuals believe in some form of the Phillips curve. Alternatively, that correlation could be entirely spurious, reflecting the fact that more pessimistic individuals expect both higher inflation and lower growth. Holding this pessimism constant, that fact that an individual is induced to believe that inflation is going to be higher in the future should not affect her expectations about growth. As a result, using growth and similar outcomes as dependent variables to estimate α would lead to wildly inaccurate conclusions. Nevertheless, we can still perform a qualitative version of this falsification exercise. For each of these outcomes, we can estimate two versions of the following regression:

$$y_{i,t+1} = \alpha + \delta\pi_{i,t+1} + \varepsilon_i \quad (7)$$

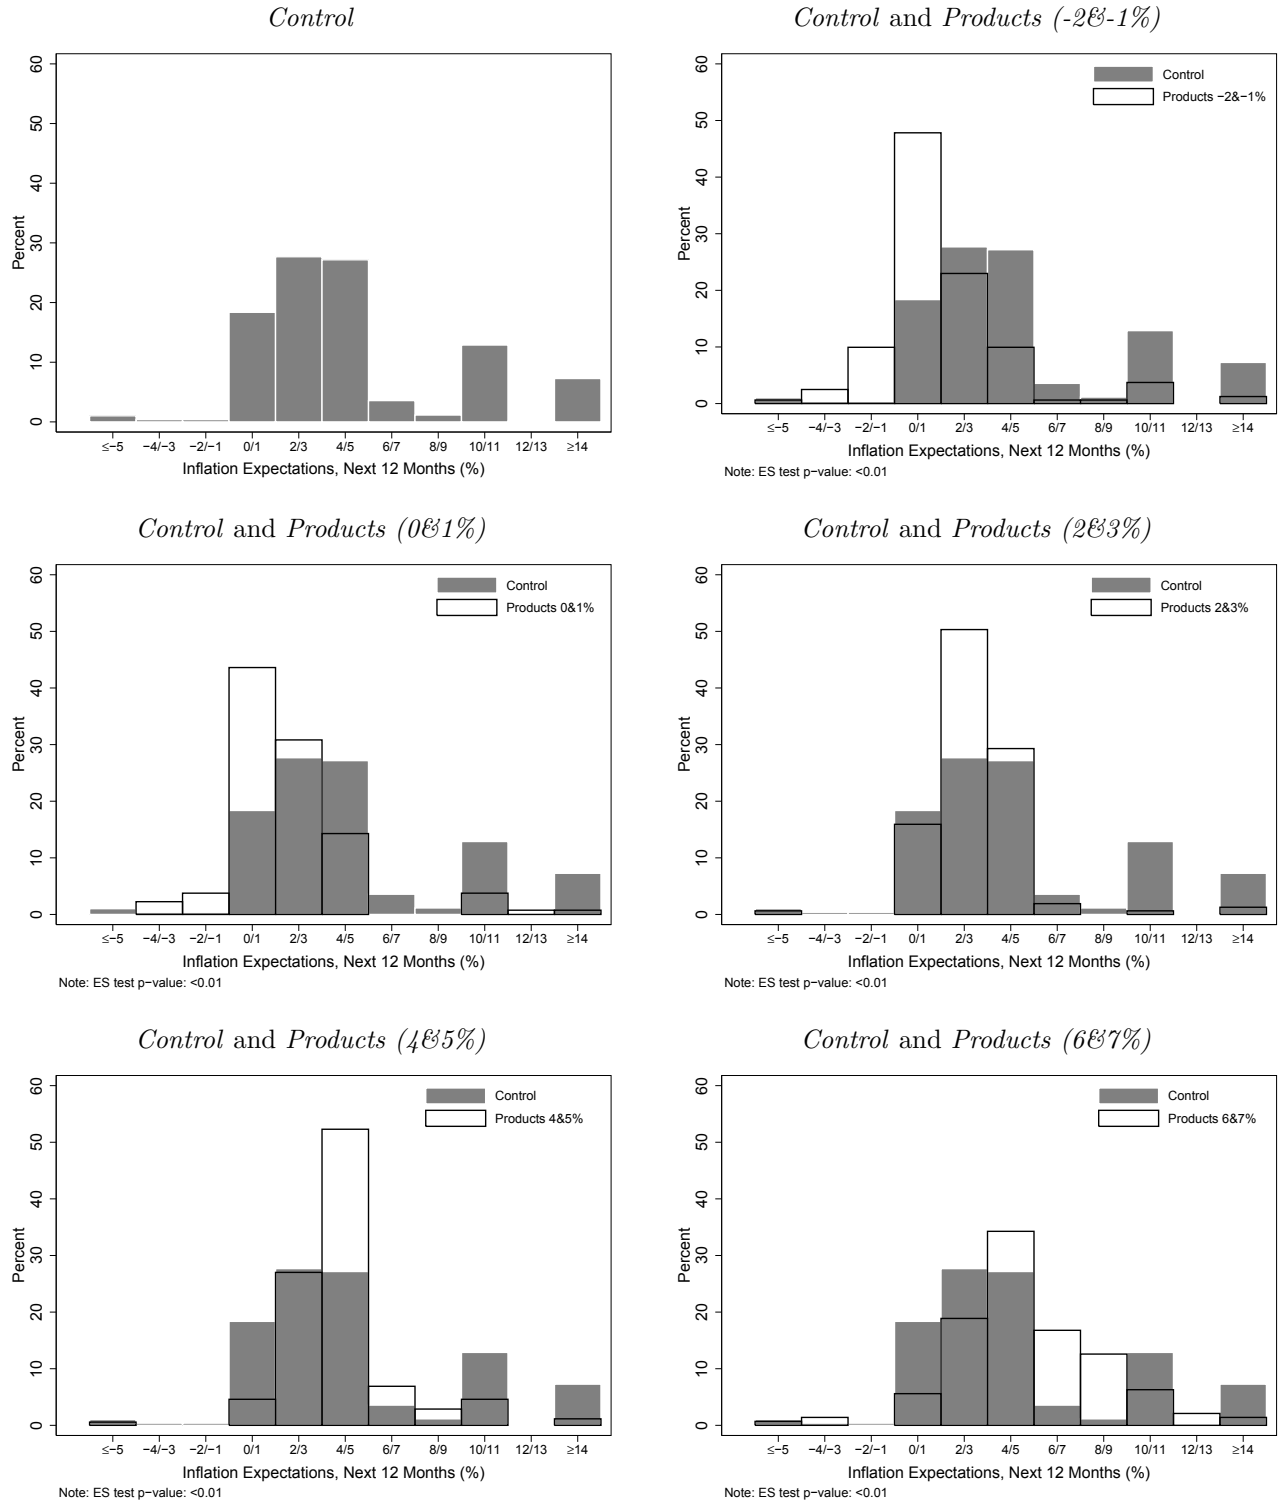
The first version, labeled as the “experimental correlation,” uses the learning equation (4) as the first stage for $\pi_{i,t+1}$ in a 2SLS estimation of (7).⁶⁶ Intuitively, this “experimental correlation” provides a measure of how much the outcome $y_{i,t+1}$ changes for every 1 percentage point increase in $\pi_{i,t+1}$ due to provision of information. Ideally, we would like to compare this experimental correlation to the true causal effect of inflation expectations on $y_{i,t+1}$ (i.e., the true δ). We denote the “non-experimental correlation” to the OLS estimate of δ from equation (7) based on subjects in the control group. Even though this non-experimental correlation may be biased with respect to the true δ because of the potential omitted variable biases discussed above, there comparison of the two correlations (the two estimates of δ) can still be informative. If the non-experimental correlations were significantly different from zero for most outcome variables but the experimental correlations were always zero, this would be a strong indication that the learning from the treatments is spurious. This would provide a qualitative rather than a quantitative test of spurious vs. genuine learning.

Figure C.5 presents these correlations for a series of additional standardized outcomes.⁶⁷ All the outcomes were constructed such as the expected correlation with inflation is positive (e.g., higher inflation should be correlated to higher interest rate). To increase the statistical power of these regressions, we pooled the three factual information treatments – the experimental correlations are statistically the same for these three treatments (see the Appendix for an illustration with the nominal interest rate). As expected, the observational correlations for the outcomes presented in Figure C.5 are all positive and significant at standard confidence levels. The experimental correlations are also positive in general, suggesting that a substantial portion of the learning was genuine. The experimental correlations, however, are lower – on absolute value – than the observational correlations. This is probably due to a combination of two factors: i. Some spurious learning; ii. Omitted-variable biases in the observational correlations. The results are thus consistent with the result presented in body of the paper that there is some spurious learning but a majority of the learning is still genuine.

⁶⁶In a 2SLS context, this corresponds to a first stage $\pi_{i,t+1} = \gamma_1\pi_{i,t}^0 + \gamma_2(\pi_{i,t}^T - \pi_{i,t}^0)$ which provides the estimated $\hat{\pi}_{i,t+1}$ to be used in the second stage $Y_i = \alpha + \delta\hat{\pi}_{i,t+1} + \varepsilon_i$.

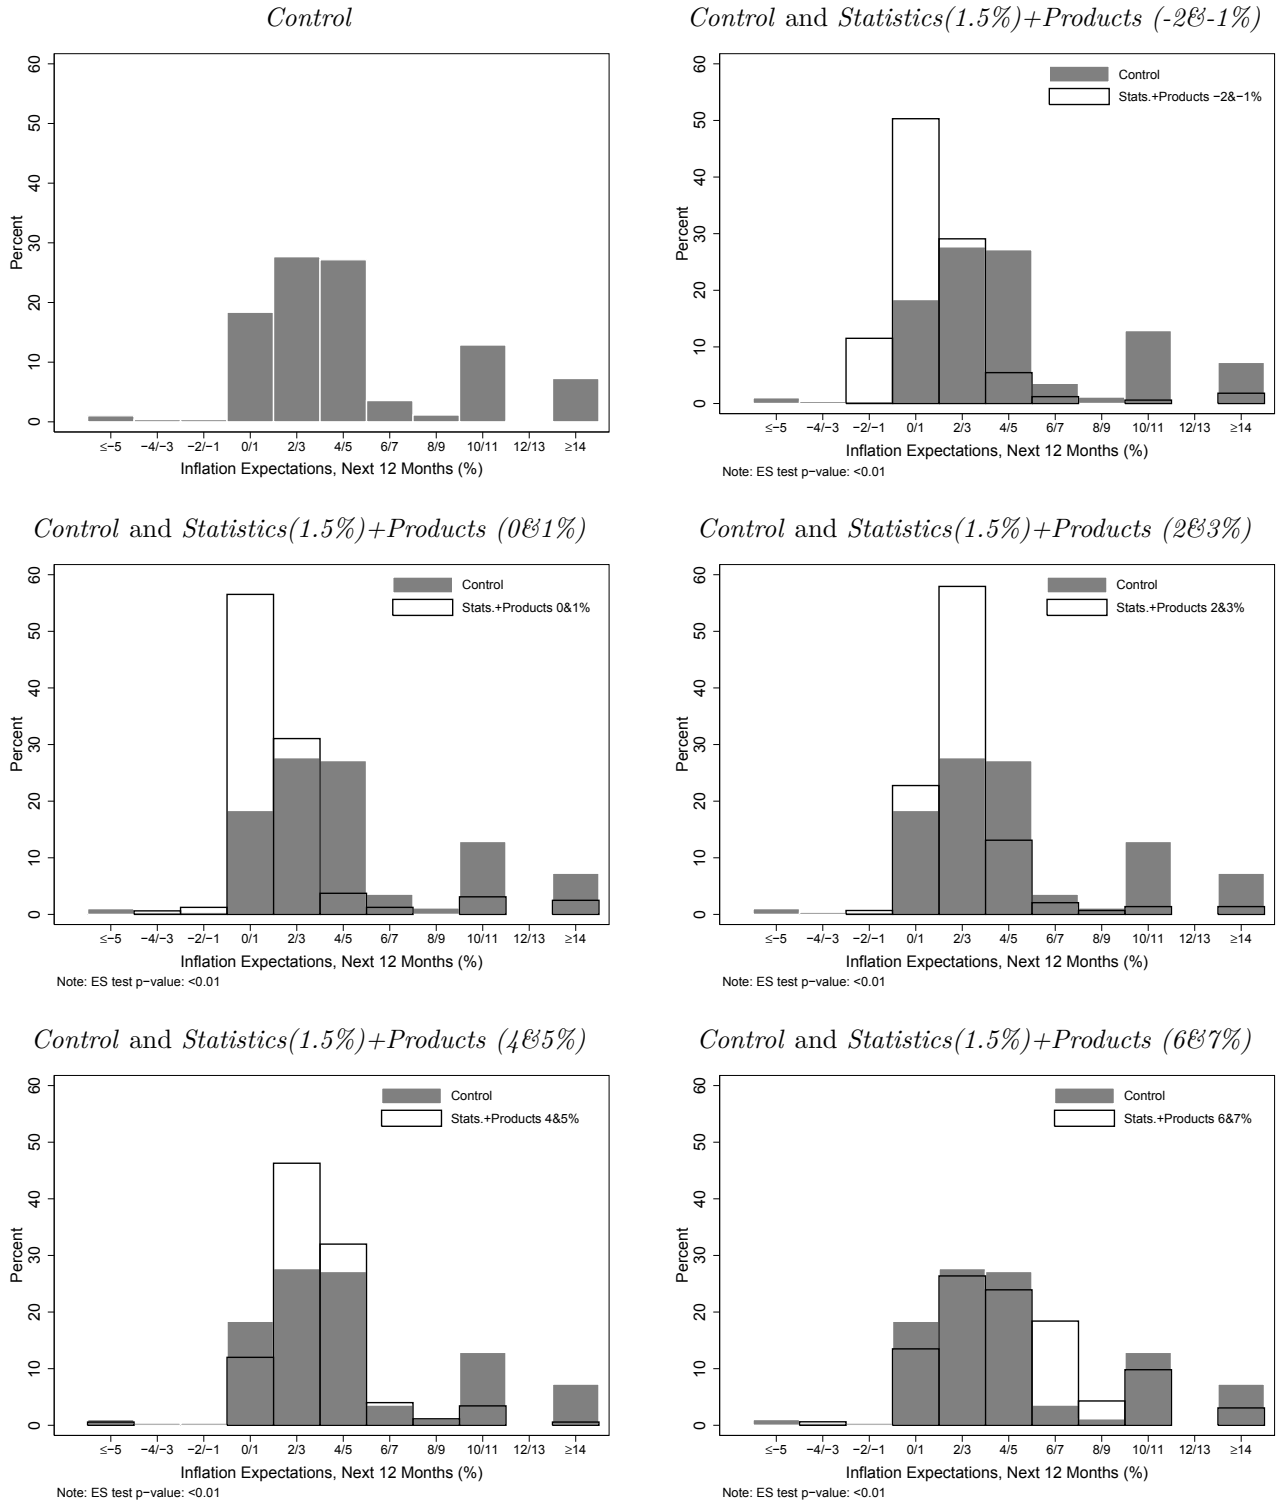
⁶⁷The categorical dependent variables presented in Figure C.5 (all but the nominal interest rate, the propensity to consume and the perceived interest rate) were rescaled according to the Probability-OLS procedure described in Van Praag and Ferrer-i-Carbonell (2007). All variables were then standardized.

Figure C.1: Inflation Expectations by Level of *Products* Treatment, *Products* Treatment Group, U.S. Online Experiment



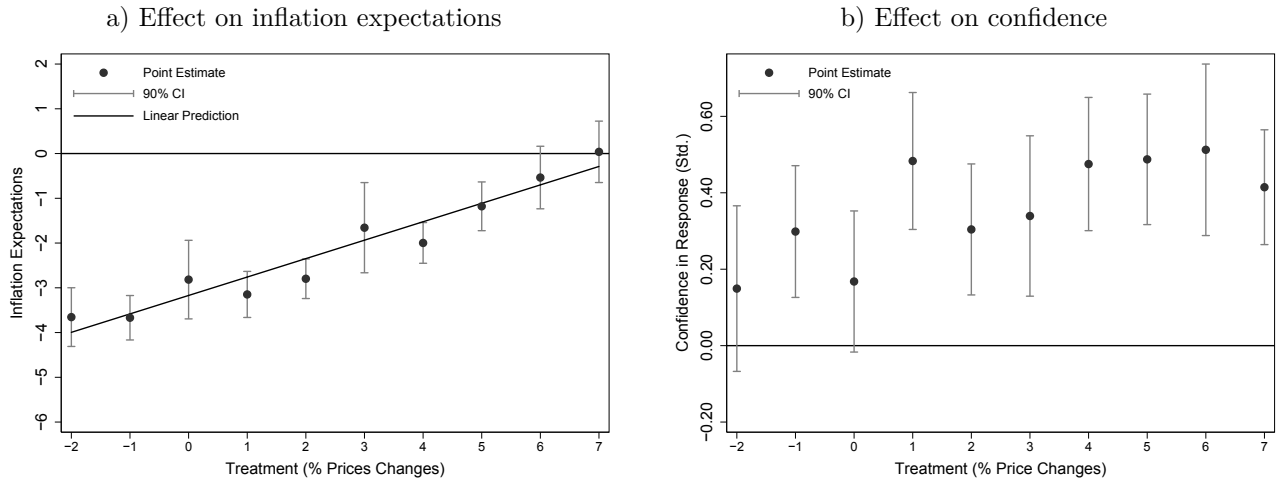
Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. ES is the Epps-Singleton characteristic function test of equality of two distributions.

Figure C.2: Inflation Expectations by Level of *Products* Treatment, *Statistics (1.5%)+Products* Treatment Group, U.S. Online Experiment




Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. ES is the Epps-Singleton characteristic function test of equality of two distributions.

Figure C.3: Treatment Effects on Inflation Expectations and Confidence about Own Expectations by Levels of *Products* Treatment, *Statistics (1.5%)+Products* Treatment Group, U.S. Online Experiment




Notes: The total number of observations is 1,732 (789 in the control group and 804 in the 10 variations of the combined specific prices and official statistics treatment). Each bar represents the point estimate of the effect of the specific price treatment compared to the control group. Robust standard errors reported.

Figure C.4: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), *Products* Treatment and *Control* Groups, by Individual Characteristics, U.S. Online Experiment



Notes: Results for the specific prices-tables treatment only. The total number of observations is 1,552 (789 in the control group and 763 in the 10 variations of the *Products* treatment). Robust standard errors reported.

Figure C.5: Observational and Experimental Correlations between Inflation Expectations and Other Economic Variables, U.S. Online Experiment



Notes: The total number of observations is 3,157 (control group and all treatments except *Hypothetical* (10%)). The observational correlations correspond to the coefficient of inflation expectations in OLS regressions of the dependent variables on inflation expectations for the *Control* group. The experimental correlations correspond to IV versions of the same models, with inflation expectations instrumented by the learning equation based on our informational treatments. The IV regressions pool the results from the three different experiments by allowing for differential levels of learning in the first stage (see Table 1). Robust standard errors reported.

D Argentina Online Experiment: Further Details and Results

D.1 Samples

The Argentina online experiment results are drawn from two different sets of respondents. The first group is comprised by a sample of economics, accountancy, business and political science graduates. This sample, with a total of 691 observations, was assigned to a control group, or to *Statistics* (24%)⁶⁸ and *Products* treatment arms, the latter with three sub-treatments with tables with average price changes of 19%, 24% and 29%. This experiment was implemented between May and June 2013 using only graduates in economics, management, accountancy, finance, international relations and political science from Argentina. We approached these subjects through mailings of graduates from the Universidad Nacional de La Plata (UNLP), Universidad Torcuato Di Tella (UTDT), and through a professional association, the Consejo de Profesionales en Ciencias Económicas of the Buenos Aires province (CPBA). About half of the individuals contacted responded to the survey resulting in a total sample of 691 respondents. Of those, 277 were accountants, 135 had a BA or MA in Economics, 89 a BA in Management, 57 an MBA or an MA in Finance, and the rest were Political Scientist and Bachelors in International Relations. All of these individuals had at least basic Economics training as part of their degrees.

The second, larger sample is based on an established public opinion research firm which carries out a quarterly online survey of adults in Argentina with the same set of basic questions since 2011. In this sample, we concentrated our efforts on a detailed version of the previously described *Products* treatment. The total of 3,653 respondents were randomly assigned to a control group (N=567) or to the *Products* treatment (N=3,086), with respondents in the latter group random assigned to one of nineteen *Products* sub-treatments with average price change in the tables of products provided ranging from 16% to 34% in one percentage point increments. Results from this periodic study are routinely used by politicians and companies. The firm relies on a stable group of respondents that participate regularly on their studies. These participants were recruited through social networking sites, and while they are not remunerated, they enter a draw for prizes, usually small household appliances. The survey has a fairly detailed questionnaire on economic and political views. We included our questions (and treatments) at the beginning of the questionnaires to minimize the attrition of respondents and also so the respondents would be more attentive when answering these questions.

Table B.1 presents some basic descriptive statistics for the main Argentina sample. This sample is not representative of the Argentine general population: while it is roughly similar in terms of

⁶⁸The value we provided for the *Statistics* treatment arm corresponds to (and was reported in the treatment as) the average of inflation estimates from private consultancies, research centers, and state-level statistical agencies, compiled and computed by opposition parties in Congress since the intervention of the national statistical agency in Argentina in 2012. See Cavallo (2013) and Cavallo, Cruces and Pérez-Truglia (2014) for more details.

age and gender composition, our sample is substantially more educated (and therefore richer) than average. This is an expected outcome from a voluntary online survey. If anything, we should expect this sample to be more informed about inflation than the average Argentine citizen.

D.2 Construction of the Informational Treatments

As in the U.S. experiment, our *Products* information provision setup consisted of displaying tables with the prices and price changes of specific products after eliciting past inflation perceptions and right before asking about respondents' inflation expectations. In the context of the Argentine experiment, we displayed a series of 19 different tables with four products each, with average price changes over the previous year (March 1 2012 to March 1 2013) ranging from 16 to 34% in one percentage point increments (see Appendix E.3 for the screen captures of the full questionnaire and for all the specific product tables).⁶⁹ The source for these tables is a database of scrapped online data from the largest supermarket chain in Argentina, and the products correspond to a subsample of four common products: olive oil, pasta, wine, and shampoos/conditioners. As in the U.S. experiment, no suggestion was made that the prices or the price changes shown in the table were representative, and that there was no deception. The text only stated that the products were selected randomly, without specifying any details about the sampling procedure.

Our information provision setup consisted of displaying tables with the prices and price changes of specific products. In the context of the Argentine experiment, in addition to the control group we displayed a series of 19 different tables with four products each, with average price changes over the previous year (March 1 2012 to March 1 2013) ranging from 16 to 34% in one percentage point increments (see two examples translated to English in Figure D.1). To construct these tables, we used a database of scrapped online data from the largest supermarket chain in Argentina. The products correspond to a subsample of four common products: olive oil, pasta, wine, and shampoos/conditioners. The tables were constructed by an algorithm to select variations of one of each product categories (e.g., Malbec wine instead of Cabernet) to obtain tables with different average levels of price changes over the preceding year. We refrained from reporting the brand names of each product because we did not want the public opinion firm to be associated with negative publicity to a particular brand. We still informed respondents that all products corresponded to well-known brands. We also attempted to hold other characteristics of the tables constant as much as possible without being deceptive (i.e., without just providing false information about products and/or their prices). With this objective in mind, the algorithm also selected products with similar initial prices within each categories. For example, consider the two olive oils in the tables with 16% and 30% average annual price changes (Figure D.1). The descriptions are identical, the initial prices are very similar, but the price changes are very different: the brand in the 30% table increased its price substantially more than the brand in the 16% table. The 750ml bottles of wine in the two

⁶⁹See two examples of these tables translated to English in Figure D.1. The accompanying text in the Appendix provides more details on the construction of these tables.

tables also have a similar initial price, but the price increase of the Malbec in the 30% table was much larger than that of the Syrah. The tables were introduced with the following text: “Before replying, please take a look at the following table. For each of the listed products, the table presents the price on March 1, 2012 and March 1, 2013 (that is, one year later). These prices were taken from the same branch from the main supermarket chain in Argentina”. It should be noted that no suggestion was made that the prices or the price changes shown in the table were representative, and that there was no deception. The text only stated that the products were selected randomly, without specifying any details about the sampling procedure.

We implemented a shorter version of the questionnaire-experiment for the sample of college graduates (see Appendix E.3 for the screen captures of the full questionnaire). The experiment had the same structure as the previous ones, and a subset of the outcomes from the larger sample Argentina experiment described above. In terms of treatments, we included three tables with specific prices (with the same format as in Figure D.1, but with dates updated accordingly – see Appendix E.3 for all the original tables included in the experiment), with average price changes of 19%, 24% and 29%. We also included a fourth treatment branch, where instead of a table, we included the following statement: “According to an average of non-official indicators produced by private firms, analysts and research centers, the annual rate of inflation with respect to the last 12 months was approximately 24%”.⁷⁰

D.3 Further Analysis

Figure 5 in the body of the paper presents the results for the online experiment for the opinion poll sample for a subset of the *Products* treatment levels. Figure D.2 presents a more detailed analysis by treatment level – lower values of average price changes in the informational treatments shifted the distribution of inflation perceptions to the left, while higher values shifted it to the right (with respect to the control group). Notably, the main effect of the middle levels of treatments (price changes between 22 and 26%) reduced the dispersion of expectations more than they affected the mean.

The Argentina opinion poll sample also allowed for a more detailed analysis of heterogeneous effects in learning. The coefficients of the learning model in Table 3 may also have different parameter values for different groups. Figure D.3 presents some evidence for differences in α between relevant groups in the population. The first two columns in the Figure present the coefficients for those with high and low levels of confidence in their inflation perceptions. In contrast to the results for the U.S., we find significant differences between the two groups: individuals who reported lower levels of confidence on their own perceptions of inflation placed a significantly higher weight on the

⁷⁰Because the government started prosecuting private sector firms and consumer associations that computed their own measures of inflation as an alternative to the adulterated official statistics, members of Parliament (who had immunity from prosecution) started compiling these private sector estimates confidentially and reported the mean of these estimates every month as the “IPC Congreso”. Our survey coincided with the April 2013 release of this indicator, with an annual inflation rate of 23.67%.

information we provided (about 0.61 compared to about 0.41). There are also similar and significant differences by education level and by age: respondents with less than a college degree and those under 40 years old place a higher weight on the information provided as part of the treatment. Females (with respect to males) also seem to learn more from the informational treatments, although this difference is not statistically significant.

Finally, as in the U.S. online experiment, we included a series of questions about other related outcomes, and we can test whether the experiment had a genuine effect on inflation expectations by comparing the observational and experimental correlations between these outcomes and inflation expectations (see section C.3 for more methodological details). These results for the main sample are summarized in Figure D.4. The results are very similar to those found in the U.S. online sample. Thus, the results are consistent with the finding reported in the body of the paper that there is some spurious learning but still a majority of the learning is genuine.

Figure D.1: Example of *Products* Treatment (Translated), Argentina Online Experiment

a) *Products* (16%)

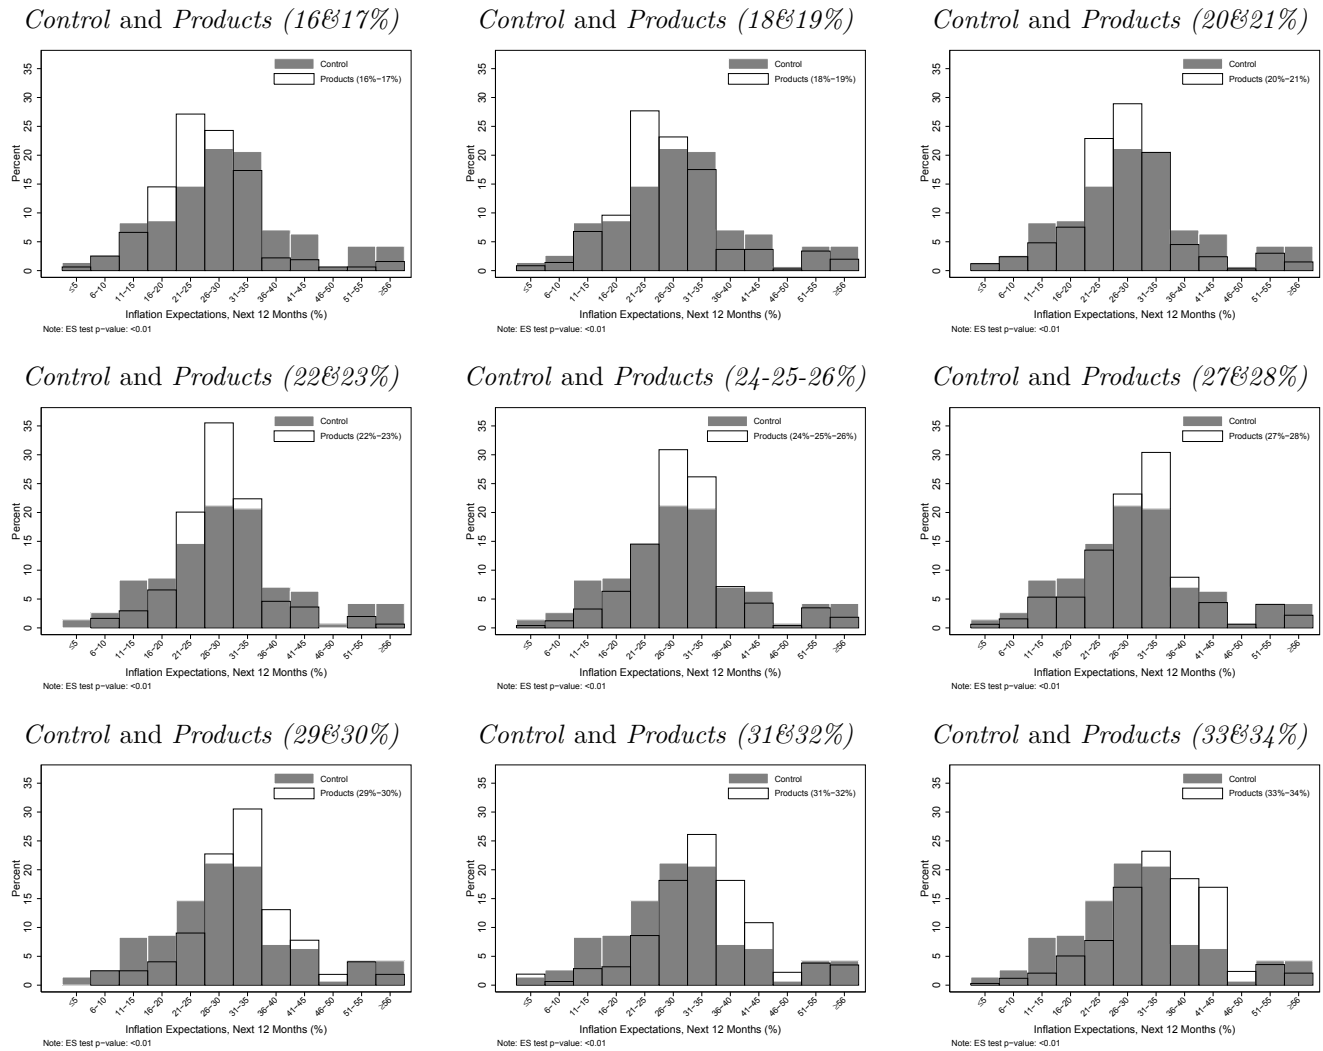
Product	Price on March-1-2012	Price on March-1-2013	Increase in %
Extra virgin olive oil 500ml	\$28 ⁸⁹	\$33 ¹⁷	14.8%
Stew noodles 500gr	\$6 ⁰⁹	\$6 ⁹⁹	14.8%
Syrah wine bottle 750ml	\$43 ⁸⁷	\$51 ²⁵	16.8%
Shampoo extra soft hipoalargenic 350ml	\$29 ³⁷	\$34 ⁵⁵	17.6%
Average increase		16.0%	

b) *Products* (30%)

Product	Price on March-1-2012	Price on March-1-2013	Increase in %
Extra virgin olive oil 500ml	\$29 ³³	\$37 ⁴⁵	27.7%
Spaghetti noodles 500gr	\$6 ⁵³	\$8 ²⁹	27.0%
Malbec wine bottle 750ml	\$42 ⁷⁹	\$56 ⁷³	32.5%
Shampoo anti age 400ml	\$29 ⁸⁰	\$39 ⁵⁹	32.9%
Average increase:		30.0%	


Notes: Prices obtained from online scrapped supermarket prices, from on of Argentina's largest supermarket chains.

Figure D.2: Inflation Expectations, *Control Group* and *Products Treatment Levels*, Argentina On-line Experiment



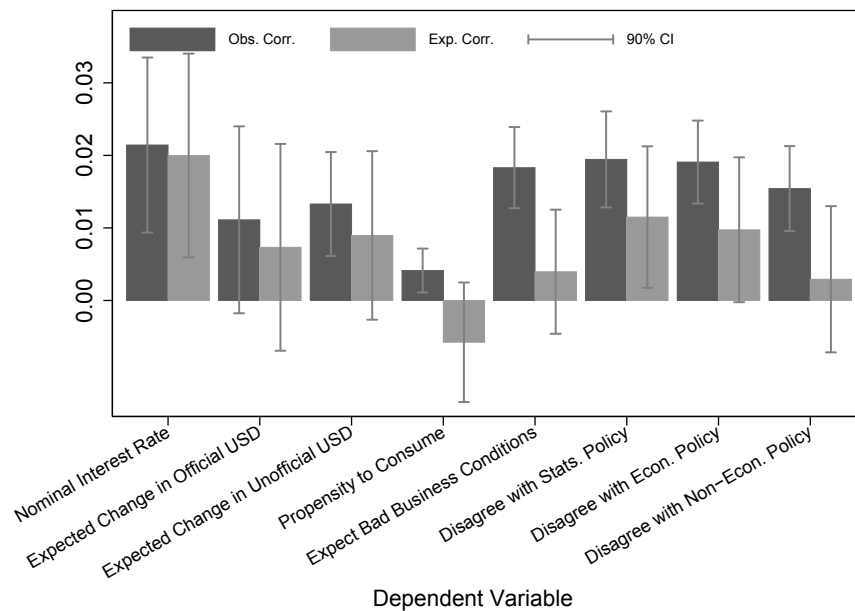
Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. ES is the Epps-Singleton characteristic function test of equality of two distributions.

Figure D.3: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), *Products* Treatment and *Control* Groups, by Individual Characteristics, Argentina Online Experiment



Notes: The total number of observations is 3,686. Robust standard errors reported.

Figure D.4: Observational and Experimental Correlations between Inflation Expectations and Other Economic Variables, Argentina Online Experiment



Notes: The total number of observations is 3,686. The observational correlations correspond to the coefficient of inflation expectations in OLS regressions of the dependent variables on inflation expectations for the *Control* group. The experimental correlations correspond to IV versions of the same models, with inflation expectations instrumented by the learning equation based on our informational treatments. Robust standard errors reported.

E Argentine Supermarket Experiment

E.1 Data Collection

The survey was carried out in June 2013 in four branches of one of Argentina’s largest supermarket chains located in the city of Buenos Aires. The subject pool were customers of the supermarket that had just made a purchase, who were invited to participate in a short survey for an academic study. About half of the individuals approached accepted to participate in the survey, and the subjects were interviewed for about 3 to 5 minutes.

The enumerators carried a handheld scanner, with which they scanned the respondents’ receipt from the supermarket purchase. These receipts contained product identifiers which could be matched to our database of scrapped online data of supermarket prices for the same chain where the study was conducted. After providing their purchase receipt for scanning, the respondents were asked 12 questions. Following our experimental design, we measure the prior belief by asking the individual about his or her perceptions of the rate of inflation over the past year. This question was followed by some randomized treatments, and then a final question about inflation expectations.

The following is an extract from the enumerators instruction manuals, translated from Spanish. Verbal statement to engage interviewees: “Hi, we are from the Universidad Nacional de La Plata. Are you willing to participate in a study on economic expectations? It will only take 5 minutes”. *To those who accept, please explain the following:* “This study attempts to relate individual shopping patterns with their economic perceptions. For this purpose, we need you to let us scan your shopping receipt. This information, the list of products, will allow us to develop the empirical analysis for our study. The receipt does not contain your name nor any sensitive information. The survey is completely anonymous. Once that we scan your receipt, we only need you to answer a brief survey that will take between 3 and 5 minutes. You can finish your participation in this study at any time.” The scanned tickets did not have identifying information (credit card receipts are processed separately and they were not scanned as part of this study). The enumerators reported high levels of interest and curiosity from the respondents, especially about the use of the handheld scanners. Appendix E.3 presents the original survey instrument, the three specific product tables, and the enumerators instruction manual.

E.2 Robustness Checks with Total Purchase Amounts Instead of Specific Product Prices

Figures E.1 and E.2 present robustness checks of the results in the main body of the paper. The previous results were based on actual and remembered price changes for products the respondents had just purchased. The survey, however, also recorded the total amount spent, and asked the respondents about their estimate of the total they would have had to pay for the same goods 12 months earlier. The results presented in this Appendix are not based on these remembered price

changes. Instead, they compare the distribution of inflation expectations (Figure E.1) for individuals for high and low remembered and actual changes in their purchase receipts total amount. Figure E.2 in turn depicts the relationship between the price changes in the receipt and inflation expectations (panel (a)), as well as the relationship between price changes in the receipt (actual and remembered).

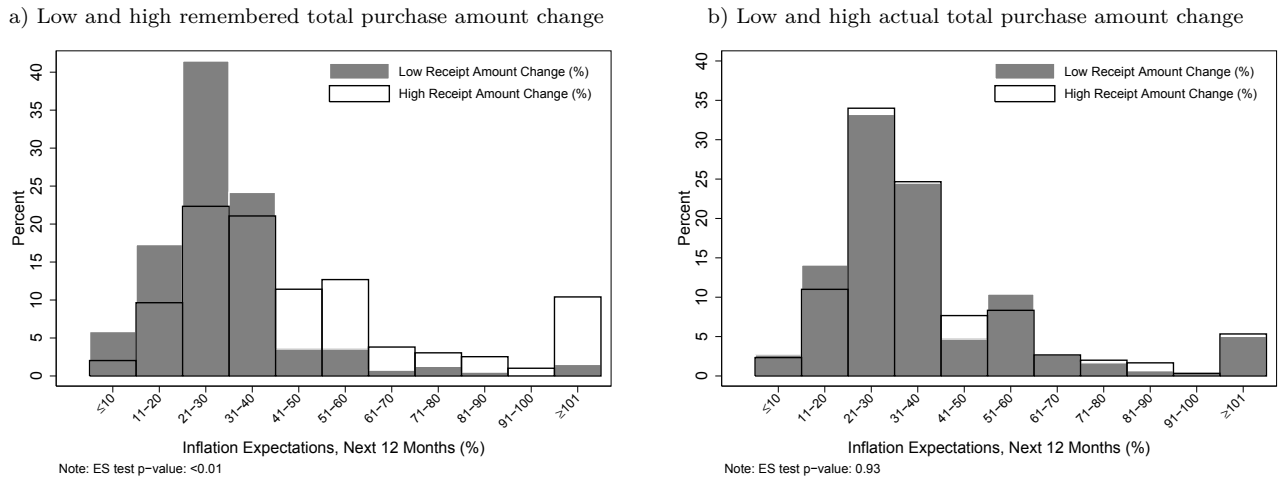
E.3 Estimating Learning Rates

The rate of learning from remembered price changes of specific products can also be depicted by means of the Bayesian learning model used before. However, we must note that, in contrast to the other informational treatments, we did not randomize the remembered price changes directly, but instead we randomized the salience for a group of products. As a result, we cannot compare the α from randomizing salience than from randomizing the information directly. Because individuals know this information and would have probably incorporated it in their inflation expectations even if we did not made it salient, the estimated α is expected to be much lower. Furthermore, we must keep in mind that in this supermarket experiment subjects were provided simultaneously with multiple pieces of information and on the spot, so we should not expect them to have as much time or interest in processing the information. For example, the table with price changes was shown to the subject for just a few seconds in a context of a street face to face survey, while in the online experiment individuals spent a median of about 40 seconds inspecting the information on the table (U.S. online experiment). Moreover, since we asked so many numerical questions, it is possible that individuals had a cognitive overload or a depleted memory for numbers. Because of these reasons, we should not expect learning rates to be as high as in the online experiments.

Table E.1 presents the estimates from the learning model described in section 2.3 for our supermarket study. The first randomly assigned information for which we compute the learning model is the average remembered price change for the four products that the respondent was asked about.⁷¹ The α coefficient is about 0.11 and strongly significant. This weight is substantially lower than the one obtained from the online experiments (about 0.5 for Argentina), but this was expected due to the reasons listed above due to the reasons listed above. This implies that individuals form their inflation expectations, in part, based on information that is mostly noise (i.e., it is not correlated with actual price changes – see Figure 7, panels (c) and (d)), as we established previously. To stress this point, in column (2), instead of using remembered price changes, we use the actual price changes in the list of randomly selected products. As expected, the estimated α is close to zero and statistically insignificant. In column (3), we present the estimates from the replication of the *Products* treatment with the three levels discussed in the previous paragraph. The α coefficient, which represents the weight given by respondents to the price information we provided, is similar in value to the α for (salient) remembered prices (although it is statistically insignificant). The last column (4) in the table pools all these alternative treatments, and the results are very similar.

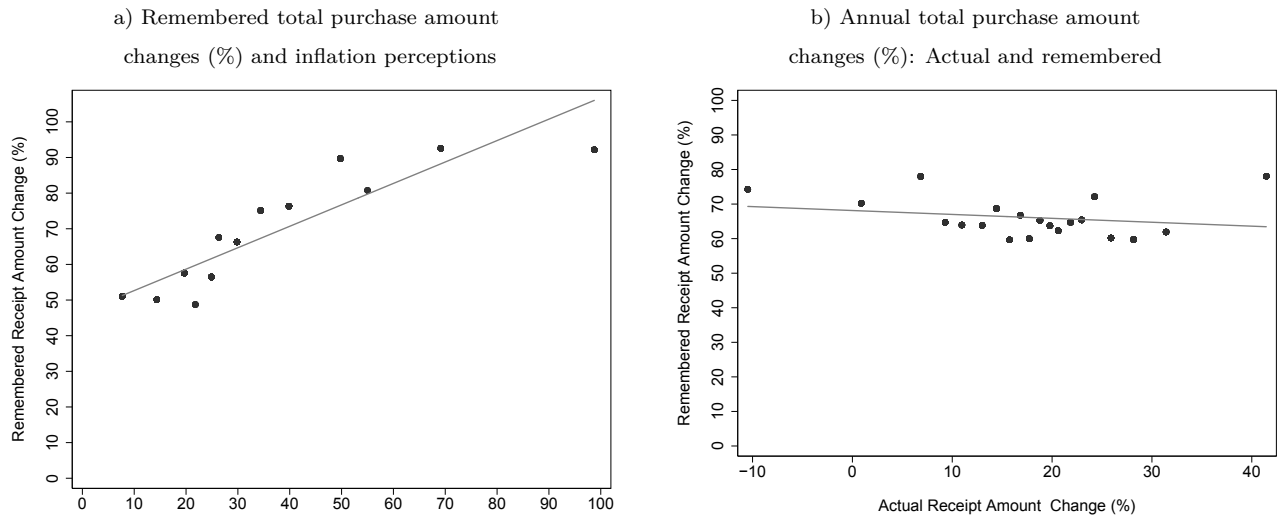
⁷¹Given the biases documented above in terms of the average price changes reported by respondents, we use here a “corrected” value using a deflation factor of 30%. The results are similar under alternative specifications.

Figure E.1: Inflation Expectations by Total Purchase Amount Changes, Argentina Supermarket Experiment



Notes: The total number of observations is 375 (lowest third of total purchased amount changes, panel (a)) and 372 (top third of total purchased amount changes, panel (b)). ES is the Epps–Singleton characteristic function test of equality of two distributions.

Figure E.2: Robustness: Implicit Price Changes from Total Purchase Amount and Inflation Expectations, Supermarket Experiment, Argentina



Notes: The total number of observations is 1,140. Panels (a) and (b) represent binned scatterplots, where the number of observations are almost identical across bins. The percentage changes in both panels are implicit – they are obtained from the total purchase amounts in pesos (AR\$) from the scanned receipt and from the estimate of the total for the same purchase a 12 months earlier as reported by the respondents.

Table E.1: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), Argentina Supermarket Experiment

	(1)	(2)	(3)	(4)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$
β	0.923*** (0.085)	0.794*** (0.084)	0.958*** (0.152)	1.005*** (0.157)
<i>Remembered Price Changes</i>				
α	0.115*** (0.035)			0.105*** (0.037)
<i>Actual Price Changes</i>				
α		-0.050 (0.053)		-0.041 (0.041)
<i>Products</i>				
α			0.130 (0.133)	0.124 (0.129)
Observations	1,070	1,070	1,070	1,070

Notes: The total number of observations correspond to 1,070 participants of the Argentina Supermarket Experiment with valid responses for inflation expectations and remembered price changes, and for which it was possible to establish the actual price changes from the scanned purchase receipts (actual price changes). The α and β coefficients are obtained from the regression given by Equation 4, section 2.3. The p-value of the difference between the α coefficient for *Remembered Price Changes* and *Actual Prices Changes* is 0.0102. Robust standard errors. *significant at the 10% level, ** at the 5% level, *** at the 1% level.