NBER WORKING PAPER SERIES

# INFLATION EXPECTATIONS, LEARNING AND SUPERMARKET PRICES: EVIDENCE FROM FIELD EXPERIMENTS

Alberto Cavallo
Guillermo Cruces
Ricardo Perez-Truglia

Working Paper 20576
http://www.nber.org/papers/w20576

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2014

We would like to thank Robert Barro, Raj Chetty and David Laibson for their valuable input at the early stages of the project; Alberto Alesina, Nageeb Ali, Rüdiger Bachmann, Christian Borgs, Sebastian Di Tella, Emmanuel Farhi, Matthew Gentzkow, N. Gregory Mankiw, Markus Mobius, Andrés Neumeyer, Roberto Rigobon, Martin Tetaz, Glen Weyl, Fernando Yu, and participants in the seminars at Harvard University, the Universidad de San Andres, the Universidad Torcuato Di Tella, Microsoft Research and the Chicago/NYU International Macro Finance Conference for their valuable comments. Julián Amendolaggine 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 field 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w20576.ack

Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments
Alberto Cavallo, Guillermo Cruces, and Ricardo Perez-Truglia
NBER Working Paper No. 20576
October 2014
JEL No. C93,D83,E31,E58

## ABSTRACT

Inflation expectations in household surveys tend to be vastly heterogeneous. The literature has been unable to distinguish empirically between alternative explanations, such as the existence of rational inattention (according to which individuals will not continuously gather costly information) and the use of information from personal shopping experiences (which can be diverse and inaccurate). To better understand the importance of these mechanisms, we use evidence from field experiments with nearly 10,000 subjects conducted in both a low-inflation country (the United States) and a high-inflation country (Argentina). We introduce a novel experimental design which, when combined with unique data sources, allows us to quantify how much weight individuals assign to each source of information about inflation. Our novel experimental framework addresses one of the most important concerns with survey experiments by separating genuine from spurious learning. We find that individuals are highly influenced by information on both inflation statistics and price changes of specific products. The results are consistent with rational inattention, since there is greater learning in a low-inflation setting where the stakes are lower (the United States), and also from information that is less costly to understand (individual supermarket prices). To further assess the importance of personal experiences, we conducted field experiments which combined data from actual products purchased by the subjects with their historical prices. We find that individuals form inflation expectations using their own memories about the price changes of supermarket products they buy, but those memories are nearly orthogonal to the actual price changes.

Alberto Cavallo
MIT Sloan School of Management
100 Main Street, E62-512
Cambridge, MA 02142
and NBER
acavallo@mit.edu

Guillermo Cruces
CEDLAS
Univesidad Nacional de La Plata
Calle 6 entre 47 y 48, 5to. piso, oficina 516
(1900) La Plata
Argentina
gcruces@cedlas.org

Ricardo Perez-Truglia
Microsoft NERD Lab Office 12073
1 Memorial Drive
Cambridge MA 02142
rtruglia@microsoft.com

A data appendix is available at:
http://www.nber.org/data-appendix/w20576

# 1 Introduction

Expectations about macroeconomic variables play an important role in economic theory and policy-making. Household inflation expectations, in particular, are key to understand consumption and investment decisions, and ultimately, the impact of monetary policies. Although Central Banks have a natural desire to influence expectations, there is little empirical evidence on how household expectations are formed and what is the best way to influence them (See Bernanke, 2007; Bachmann et al., 2012; Coibion and Gorodnichenko, 2013).

Inflation expectations in household surveys tend to be much more heterogeneous than those of professional forecasters (Ranyard et al., 2008; Armantier et al., 2013). Two main explanations have been given in the literature for this degree of dispersion. Some authors attribute it to rational inattention, according to which individuals only partly incorporate information on topics such as inflation statistics because acquiring that information is costly. 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). These explanations are hard to distinguish empirically because they are not mutually exclusive. Individuals may choose to be rationally inattentive and, at the same time, use their personal shopping experiences as a low-cost source of information about price changes. We present evidence from a series of experiments specifically designed to disentangle some of these effects.

Using both online and offline survey experiments, we randomly provided subjects with information related to past inflation, such as inflation statistics and the historical prices of specific supermarket products. On the basis of that experimental variation, we use a learning model to infer – from the treatment effects on the distribution of inflation expectations – how much weight subjects gave to a given piece of information relative to their prior beliefs.[1] In order to assess the role of the rational inattention model, we conducted field experiments in 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] Our experimental design addresses one of the most common criticisms of this type of survey experiments by disentangling how much of the reaction to the information reflects genuine learning rather than spurious learning. Additionally, our experiments introduce some unique features, such as exploiting scraped data on supermarket prices for

[1]Note that instead of asking subjects about past inflation twice, which naturally tends to exacerbate spurious learning, we instead elicited beliefs about past inflation, provided information related to past inflation, and then estimated the learning rate from the treatment effects on the distribution of beliefs about future inflation.

[2]We do not use official inflation statistics for Argentina, since these are widely discredited. We use instead an indicator compiled by the private sector, which is well known and trusted.

individual products,[3] re-interviewing subjects months after the information provision experiment and conducting survey experiments with shoppers that use accurate data about the products that they purchased and the corresponding historical prices.

Our results indicate that information related to past inflation has a major impact on inflation expectations. We find that, when confronted with information about past inflation that is different from their priors, individuals will assign a weight of between 50% to 80% to the new data to update their beliefs. This happens both when we provide information about aggregate inflation statistics and when we provide information about the historical prices of a few individual supermarket produces. This evidence is consistent with the existence of largely *inattentive* consumers that learn from new information. Furthermore, the results across countries suggest this inattention is *rational*. Relative to her prior belief, an individual in a low-inflation context assigns a weight of roughly 85% to the information on recent inflation statistics, whereas an individual in a high-inflation context assigns a weight of roughly 50%. The differences are similar when comparing the weights assigned to information about supermarket prices rather than inflation statistics. The fact that learning rates were 70% higher in the low-inflation context is consistent with the rational inattention model, which predicts that individuals in a context of higher inflation are more informed because the cost of misperceiving inflation is greater (Mankiw et al., 2003; Carroll, 2003).

Another treatment arm provided individuals with information on inflation statistics and simultaneously with information on historical prices for a handful of supermarket products. Subjects still assigned significant weight to the prices of specific products – even a higher weight than that assigned to inflation statistics. In other words, subjects were more prone to incorporating information about the price changes of a few familiar products, such as bread and milk, than to statistics on the price changes of thousands of products. One possible interpretation, still consistent with rational inattention, is that it is less costly for people to incorporate information on individual prices, simply because they may be easier to understand. This interpretation would imply that, even when accurate information is readily available, individuals would still prefer to use less accurate information such as their own price memories from their shopping experiences. Indeed, this interpretation is consistent with prior survey data on price memories and inflation expectations.[4] In our own surveys, we found that 64.4% of the subjects in the U.S. reported trying to recall the prices of specific products when answering about inflation expectations, twice as many as those who report trying to recall inflation statistics. Even in Argentina, where credible (not official) inflation statistics were readily available and widely reported in the media on a daily basis, 74.9% of respondents reported trying to recall prices of specific products when asked about past inflation.

To better understand how past shopping experiences affect inflation expectations, we conducted a consumer intercept survey experiment at a supermarket chain in Argentina. We recorded con-

[3]The 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.

[4]For example, according to survey evidence presented by Bruine de Bruin et al. (2011), when asked about their perceptions of inflation most individuals report that they try to recall prices of specific products.

sumer purchases by scanning the supermarket purchase receipts of participants, which were 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 set of items that they had just purchased, which allowed us to generate exogenous variations in the salience of their own price memories. We find that inflation expectations are not related to the actual inflation rate experienced by the products recently bought by the subjects. Our experimental evidence suggests that individuals try to use their own price memories in forming inflation expectations, but that these memories about price changes are nearly orthogonal to actual price changes. Far from correcting a representativeness bias in aggregate inflation statistics, the use of price memories as inputs in the formation of inflation expectations seems to induce significant errors in inflation expectations. This evidence suggests that, even after accounting for rational inattention, some of the excess heterogeneity in household inflation expectations is due to the use of (less accurate) information from personal experiences (Madeira and Zafar, 2014).

Methodologically, our paper contributes to the literature on inflation expectations by addressing a common concern in survey experiment studies: the existence of spurious learning.[5] For instance, when an individual 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: a desire to agree with the interviewer because of a desirability bias (Goffman, 1963); a fear of being deemed ignorant; an unconscious numerical anchoring (Tversky and Kahneman, 1974).[6] As part of our experimental design, we developed a methodology for measuring how much of the reaction to the information provided can be attributed to genuine learning and how much can be attributed to spurious learning. Our first strategy consisted of estimating the learning model with data from a follow-up survey conducted months after the information was first provided to the subjects. Intuitively, the importance of the spurious reaction – such as unconscious anchoring or interviewer pressure – should disappear or diminish considerably months after the information is provided (along with the salience effect of providing information that was already known by the subject). A second strategy to disentangle spurious from genuine learning is to estimate the learning model using the effect of information about inflation on expectations regarding other nominal variables. Intuitively, if a piece of information causes a subject to expect a higher inflation rate, that same information should also affect expectations about the nominal interest rate (or other related nominal variables). A third strategy included a treatment arm in which – under the pretext of a cognitive test – individuals were knowingly given information about price changes of fictitious products. While the results are suggestive of some spurious learning of between 10% to 50% of the total effect, we still find that the majority of the learning is genuine rather than spurious. Overall, these findings indicate that concerns about spurious effects in the context of

[5]This is a common criticism of survey experiments in general. Our methodology can be used to test the provision of information and survey experiments beyond this specific context.

[6]See 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.

survey experiments are justified and must be taken seriously. Indeed, the methodology that we propose for disentangling genuine learning could be applied to survey experiments on topics other than inflation expectations.

Our findings have implications for macroeconomic theory. The idea that monetary policy can have real effects due to imperfect information 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. These models can explain a number of stylized facts, such as why the price level responds slowly to monetary policy shocks even though individual prices change fairly frequently and substantially, but they are focused mostly on firms' price-setting decisions. Our findings suggest that informational frictions are also worth exploring in the case of households. An example of a study taking this approach is Coibion and Gorodnichenko (2013), which uses a Phillips curve that is augmented by incorporating household inflation expectations to explain why disinflation was not more intense during the Great Recession. Consistent with our findings, they argue that household inflation expectations rose significantly during that period because individuals based their expectations on their own experience with price changes of familiar products such as gasoline.

Our findings are also related to recent debates on central bank transparency and communication strategies. The ability of central banks to correct households' misperceptions could help households make better financial decisions (Armentier et al., 2013).[7] It is less clear whether correcting these biases in perceptions is desirable from the perspective of the monetary authority. For example, some authors argue that information disclosure is welfare-enhancing (Hellwig, 2005), while others argue the opposite (Morris and Shin, 2002). In practice, however, central banks will often try to use information to affect household inflation expectations. At the time of writing, for example, Japan's Central Bank is eagerly trying to increase household inflation expectations, while high-inflation countries such as Argentina and Venezuela are desperately trying to reduce them. Our evidence suggests that, in addition to the dissemination of aggregate statistics, central banks could also consider incorporating information about individual product's inflation rates - for instance, by constructing tables with historical and current prices for goods and services that people buy often and may find easier to understand.[8]

Our paper is related to several strands of literature. First, it is related to a group of studies

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

[8]Our survey data from the United States and Argentina shows that a substantial share of respondents do not trust inflation statistics and/or do not believe those statistics to be representative. These credibility issues are, if anything, more widespread and severe in developing economies like Argentina, where official statistics are widely discredited (e.g., Cavallo 2014). Our findings suggest that central banks could also consider communicating how objective, accurate, and representative aggregate inflation statistics are. In recent years, for example, the European Central Bank and the French statistical agency have made considerable efforts to create user-friendly online tools to explain how inflation statistics are collected and processed. See 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.

that measure the role of inflation statistics in household inflation expectations, exploiting media coverage of statistics (Lamla and Lein, 2008; Badarinza and Buchmann, 2009; Drager, 2011), the publication of official statistics (Carrillo and Emran, 2012), and information-provision experiments (Roos and Schmidt, 2012; Armantier et al., 2014). Other studies have looked at the role of personal experiences in the formation of inflation expectations. 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, 2013) and that individuals place excessive weight on information about inflation rates that occurred during their lifetime (Malmendier and Nagel, 2013). Despite substantial research, the understanding of these two mechanisms – inflation statistics and personal experiences – remains fragmentary and incomplete (Ranyard et al., 2008). Our experimental design provides a unified framework for measuring and comparing the contribution of these two factors to the formation of inflation expectations. Additionally, our paper builds upon a recent group of studies that employ survey experiments to investigate inflation expectations. Roos and Schmidt (2012) and Armantier et al. (2014), for instance, study how individuals react to information about U.S. inflation statistics by adjusting their own self-reports about inflation. Bruine de Bruin et al. (2011) in turn show that subjects who are asked to think about products with extreme price changes tend to report higher inflation expectations. While these studies provide the groundwork for our analysis, our paper makes a contribution to this literature by providing an original experimental design capable of disentangling genuine from spurious learning, by replicating the experiments in contexts of low and high inflation, and by exploiting unique sources of data (e.g., individuals' purchases, their price memories for the products they purchased, and actual historical prices).

The paper proceeds as follows. In section 2, we describe the experimental design that is used in the rest of the paper. Section 3 presents evidence from a series of online experiments conducted in the United States and Argentina. Section 4 presents evidence from a field experiment conducted in supermarkets in Argentina. The last section presents some conclusions.

# 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. A growing number of studies have relied on survey experiments to examine the formation of inflation expectations. These studies typically involve providing subjects with some information related to inflation and then comparing the elicited inflation expectations of subjects who are randomly assigned to receive the information and those who are assigned not to receive it (e.g., Bruine de Bruin et al., 2011; Roos and Schmidt, 2012; Armantier et al., 2014). Our study builds on this principle, but addresses some important problems faced by the existing experimental designs. Most importantly, we address the criticism that the information provided in the experiment may not generate genuine learning, but instead induce spurious effects on self-reported variables. Our framework provides estimators to quantify how much of the reaction to the information provided reflects genuine learning and how much reflects spurious learning.

Standard survey experiments entail asking the individual about inflation (e.g., "what was the annual inflation rate over the past twelve months?"), providing individuals with information about inflation (e.g., "according to source X, inflation was actually Y%"), and then asking the same question about inflation again in order to measure whether individuals revised their reported expectations in the direction of the signal provided (e.g., Roos and Schmidt, 2012; Armantier et al., 2014). Our framework is a variation on this principle insofar as it does not rely on asking the same question twice. The basic structure of the survey experiments consists of collecting some background information and then:

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 ($\pi^0_{i,t}$ 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 ($\pi^T_{i,t}$). In the case of the control group with no information, there is no signal. Some of the treatments provided information on average price changes from one or more sources, such as recent official inflation statistics or a table with the historical prices of specific products. Each type of information provision treatment is described in the following sections when presenting the corresponding results.

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, $i_{i,t+1}$).

The goal of the empirical model that follows is to infer how much weight individuals assign to a particular type of information (e.g., inflation statistics) from the joint distribution of $\left\{\pi^0_{i,t}, \pi^T_{i,t}, \pi_{i,t+1}\right\}$, even if more than one signal is provided simultaneously.

## 2.2 A Framework to Estimate Learning Rates

The results presented in the following sections consist mostly of reduced-form evidence on how individuals react to randomly assigned information. The main advantage of this model-free approach is its transparency. However, we need to estimate a learning model to establish the relative importance that individuals place on the different sources of information provided. The simple model developed here allows us to summarize the learning rate on the basis of a single parameter that can be compared across experimental samples and treatments.

Let $\pi_{i,t}$ 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):

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

Significantly, this model of inflation forecasting does not make any assumptions about the agent's rationality. The fact that individuals use information about the past to estimate future inflation is 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), which also seems to be consistent with the data (Atkeson and Ohanian, 2001).

On the basis of the data, in the rest of the paper we use a linear specification for $f()$: i.e., $\pi_{i,t+1} = \mu + \beta\pi_{i,t}$, where $\beta$ is the degree of pass-through from inflation perceptions to inflation expectations. The strong linear relationship between perceptions of past inflation and expectations of future inflation has been documented in previous studies (Jonung, 1981). This same relationship is present in our data, as depicted by Figure 1, which presents the relationship between perceived past inflation and expected future inflation for our online samples in the United States (panel (a)) and Argentina (panel (b)).[9] A great deal of the variation in inflation expectation can be explained by variation in inflation perception: in our U.S. sample, 29% of the variation in inflation expectation is due to variation in inflation perception, whereas the equivalent figure for our Argentine sample is 60%. These proportions would undoubtedly be even higher if they took into account the significant measurement error in the reporting of these variables. In other words, the dispersion and bias in inflation expectations appears to be mostly a mechanical product of individuals' uncertainty about past inflation (see also Blanchflower and MacCoille, 2009). There is some evidence that extrapolating on the basis of past inflation is rational: for example, Atkeson and Ohanian (2001) report that, since 1984, the one-year-ahead inflation forecast of professionals in the U.S. has been no better than the "naïve" forecast of the inflation rate over the previous year. We should not expect households to come up with more accurate predictions on future inflation than those naïve estimates. The real question, then, is not why future inflation expectations are so dispersed, but why perceptions about past inflation are so dispersed.

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

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

$$\pi_{i,t} = g\left(\pi^0_{i,t}, \pi^T_{i,t}\right) \tag{2}$$

In our setup, we have information on these three elements $\pi^0_{i,t}, \pi^T_{i,t}, \pi_{i,t+1}$: $\pi^0_{i,t}$ is the respondent's stated past inflation perception (pre-treatment), $\pi^T_{i,t}$ 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). Under certain functional form assumptions, we can estimate the parameters behind $f()$ and $g()$.

There are several plausible functional forms for $g()$. A simple and parsimonious alternative is to assume a Gaussian model. The prior belief is then normally distributed with mean $\pi^0_{i,t}$ and standard deviation $\sigma^0_{i,t}$. 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, $\pi^T_{i,t}$, which represents the price change for one product randomly drawn from the universe of products (the results would be equivalent if the price change for that randomly selected product and the average price change over multiple randomly-drawn products were the same). The population of price changes for all possible products follows a normal distribution with mean $\pi_{i,t}$ and standard deviation $\sigma^T_{i,t}$; this functional form is also roughly consistent with the actual distribution of price changes. By construction, $\pi^{TRUE}_{i,t}$ 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 $\sigma^T_{i,t}$, which is assumed to be known. Under these assumptions, the posterior belief is distributed normally with the following mean and variance:

$$\pi_{i,t} = \frac{\left(\frac{1}{\sigma^0_{i,t}}\right)^2}{\left(\frac{1}{\sigma^0_{i,t}}\right)^2 + \left(\frac{1}{\sigma^T_{i,t}}\right)^2}\pi^0_{i,t} + \frac{\left(\frac{1}{\sigma^T_{i,t}}\right)^2}{\left(\frac{1}{\sigma^0_{i,t}}\right)^2 + \left(\frac{1}{\sigma^T_{i,t}}\right)^2}\pi^T_{i,t},\ \sigma_{i,t} = \sqrt{\frac{\left(\sigma^0_{i,t} \cdot \sigma^T_{i,t}\right)^2}{\left(\sigma^0_{i,t}\right)^2 + \left(\sigma^T_{i,t}\right)^2}}$$

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

$$\pi_{i,t} = (1 - \alpha_{i,t})\pi^0_{i,t} + \alpha_{i,t}\pi^T_{i,t} \tag{3}$$

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

$$\pi_{i,t+1} = \gamma_0 + \underbrace{\gamma_1}_{\beta}\pi^0_{i,t} + \underbrace{\gamma_2}_{\alpha\beta}\left(\pi^T_{i,t} - \pi^0_{i,t}\right) \tag{4}$$

Since $\pi_{i,t+1}$, $\pi^0_{i,t}$ and $\pi^T_{i,t} - \pi^0_{i,t}$ are all observed in our experimental data, we can estimate $\hat{\alpha}$ and $\hat{\beta}$ by simply running the above linear regression.[10] The parameter $\beta$ represents the rate of pass-through from perceptions of past inflation to future inflation expectations. The parameter $\alpha$ 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 $\pi^0_{i,t}$ and the informational treatment provides a signal that inflation is $\pi^T_{i,t}$, the posterior belief can be expected to be between $\pi^0_{i,t}$ and $\pi^T_{i,t}$, and the parameter $\alpha$ reflects how much closer $\pi_{i,t}$ is to $\pi^T_{i,t}$ relative to $\pi^0_{i,t}$.

As in Armentier et al. (2013), we could have asked about inflation, then provided information and asked about inflation again to see whether the information made individuals revise their self-reports. However, this setup has the limitation that individuals may feel pressured and/or primed to revise their expectations. To avoid asking the same question twice, we rely instead on the relationship between perceptions of past inflation and expectations of future inflation. The following example illustrates the intuition behind our identification strategy. 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%=0.8×10% + 0.2×20%.

This model of Bayesian learning makes a number of additional predictions that can be directly tested with the data.[11] It 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 ($\sigma^T_{i,t}$), the effect of providing a signal on $\sigma^T_{i,t}$ should be independent of the particular value of the signal that was drawn ($\pi^T_{i,t}$). Also, $\alpha$ can differ according to the individual's reported degree of confidence in her perceptions of past inflation, $\sigma^0_{i,T}$. The Bayesian learning model predicts that individuals with a stronger prior belief should have a lower $\alpha$ value. The 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

[10] One assumption is that the above OLS regression yields an unbiased estimate for $\beta$. Since $\pi^0_{i,t}$ is not randomized, at least in principle $\beta$ could suffer from omitted variable bias, which in turn could bias the estimation of $\alpha$. In unreported results (available upon request), we conducted an auxiliary experiment and found strong evidence that this is not a cause for concern. In a nutshell, the auxiliary experiment consists of providing information before eliciting inflation perceptions $\pi^0_{i,t}$, running a 2SLS regression of $\pi_{i,t+1}$ on $\pi^0_{i,t}$ using the random treatment assignment as an instrumental variable, and then performing the Haussman test of whether the OLS and IV estimates of $\beta$ are equal.

[11] See also Armantier et al. (2014), who provide clever tests of Bayesian learning in the context of household perceptions about inflation.

whether this prediction is accurate by estimating the basic model including an additional quadratic term, $\pi_{i,t+1} = \gamma_1 \pi^0_{i,t} + \gamma_2 \left(\pi^T_{i,t} - \pi^0_{i,t}\right) + \gamma_3 \left(\pi^T_{i,t} - \pi^0_{i,t}\right)^2$, 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 $\pi_{i,t+1} = \gamma_1 \pi^0_{i,t} + \gamma_+ \cdot 1\left\{\pi^T_{i,t} > \pi^0_{i,t}\right\} \cdot \left(\pi^T_{i,t} - \pi^0_{i,t}\right) + \gamma_- \cdot 1\left\{\pi^T_{i,t} < \pi^0_{i,t}\right\}\left(\pi^T_{i,t} - \pi^0_{i,t}\right)$ and then testing whether $\hat{\gamma_-} = \hat{\gamma_+}$. Some of these additional robustness checks are provided in the empirical exercises presented below.

## 2.3 Disentangling Genuine from Spurious Learning

A further 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, respondents may react in order to show agreement with the interviewer due to a desirability bias (Goffman, 1963), a fear of being deemed ignorant, or unconscious numerical anchoring (Tversky and Kahneman, 1974). These spurious effects are a major concern for our experiments and for information provision experiments in general. This does not mean, however, that these exercises are invalid: our framework attempts to quantify how much of the estimated $\alpha$ reflects genuine, rather than spurious, learning, and to disentangle the two. To this end, we make use of data on expectations about other nominal variables and of information on the evolution of expectations obtained through follow-up surveys.

The key to gauging how much of the reaction to the information provided reflects genuine learning lies in variations in the forward-looking equation (1). The first and simplest exercise consists of conducting follow-up interviews with the same subjects several months after the initial experiments. This allows us to measure the persistence of the informational treatments' effects. Let $\pi^{follow-up}_{i,t+1}$ be the inflation expectations in the follow-up survey. The forward-looking equation (1) then becomes: $\pi^{follow-up}_{i,t+1} = \mu + \beta \pi_{i,t}$, where $\beta$ is the degree of pass-through from inflation perceptions as stated in the original survey to inflation expectations stated in the follow-up survey. Combined with the learning equation in (3) we obtain:

$$\pi^{follow-up}_{i,t+1} = \gamma_0 + \underbrace{\gamma_1}_{\beta} \pi^0_{i,t} + \underbrace{\gamma_2}_{\alpha\beta} \left(\pi^T_{i,t} - \pi^0_{i,t}\right) \tag{5}$$

In other words, we use the same estimation procedure, but instead of using $\pi_{i,t+1}$ as the dependent variable we use $\pi^{follow-up}_{i,t+1}$. One concern with spurious learning is that individuals may report expectations closer to the information provided by the experimenter due to numerical anchoring. Numerical anchoring is, by definition, very short-lived, so we would not expect that to be a concern in the context of this experiment. Regarding interviewer pressure, it is unlikely that the subject will feel any pressure months after the information is received. Indeed, it is unlikely that the subject would remember the exact value of the signal provided by the experimenter.[12] Therefore,

[12]Significantly, most of the survey experiments discussed below were carried out online, which reduces the potential

the estimate of $\alpha$ in this specification would not be subject to these sources of spurious learning (or, at least, not to the same degree). Thus, the ratio between the $\alpha$ estimated with $\pi_{i,t+1}^{follow-up}$ and the $\alpha$ estimated with $\pi_{i,t+1}$ can provide an estimate of the amount of spurious, rather than genuine, learning.

The second test of spurious versus genuine learning is based on individuals' perceptions and expectations regarding other economic indicators closely related to inflation. In the context of 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 basically says that individuals partially understand the Fisher equation (Behrend, 1977). If, as a consequence of an informational treatment, an individual truly believes that future inflation will be 1 percentage point higher, we should also observe that this individual expects an interest rate that is 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.

Let $i_{i,t+1}$ denote the expectation about the nominal annual interest rate. Formally, this test consists of replacing the forward-looking equation (1) by $i_{i,t+1} = \mu + \beta\pi_{i,t}$, where $\beta$ is the degree of pass-through from inflation perceptions today to interest rate expectations. Combined with the learning equation in (3) we obtain:

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

Again, this corresponds to using $i_{i,t+1}$ as dependent variable instead of $\pi_{i,t+1}$.[13] Once again, by comparing how much smaller $\alpha_i$ obtained from the learning equation with $i_{i,t+1}$ as the dependent variable is than $\alpha_\pi$ obtained from the learning equation with expected inflation as the dependent variable, we can quantify how much of the learning is genuine rather than spurious.

of perceived interviewer pressure.

[13]Once again, an implicit assumption is that the coefficient on $\pi_{i,t}^0$ – which is identified solely with non-experimental variation - is not subject to omitted-variable bias. This assumption can also be tested with an ancillary experiment, as discussed in footnote 10.

# 3 Results from Online Experiments in the United States and Argentina

## 3.1 Evidence from the United States

### 3.1.1 Subject Pool, Descriptive Statistics and Informational Treatments

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).

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 1.a depicts the relationship between the two variables by means of a binned scatterplot. There is a strong positive association between the two, which appears to be linear. The coefficient of perceptions in a regression with expectations as the dependent variable is 0.782 for the control group (p-value$<$0.01), with an r-square of 0.29.

After eliciting past inflation perceptions (inflation over the previous twelve months), all respondents saw the following message : "We will now ask you about future inflation." The treatments in the U.S. online experiment differed in the type of information provided, with a control group and four treatment arms randomly assigned with equal probability. Appendix E.3 provides a snapshot of the informational treatments as seen by the respondents. Individuals in the control group were presented with a direct question about inflation expectations for the next twelve months. Panel (c) of Figure 2 illustrates the *Statistics (1.5%)* treatment arm, which 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). 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. The average of the three statistics indicated an annual average inflation rate of 1.5%, which was also displayed on the table.

The *Products* treatment arm, in turn, presented respondents with a table containing the prices of six products at the time of the survey and one year earlier, as well as the price change (in percentage) 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 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 (see Figure 2). Respondents in this treatment arm were randomly assigned one of ten tables with average price changes ranging from -2% to 7% in 1 percentage point increments, indicated in parentheses after the *Products* treatment arm name in the rest of this paper. 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 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 2). If the empirical results show a significant degree of "learning" from this information, this would indicate that some of the respondents had a very limited level of knowledge and, thus, they used this fictitious scenario as a benchmark.

### 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 the control group as a benchmark. The average reported inflation expectations for the following twelve months for this group is 5.08%,

[14]The 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. 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 also selected products with similar initial prices within each of the product categories.

with a standard deviation of 5.81. The distribution of this variable is concentrated between 0 and 6% (about 75% of the observations), although there is also a substantial fraction reporting values of 10% or more (the histograms accumulate the observations below -5% and above 15% in the extremes).

Our informational treatments should affect both the mean and the distribution of inflation expectations. The tables depicting products with average price changes of, for instance, 2%, should change the mean inflation expectation, compressing its distribution around that value. Figure 3 shows that this is the case. 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 (with an average annual rate of 1.5% – see panel (c) in Figure 2 for the actual information provided). As expected, this informational treatment substantially reduced the average inflation expectations with respect to those of the control group (2.13% compared to 5.08%) and substantially changed the distribution of this variable (standard deviation of 1.98 compared to 5.81). The histograms indicate that the distribution is considerably more concentrated in the 0 to 4% range (over 85% of those in this group compared to about 40% of those in the control group).[15] This indicates that our experimental subjects reacted substantially and incorporated the information about official statistics on inflation we provided into their inflation expectations.

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 (see panels (a) and (b) in Figure 2 for two examples of the actual information provided). In panel (c), we verify that individuals who were shown tables with average price increases of 0 or 1% exhibit lower average inflation expectations than those in the control group; the distribution is compressed around 0-1%, as expected and the whole distribution is shifted to the left. In panel (d), we observe that those who were shown tables with average price increases of 2 or 3 percent per year also have lower average inflation expectations than those in the control group, whose inflation expectations are concentrated in the 2-3% range. All the specific prices treatments had a significant effect on the distribution of inflation expectations in the expected direction: respondents' reported expectations exhibit a positive correlation with the average price chances in the tables we presented (see the discussion of Figure 4 and Appendix C for more detailed results on the other *Products* sub-treatments). These results indicate that individuals incorporate the prices of specific products in forming their inflation expectations if this information is available to them.

As in the *Statistics* treatment arm, 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%. Immediately afterward, those subjects were presented with one of the *Products* tables with

[15] Each panel in Figure 3 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%).

the price changes of actual products. It would be expected that the inflation expectations of subjects who receive nationally representative, aggregate official statistics on inflation would not be affected by information on the price changes of a few arbitrarily selected products in. 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: there is a substantial difference in the distribution of inflation expectations for the treatment groups *Statistics (1.5%)+Products* with the 0-1% average price change tables and the *Statistics (1.5%)+Products* with 2-3% price change tables. Moreover, these distributions also vary from the distribution of inflation expectations for the *Statistics (1.5%)* treatment arm (panel (a)) and for the *Products* treatments with the same range of price changes (panels (c) and (d)). The distribution of expectations for the groups receiving the *Statistics (1.5%)+Products* with the 0% and 1% average price change tables also shifted substantially to the left, with a higher concentration in the 0-1% range; the mass of the interval for the *Statistics (1.5%)+Products* group with 2% and 3% tables increased even more. This implies that individuals change their inflation expectations on the basis of information about the price changes of specific products even when aggregate representative statistics are readily available. One possible interpretation, consistent with rational inattention, is that it is less costly for people to incorporate information on individual prices, simply because they may be easier to understand.

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. Average inflation expectations were not affected by this treatment (the rate was 5.06% for this group and 5.08% for the control group). The ES test does, however, indicate a significant difference between the distribution of inflation expectations for this treatment group and for the control group. This can be attributed to a mass in the distribution at the 10-11% range.

Figure 3 provides summaries of the different information levels within each treatment arm. The results presented in panel (a) of Figure 4 depict 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. The results from the *Statistics (1.5%)+Products* treatment was very similar (see Figure C.3, panel (a)). The effect of the *Statistics (1.5%)* treatment on inflation expectations (which is not depicted in Figure 4) was -2.94, a relatively large figure (significant at the 1% level). The *Hypothetical* treatment, meanwhile, had virtually no effect on mean inflation expectations, with a coefficient close to zero (-0.0155) and not statistically significant (p-value of 0.952).

#### 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 2.2. The main estimates from the learning model for the U.S. online experiment are presented in Table 1. The table reports the values of $\alpha$ and $\beta$ from equation 4. As discussed above, $\beta$ can be interpreted as the degree of pass-through between perceptions of past inflation and expectations of future inflation, and $\alpha$ 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.

We estimated this regression separately for the four different treatments. The first pattern that emerges from Table 1 is that, consistent with panel (a) in Figure 1, there is a high correlation between inflation perceptions and inflation expectations, reflected in a relatively high value for $\beta$ which varies from 0.718 to 0.829, 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 $\alpha$ for the factual informational treatments in columns (1) to (3). The weight given to the information in the *Statistics (1.5%)* treatment was 0.847, whereas the weight given to its equivalent in the *Products* treatment was 0.710 (the difference between the two is statistically significant at the 1% level). Moreover, the sum of the two values of $\alpha$ for the *Statistics (1.5%)+Products* treatment, in which we combined the two pieces of information, is 0.732, that is, between the values of $\alpha$ reported in the two previous columns.

An estimated $\alpha$ of about 0.7-0.85 means that, in forming their posterior beliefs about inflation expectations, individuals in our sample assign a 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. 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.[16] 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, etc.[17]

A second notable result from Table 1 is that both the information from the specific product

[16] Significantly, 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.

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

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 $\alpha$ coefficients. The combined treatment consisted of displaying first the table with three official annual inflation indicators, and then the table with six actual products and their price changes. The latter information not only had an effect on inflation expectations even when inflation statistics were made available, its effect was also stronger than that of statistics: the $\alpha$ coefficient for the tables is 0.449, substantially higher than the $\alpha$ of 0.283 for the official statistics in the same experiment (the difference is statistically significant at the 1% level). These results suggest that 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 official 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. 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. One prediction yielded by this model is that providing relevant information will increase the accuracy of the later belief, $\sigma_{i,t}$. We 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. The mean standardized confidence for the control group is -0.191; the confidence is significantly higher for the *Products*, *Statistics (1.5%)* and the *Products+Statistics (1.5%)* treatments, with virtually no difference for the *Hypothetical* treatment arm.[18] Moreover, the learning model also predicts that all signals, regardless of value, should be equally informative for 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 had similar effects on respondents' confidence in their stated expectations, although with a slight asymmetry. We can reject at standard levels the equality of all ten coefficients (p-value 0.0475). This is evidence that individuals are less prone to incorporate information about price decreases than about price increases. This survey data also reveals that even individuals with biased inflation expectations report confidence about their stated expectations. For individuals in the control group, 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

[18]The 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).

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. This is consistent with the view that a substantial portion of the observed heterogeneity in reported inflation expectations reflects actual heterogeneity in beliefs rather than measurement error or rational inattention.

A related test suggested in section 2.2 is that there may be non-linearities or asymmetries in the reaction to the information provided (e.g., individuals may learn more from signals that are closer to their prior belief). Columns (1) to (3) in Table 2 present some robustness tests of the learning results for the *Statistics (1.5%)* treatment arm, and columns (4) to (6) present the results for the *Products* treatment. The coefficients in columns (2) and (5) present a specification with a quadratic term (as discussed at the end of section 2.2). The corresponding estimates for this coefficient are virtually zero (0.007 and -0.003, respectively), and the linear terms for $\alpha$ and $\beta$ are very similar to those presented in columns (1) and (4), which reiterate the baseline estimates of Table 2. This supports the validity of the learning model. Columns (3) and (6) 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 $\alpha$ of 0.632 (*Statistics*) and 0.606 (*Products*) for those with $\pi_{i,t}^{T} - \pi_{i,t}^{0} \geqslant 0$, 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 statistically significant for the *Products* treatment (p-value of 0.22). While the quadratic coefficient for *Statistics* is small and not statistically significant, there is evidence of mild asymmetry for this treatment, which is reflected in the significant difference between the coefficients for positive and negative differences.[19] Appendix C.3 presents further tests of the Bayesian learning model; the results generally support 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 more pressing concern is whether or not the learning induced by our experimental setup is spurious. As discussed in section 2.3, 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. Under the pretext of a cognitive test, a randomly selected group of our subjects were knowingly given information about the price changes of fictitious products. The respondents were asked to "eyeball" the price change of a product that, over the course of a one-year period, experienced an increase

[19] A possible interpretation of this difference is that individuals with $\pi_{i,t}^{T} - \pi_{i,t}^{0} < 0$ are those who have perceived past inflation $\pi_{i,t}^{0}$ to be higher; these same individuals are likely to be less informed and thus more prone to incorporate the information provided. This evidence could still suggest small deviations from rational learning having to do with an asymmetric reaction to information about price changes.

of about 10% (panel (d) of Figure 2). 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 (4), Table 1. The coefficient $\alpha$ for the *Hypothetical* treatment is substantially lower than for the other treatments; that value stands at 0.185 (statistically significant at the 1% level). Though statistically significant, this rate is economically insignificant when compared to the learning rates of 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 willing to use inflation figures from a hypothetical exercise as a benchmark.

We present two further tests following the discussion in section 2.3. For the first test, we estimated the learning model using the inflation expectations in the follow-up survey to weed out spurious learning. We used data on a subsample of 513 subjects who we re-interviewed two months after the original online experiment. This subsample was asked again about their perceptions of past inflation and their inflation expectations, but they were not subjected to any type of informational treatment.[20] Column (5) in Table 1 presents the results of the basic regression with inflation expectations in the original survey as the dependent variable for the *Products* treatment group, but only for the subsample of those who later participated in our follow-up survey. The $\alpha$ and$\beta$ coefficients are very similar to those presented in column (2) for the full sample (0.704 compared to 0.710 and 0.845 compared to 0.777, respectively). Column (6) presents the regression for the same follow-up subsample, but in this case inflation expectations as reported in the follow-up survey are the dependent variable. The $\alpha$ coefficient of 0.325 is statistically significant at the 10% level. Though it is less than half as large as the figure in column (5), this coefficient indicates that 46.2% of the effect of the information provided can be attributed to genuine, rather than spurious, learning. The results for inflation expectations reported in the follow-up survey are similar to the results for other treatment arms with the exception of the *Hypothetical* treatment for which the $\alpha$ coefficient is close to zero and statistically insignificant (results not reported). This is consistent with the interpretation of short-term anchoring effects. While some of the learning seems to be spurious, about half of it can be attributed to a genuine effect of the information provided.

The second test consists of measuring learning rates based on the indirect effect of the information provided on a related outcome, the expected nominal interest rate. Behrend (1977) 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. We report results from this exercise in Columns (7) and (8) of Table 1, where the dependent variable is an individual's expectation for interest rates for the following twelve months. Column (7) presents the results for the *Products* treatment. The positive and significant $\beta$ coefficient indicates a positive relationship between the two variables. Notably, the $\alpha$ coefficient is statistically significant and,

[20]Multiple tests suggest that selective attrition is not a concern (results not reported).

although lower than the corresponding coefficient for inflation expectations shown in column (1), it is relatively high (0.549). When the learning rate for interest rates is compared to the learning rate for inflation expectations (a figure of 0.710, as shown in column (2), *Products* treatment), it appears that almost 80% of the learning is genuine. The results in column (8), on the other hand, indicate that the *Hypothetical* prices treatment, which provided a non-factual signal, did not have a significant effect on individuals' expected interest rates. This is the case despite the small but significant effect of this treatment on inflation expectations (column (4)). This difference can be interpreted as evidence that this non-factual treatment did not induce genuine learning on participants. Interestingly, the difference of 0.161 between the $\alpha$ coefficient for inflation expectations and that coefficient for the expected nominal interest rate for the *Products* treatment (columns (1) and (7)), which represents spurious learning for that treatment, is very close to the (seemingly spurious) learning coefficient for inflation expectations in the *Hypothetical* treatment, a figure of 0.185 (column (4)).[21] The results for the nominal interest rate also support the 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 levels. However, it may be argued that individuals may respondent about their own idiosyncratic experience – i.e., the price change of their own consumption basket. The results described in this paragraph show us that this is not 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 are aggregate and not idiosyncratic variables.

The results from the follow-up survey indicate that about half of the learning was spurious, whereas spurious effects account for a much smaller fractions for other expectations of nominal variables such as the interest rate (which were collected at the time of the original survey). We can rationalize this difference as follow. Providing information can have two effects in the short run. The first effect is learning: those who did not know that inflation was around 1.5% incorporate this information. A second effect is salience: even those already aware of this information assign a higher weight to it. The nature of the salience effect implies that it will be short-lived, and as such it is likely to have disappeared by the time of the follow-up survey.

## 3.2 Evidence from Argentina

### 3.2.1 Subject Pool, Descriptive Statistics and Informational Treatments

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, on average, 1.8%, in Argentina the average rate for the same time period was around 22.5%. As a result, the cost of ignoring inflation in Argentina was substantially higher. For

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

example, individuals must rely on good information on inflation prospects in drawing up contracts (labor, real estate, etc.), because it is illegal to index such contracts or to use more stable foreign currencies.[22] Opinion polls in Argentina at the time of the survey systematically indicated inflation as one of the population's primary concerns.[23] Inflation statistics were mentioned on offline and online news outlets on a daily 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), who were interviewed in June 2013. This sample, which yielded a total of 691 observations, was assigned to a control group, or to the *Statistics (24%)*[24] 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; the survey has contained the same set of basic questions since 2011.[25] For this sample, which was interviewed in March 2013, 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); 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.

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 1 depicts the relationship between the two variables, with a binned scatterplot. As in the U.S. survey, there is a strong, linear, and positive association between the two. The coefficient of perceptions in a regression with expectations as the dependent variable is 0.883 (p-value<0.01), with an r-square of 0.60 (control group only).

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

[23] For 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.

[24] The 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 daily basis (for more details, see Cavallo, Cruces and Perez-Truglia, 2014).

[25] This sample is 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 average Argentine. See the Appendix for comparative descriptive statistics of our samples and the Argentine population.

### 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, providing these two types of information had a significant effect on the distribution of inflation expectations, which are more densely concentrated than those for the Control group. The p-values for the difference of distributions between each of the two groups and the control group indicate that both treatments had statistically significant effects on the distribution of inflation expectations. For 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).

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%)* and *Products (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 distribution to the left of the control group's. Conversely, inflation expectations of the respondents in the *Products (31%)* and *Products (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). The full set of results follows the expected pattern: respondents' reported inflation expectations clearly increased in conjunction with the average price changes in the table they were presented as part of the treatment (see Appendix Figure D.2 for a more detailed analysis by treatment level). 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; the analysis of the learning model below explores this question more formally. In terms of the size of the effect, 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

This section presents the results for regressions yielded by the learning model of Equation 4 for the Argentina online experiments. Columns (1) and (2) of table 3 present the results of the learning model for the *Statistics* and *Products* treatments based on the sample of college graduates. In column (1) (in which we provided respondents with inflation statistics) and in column (2) (the *Products* treatment with average price changes of 19%, 24%, and 29%), the value of $\alpha$ is close to 0.5: 0.506 for *Statistics* and 0.459 for *Products*. Notably, the coefficient $\alpha$ in the estimate for the *Products* treatment in the larger opinion poll sample (column (3)) is 0.494, that is, very close to 0.5.[26]

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. The results in the table also indicate a very high pass-through from perceptions of past inflation to expectations of future inflation. While substantial, the weight individuals in Argentina assigned to the information about prices changes for specific products is substantially less than the value of around 0.71 that what we found for our U.S. sample. The fact that learning rates were 40% higher in the low-inflation U.S. context is consistent with the prediction of the rational inattention model whereby 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).

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 a first test, panel (b) in Figure 6 compares the impact of each treatment level on the standardized level of self-reported confidence about the answer to the question regarding inflation expectations. The results suggest that respondents in all treatment levels gleaned the same amount of information, which is compatible with the prediction of the learning model. This model might still be too simple, though. The result of a second test which entails an alternative specification with a quadratic term is provided in column (4) of Table 3. The results indicate that the linear terms for $\alpha$ and $\beta$ are very similar to those presented in column (3), while the coefficient for the quadratic term is not statistically significant (it is virtually equal to zero). This indicates that our learning model is a good fit for the data. Column (5) in Table 3 presents the results of an alternative specification that contemplates differential learning for positive and negative differences between the signal and the prior. The coefficient $\alpha$ of 0.484 for those with $\pi^T_{i,t} - \pi^0_{i,t} \geqslant 0$ and of 0.497 for those with $\pi^T_{i,t} - \pi^0_{i,t} < 0$ suggests that learning is perfectly symmetric. This result contrasts

[26] One 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.

with the evidence in the U.S. sample of small deviations from rationality in the form of asymmetric responses to signals about inflation.

#### 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.[27] 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.[28] Column (6) in Table 3 presents the results of the basic regression where inflation expectations in the original survey are the dependent variable for the subsample of the *Products* treatment group that later participated in the follow-up survey. The $\alpha$ and $\beta$ coefficients are very similar to those presented in column (3) for the full sample (0.963 compared to 0.902 and 0.456 compared to 0.494, respectively). Column (7) presents the regression for the same follow-up subsample, but in this case inflation expectations as reported in the follow-up survey were the dependent variable. The $\alpha$ 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 46.1% in the context of a similar follow-up survey. In sum, while there is a significant level of spurious learning, about half of the total can be attributed to a genuine effect of the information provided.

In column (8) 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 $\alpha$ 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 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; some contracts are indexed according to the exchange rate, and investment is also related to the value of the local currency against the dollar. Most individuals are aware of the market value of this exchange rate (the 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) and have expectations about its future evolution. The $\alpha$ coefficient from

[27] This 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.

[28] There was no significant difference in the probability of showing up at the follow-up sample between the treatment and the control groups. Furthermore, we estimated the learning regression with an attrition indicator as the dependent variable and neither $\alpha$ nor $\beta$ was statistically significant.

this estimation, presented in column (9), is 0.435, that is, very close to the figure for the nominal interest rate (column (8), 0.468) and for inflation expectations (column (3), 0.494).

# 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. Indeed, 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. Nearly 63.4% of respondents in our U.S. online experiment and 74.9% of respondents in our Argentina online experiment reported to have thought about the prices of specific products when they were asked how they learned about inflation. While 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. Furthermore, the consequences of using price memories to form inflation expectations depends on how good an individual's memory of price histories is. For example, if a price memory is very accurate, then it would be optimal for an individual to incorporate the information it provides into her expectation, as it could correct for biases in inflation statistics arising from differences between the individual's own consumption basket and the basket used to compute inflation statistics. The data requirements for testing these additional hypotheses are onerous. 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 would be required. Moreover, experimental variation in the price memories of subjects would have to be created. The unique consumer intercept survey we conducted at the main exit of several supermarkets in Buenos Aires meets 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 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 next twelve months. The first treatment, which was aimed at generating random variation in the salience of the individual's own price memories about specific products, consisted of asking the individual to recall present and past prices for a set of four products that were randomly selected from the subject's receipt. The second treatment, identical to the one used in the online experiments, consisted of showing the individual the actual price histories for six randomly-chosen products.

## 4.3 Effects of the Use of Actual and Remembered Price Changes on the Formation of Inflation Expectations

The informational treatments in the supermarket experiment were similar to those in the online experiments. As with the *Products* treatments described in the previous two sections, in the supermarket experiment one piece of information consisted of a table of price histories for four products. We randomly assigned one of three tables with levels of average price changes of 19%, 24%, and 29%. Panels (a) and (b) in Figure 7 present the distributions of inflation expectations in pairwise comparisons between treatments. While there is no 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 do differ significantly: average inflation expectations are clearly higher for the latter treatment.

Respondents were also asked to recall the current price, and the price twelve months earlier, of two specific products they had just purchased. These products were randomly selected by the interviewer from the products listed on the respondent's receipt. The interviewers selected two additional products from the receipt, read each of their prices out loud, and asked the respondents what they thought the prices of these two products had been twelve months earlier. The experimental intervention thus consisted of randomizing the products from the purchase receipt whose current and past prices the individual was asked to provide. Some individuals were asked about products that had experienced larger changes in price and others were asked about products with smaller changes in price. Rather than providing a table with product prices and price changes as the online experiments does, this study effectively asked respondents to "fill in the table" by recalling past and current prices for the products they had just purchased and which, therefore, were relevant to them. Panel (c) in the Figure 7 presents a comparison of the distribution of inflation expectations between two groups defined according to "high" and "low" levels of price changes on the basis of respondents' memories.[29] The evidence indicates that respondents who reported higher proportional price changes for some of the products they had purchased also reported significantly higher inflation expectations. Panel (d) in Figure 7 presents a comparison of the distribution of inflation expectations between groups defined according to "high" and "low" levels of price changes on the basis of the actual price changes of the products that the individual was asked about. The evidence indicates that respondents who were asked about products with higher price changes did not expect a significantly higher level of inflation. In other words, whatever memory information individuals were calling on was completely orthogonal to the actual price changes: what individuals were adding to their memories was pure noise. Interestingly, the results are the same 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).

In 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 an $\alpha$ with the usual regression would not yield the same interpretation in terms of rate of learning as it had in the case of other information treatments.[30]

## 4.4 Effect of the Accuracy of Memories about Current and Past Prices

The previous discussion establishes that, insofar as individuals' price memories are accurate, relying on those memories would not exert a substantial bias on inflation expectations. The evidence presented in the previous section, though, indicates that, when the salience of a group of products was randomized, individuals reacted to the recalled price changes but not to the actual price changes of those products. This section provides more direct evidence on the accuracy of these recalled prices.

Panel (a) in Figure 8 presents a scatterplot of prices for products respondents had just purchased; it shows the prices they reported paying for those items without looking at the receipt and the prices they actually paid for them. The relationship between the two variables seems to be linear, with most observations clustered around the 45 degree line, indicating that individuals' memories of the prices of the products they had just purchased were fairly accurate. Panel (b) in Figure 8 presents the results of a more taxing exercise for respondents' memory: we present a scatterplot of respondents' reported recollections of the prices of the same goods one year earlier and of the prices in 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.[31] Individuals

[29] We divided products into three equally sized groups according to price change: low, middle and high.

[30] Appendix E presents regression for the corresponding rate of learning, although these results should be interpreted with these caveats in mind.

[31] Bates and Gabor (1986) and Kemp (1987) also find that individuals' implicit price changes overestimate the

seem to underestimate the past prices of the products they had purchased. This may, in part, be because individuals struggle with the operation of projecting percentage changes into the past or because individuals who have an accurate notion of aggregate inflation[32] are not necessarily able to translate a current price and an inflation rate into an accurate past price.

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. Individuals might, for instance, be reporting prices for twenty months earlier rather than for twelve months earlier. Panel (c) in Figure 8 presents respondents' perceptions of aggregate inflation over the previous twelve months and the implicit average percentage price change of the products for which we requested this information. As expected, the correlation is positive and significant. Although the recalled price changes are consistently larger than individuals' reported perception of inflation, those recalled changes provide information about individuals' beliefs regarding the annual inflation rate. Panel (d) in Figure 8, in turn, presents a comparison of the implicit average percentage price change for products purchased by respondent and the actual price changes observed in our database of supermarket prices. There is no correlation whatsoever between the actual price changes and the price changes remembered by individuals in our sample. In other words, individuals' memories of price changes for specific products appear to be mostly noise. Though 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 just purchased.[33] We included one question in our study to test this very hypothesis. Immediately after asking about 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. The results are qualitatively identical to those for individual products (see Appendix for results), which indicates that respondents do not seem to fare any better when asked about total purchase amounts. All in all, this evidence suggests that – far from correcting a representativeness bias – using price memories as inputs for the formation of inflation expectations is bound to induce large errors in beliefs.

---

actual price changes.

[32]See, for example, the discussion about implicit memory in Monroe and Lee (1999).

[33]Alternatively, individuals may follow the evolution of prices for a fixed set of products (e.g., their own consumption basket), and their memories for these prices may be accurate. With this caveat in mind, we show in Appendix A that even with accurate recollections, moderate memory limitations (for instance, in the number of products an individual may keep track of) can explain the substantial heterogeneity in household inflation expectations.

# 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 a learning model. We find that individuals are highly influenced by information on both inflation statistics and price changes of specific products. The results are consistent with rational inattention, as there is greater learning in a low-inflation setting where the stakes are lower (the United States compared to Argentina) and from information that is less costly to understand (supermarket prices). To further assess the importance of personal experiences, we conducted experiments at supermarkets which combined data from actual products purchased by the subjects and their respective historical prices. We find that individuals use their own memories about the price changes of the products that they buy, but that those memories are nearly orthogonal to the actual price changes. This evidence suggests that using price memories as inputs for the formation of inflation expectations is bound to induce biases and errors in beliefs.

Our findings have a number of implications for macroeconomic theory and for 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, 2013). From a more practical perspective, our findings imply that central banks should take these insights into account in designing communication strategies for disseminating inflation statistics, strategies that make that information more credible and facilitate understanding on the part of households. Central banks may, for instance, follow the examples of the European Central Bank and the French statistical agency, which created online tools to teach households about the accuracy and the representativeness of official inflation statistics.

Our findings 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.[34] Other authors argue that the failure to incorporate public information is a natural outcome of rational inattention (Mankiw et al., 2003). But this would imply that survey data on expectations has limited value since individuals with inaccurate expectations merely reveal that they do not care about inflation. Our evidence suggests that individuals report biased beliefs on inflation also because they use private sources of information

[34]Even if, for example, individuals have an accurate intuitive idea about future inflation levels, an imperfect elicitation process may result in unreliable stated expectations (e.g., confusing annual and monthly rates). Indeed, 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.

(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 and/or rational inattention.[35]

[35]Consistent with this interpretation, our survey data indicates that individuals with biased inflation perceptions and expectations describe themselves as confident about their stated beliefs, sometimes even more confident than those with more accurate perceptions of inflation.

# 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: Microeconometric 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. (2013), “Is The Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation,” NBER Working Paper No. 19598.

[20] 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).

[21] 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.

[22] Drager, L. (2011), “Inflation Perceptions and Expectations in Sweden - Are Media Reports the ’Missing Link’?” KOF Swiss Economic Institute Working Paper No. 273.

[23] Eurobarometer (2008), “Europeans’ knowledge of economic indicators,” Special Eurobarometer 323, Wave 67.2 – TNS Opinion & Social, European Commission, Brussels.

[24] Fuster, A.; Laibson, David and Mendel, Brock (2010), “Natural Expectations and Macroeconomic Fluctuations,” Journal of Economic Perspectives, Vol. 24 (4), pp. 67-84.

[25] 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.

[26] Goffman, E. (1963), “Stigma: Notes on the Management of Spoiled Identity,” New Jersey: Prentice-Hall.

[27] Hellwig, C. (2005), “Heterogeneous Information and the Benefits of Public Information Disclosures,” Working Paper.

[28] Jonung, L. (1981), “Perceived and expected rates of inflation in Sweden,” The American Economic Review, Vol. 71 (5), pp. 961-968.

[29] 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.

[30] Kemp, S. (1987), "Estimation of past prices," Journal of Economic Psychology, Vol. 8, pp. 181–189.

[31] Lucas, R. E., Jr. (1972), "Expectations and the Neutrality of Money," Journal of Economic Theory 4 (2): 103–124.

[32] Madeira, Carlos and Zafar, Basit (2014) "Heterogeneous Inflation Expectations, Learning, and Market Outcomes." Forthcoming, Journal of Money, Credit, and Banking.

[33] Manski, C. (2004), "Measuring expectations," Econometrica 72 (5), pp. 1329–1376.

[34] 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.

[35] 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.

[36] Malmendier, U. and Nagel, S. (2013), "Learning from Inflation Experiences," Working Paper, Berkeley.

[37] 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.

[38] 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.

[39] Morris, S. and Shin, H.S. (2002), "Social Value of Public Information," The American Economic Review, Vol. 92 (5), pp. 1521-1534.

[40] Phelps, E.S. (1969), "The New Microeconomics in Inflation and Employment Theory," American Economic Review: Papers and Proceedings, 59.

[41] 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.

[42] Rosenthal, R. (1966). Experimenter effects in behavioral research. New York : Appleton-Century-Crofts.

[43] 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.

[44] Sargent, T.J. (1993), "Bounded Rationality in Macroeconomics," Oxford: Oxford University Press.

[45] Sims, C. (2005), "Rational inattention: a research agenda," Discussion Paper Series 1, Deutsche Bundesbank.

[46] Tversky, A., and Kahneman, D. (1974), "Judgement under uncertainty: Heuristics and biases," Science, Vol. 185, pp. 1124–1130.

[47] Van Praag, B., and A. Ferrer-i Carbonell (2007), "Happiness Quantied: A Satisfaction Calculus Approach." Oxford: Oxford University Press.

[48] Veldkamp, L. (2011), "Information Choice in Macroeconomics and Finance," New Jersey: Princeton University Press.

[49] Zizzo, D. J. (2010). "Experimenter demand effects in economic experiments," Experimental Economics, Volume 13, Issue 1, pp 75-98.

Figure 1: Past Inflation Perceptions and Future Inflation Expectations, Individuals in the *Control* group, U.S. and Argentina Online Experiments

a) U.S. b) Argentina

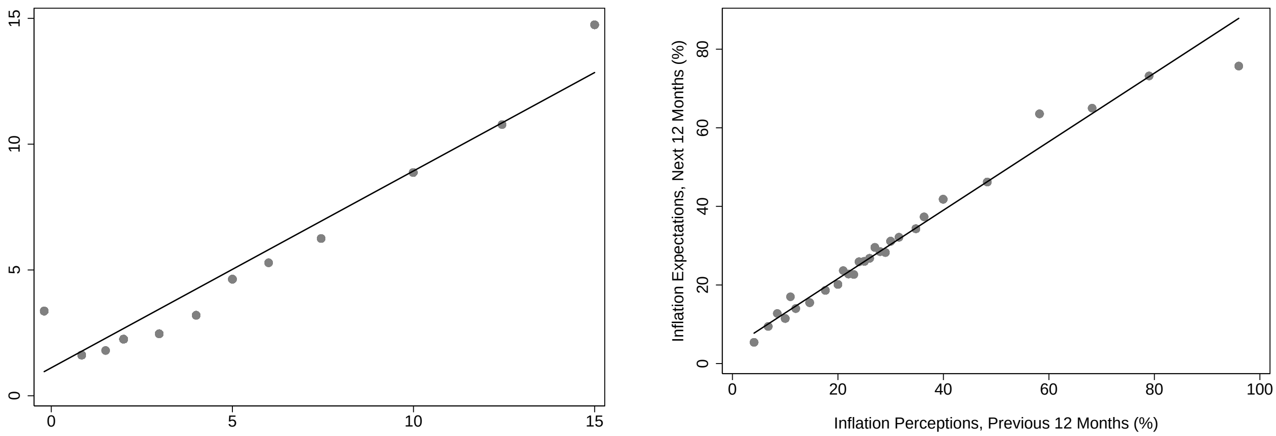


Notes: The total number of observations are 783 for the U.S. and 567 for Argentina (control group only). The darker markers represent the average inflation expectations for quantiles of inflation perceptions.

Figure 2: Example of *Products* (various levels), *Statistics (1.5%)* and *Hypothetical (10%)* Treatments, U.S. Online Experiment

a) *Products (-2%)*

| Product | Price on August 1, 2012 | Price on August 1, 2013 | Price change in % |
|---|---|---|---|
| **Infant Formula** (Enfamil Gentlease) | $18.69 | $18.69 | 0.0% |
| **Bread** (Anzio & Sons Sub Rolls) | $3.59 | $3.59 | 0.0% |
| **Pasta Sauce** (Barilla Marinara) | $2.79 | $2.80 | 0.4% |
| **Cereal** (Cheerios Honey Nut) | $5.29 | $4.99 | -5.7% |
| **Soda** (Schweppes Ginger Ale) | $1.79 | $1.67 | -6.7% |
| **Body Wash** (Dial Spring Water) | $6.09 | $6.09 | 0.0% |
| | | **Average change:** | **-2.0%** |

b) *Products (2%)*

| Product | Price on August 1, 2012 | Price on August 1, 2013 | Price change in % |
|---|---|---|---|
| **Infant Formula** (Similac with Iron) | $7.29 | $7.59 | 4.1% |
| **Bread** (Pepperidge Farm Sliders) | $3.00 | $2.99 | -0.3% |
| **Noodles** (No Yolks) | $2.79 | $2.79 | 0.0% |
| **Cereal** (Natures Path Envirokidz) | $4.99 | $5.39 | 8.0% |
| **Soda** (Dr Pepper) | $1.79 | $1.79 | 0.0% |
| **Body Wash** (Dial Spring Water) | $6.09 | $6.09 | 0.0% |
| | | **Average change:** | **2.0%** |

c) *Statistics (1.5%)*

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

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

d) *Hypothetical (10%)*

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

| | |
|---|---|
| **Price on January 1st 2012:** | $9.99 |
| **Price on January 1st 2013:** | $10.99 |

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

- o About 1%
- o About 5%
- o About 10%
- o 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 3: Inflation Expectations by Informational Treatments, U.S. Online Experiment

a) *Control* and *Statistics (1.5%)*

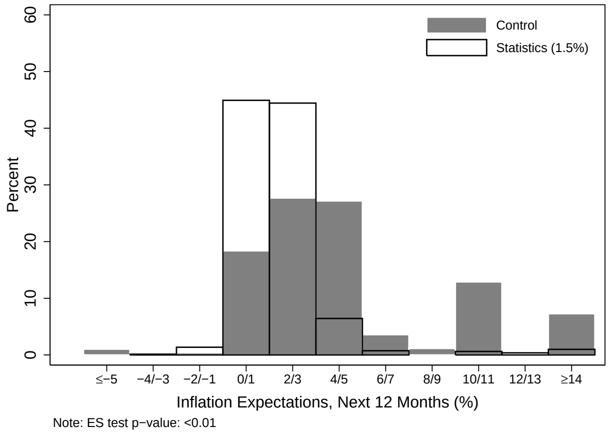


b) *Control* and *Hypothetical (10%)*

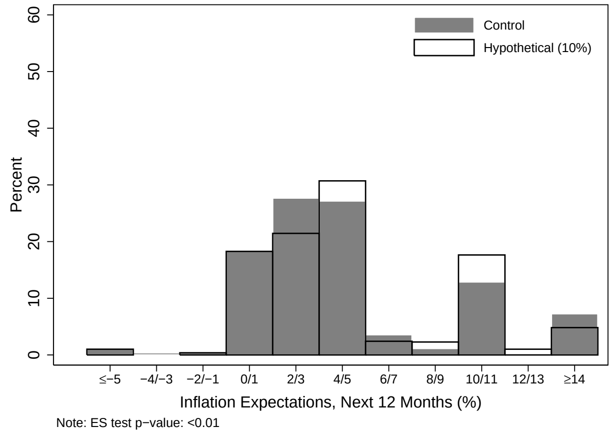


c) *Control* and *Products (0&1%)*

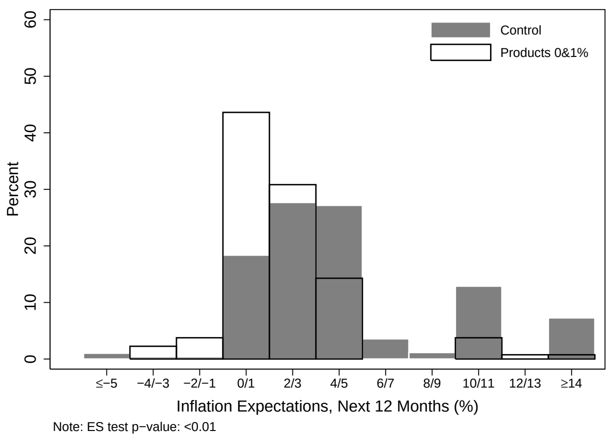


d) *Control* and *Products (2&3%)*

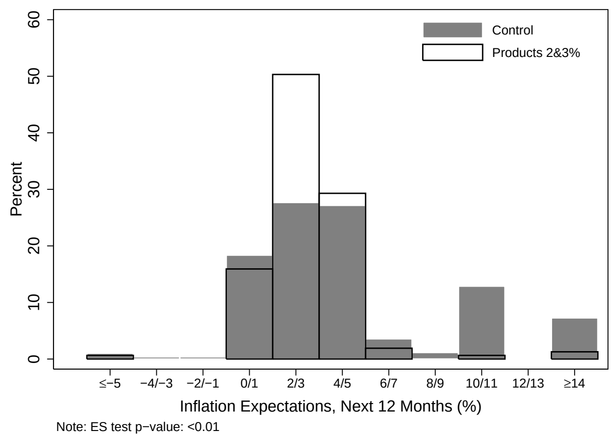


e) *Control* and *Statistics (1.5%)+Products (0&1%)*

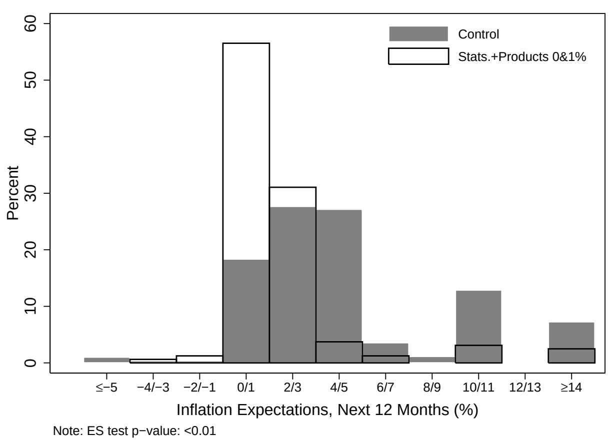


f) *Control* and *Statistics (1.5%)+Products (2&3%)*

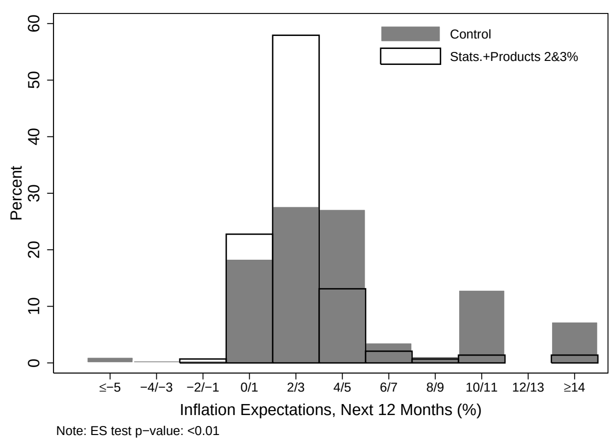


Notes: The total number of observations is 3,945, with 783 in the *Control* group, 807 in the *Statistics (1.5%)* treatment, 763 in the *Products* treatment (10 tables with average price changes from -2 to 7% in 1 percentage point increments within this treatment), 804 in the *Products+Statistics (1.5%)* 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

a) Effect on inflation expectations

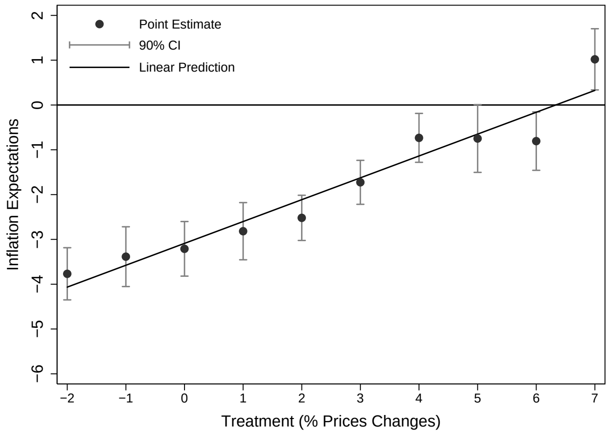


b) Effect on confidence

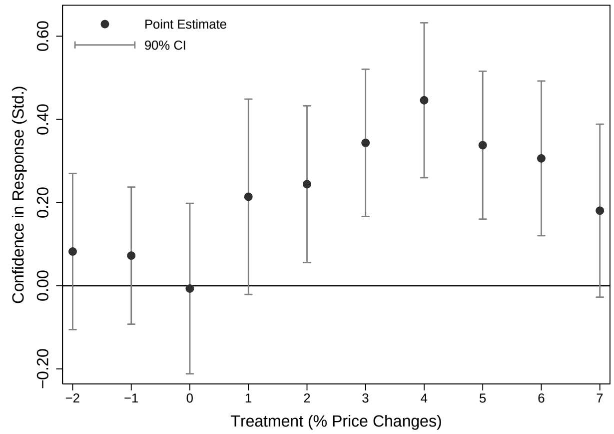


Notes: The total number of observations is 1,552 (783 in the control group and 763 in the 10 variations of *Products* treatment). Each bar represents the point estimate of the effect of the specific sub-treatment (average product price changes in the table presented) compared to the control group. Robust standard errors reported.

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

a) *Control* and *Statistics (24%), sample I*

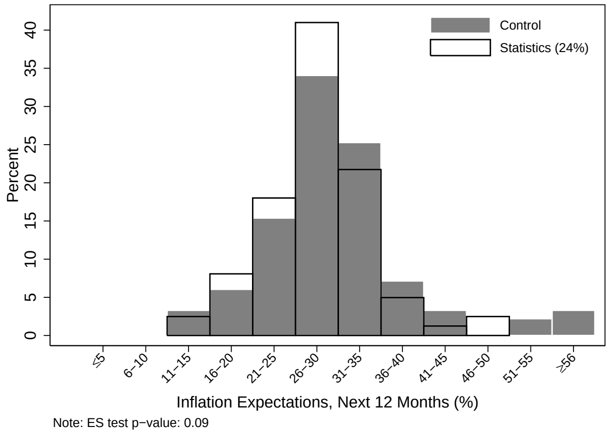


b) *Control* and *Products (24%), sample I*

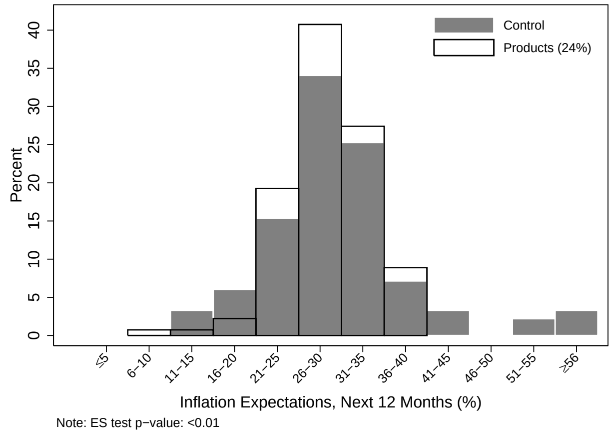


c) *Control* and *Products (18&19%), sample II*

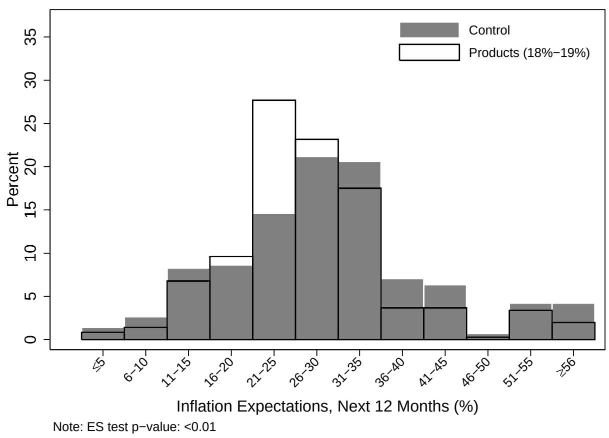


d) *Control* and *Products (31&32%), sample II*

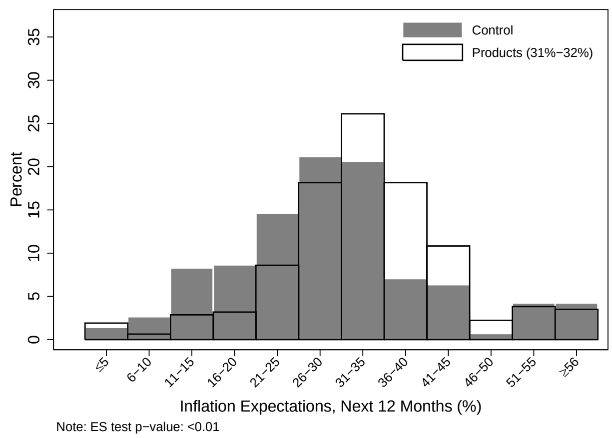


Notes: Panels (a) and (b) present results for the college graduates online experiment sample (sample I). The total number of observations is 641, with 174 in the *Control* group, 127 in the *Products (24%)* group, and 146 in the *Statistics (24%)* group. Panels (c) and (d) present results for the opinion 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

a) Effect on inflation expectations b) Effect on confidence

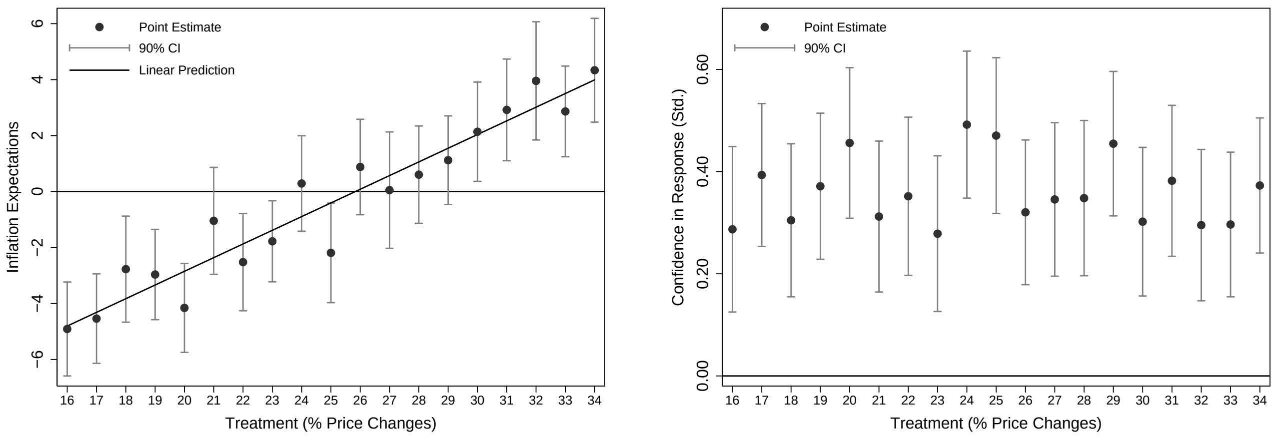


Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 *Products* treatment groups. Each bar represents the point estimate of the effect of the specific sub-treatment (average price changes for each product in the table presented) compared to the control group. Robust standard errors reported.

Figure 7: Inflation Expectations by *Product* Treatment Levels and by Remembered and Actual Price Changes, Supermarket Experiment, Argentina

a) *Products (19%)* and *(24%)*

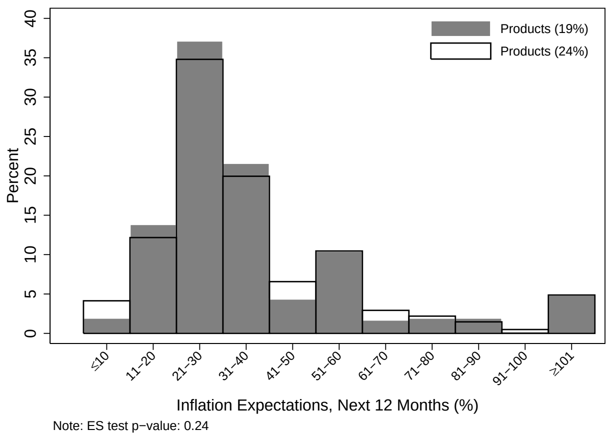


b) *Products (19%)* and *(29%)*

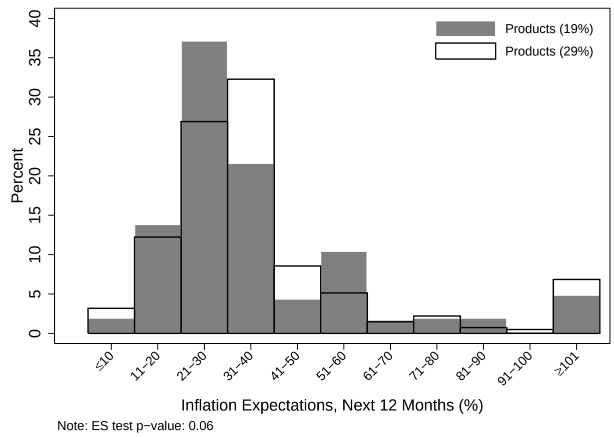


c) Remembered price changes

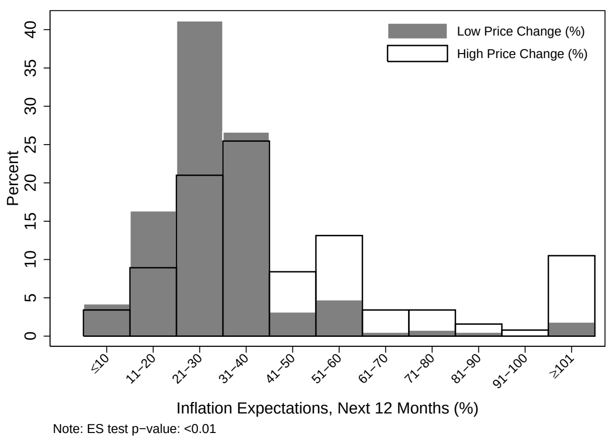


d) Actual price changes

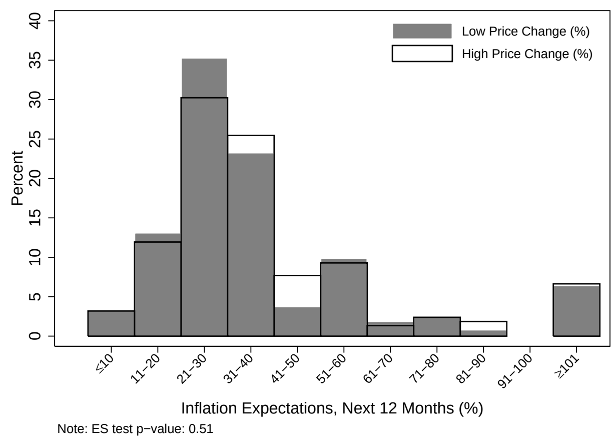


Notes: The total number of observations is 1,232 for panels (a) and (b) (412 in the *Products (19%)* group, 411 in the *Products (24%)* group and 409 in the *Products (29%)* group). The number of observations in panels (c) and (d) are 379 (lowest third of remembered price changes) and 381 (top third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

Figure 8: Remembered and Actual Past Prices, Implicit Price Changes and Inflation Expectations, Supermarket Experiment, Argentina

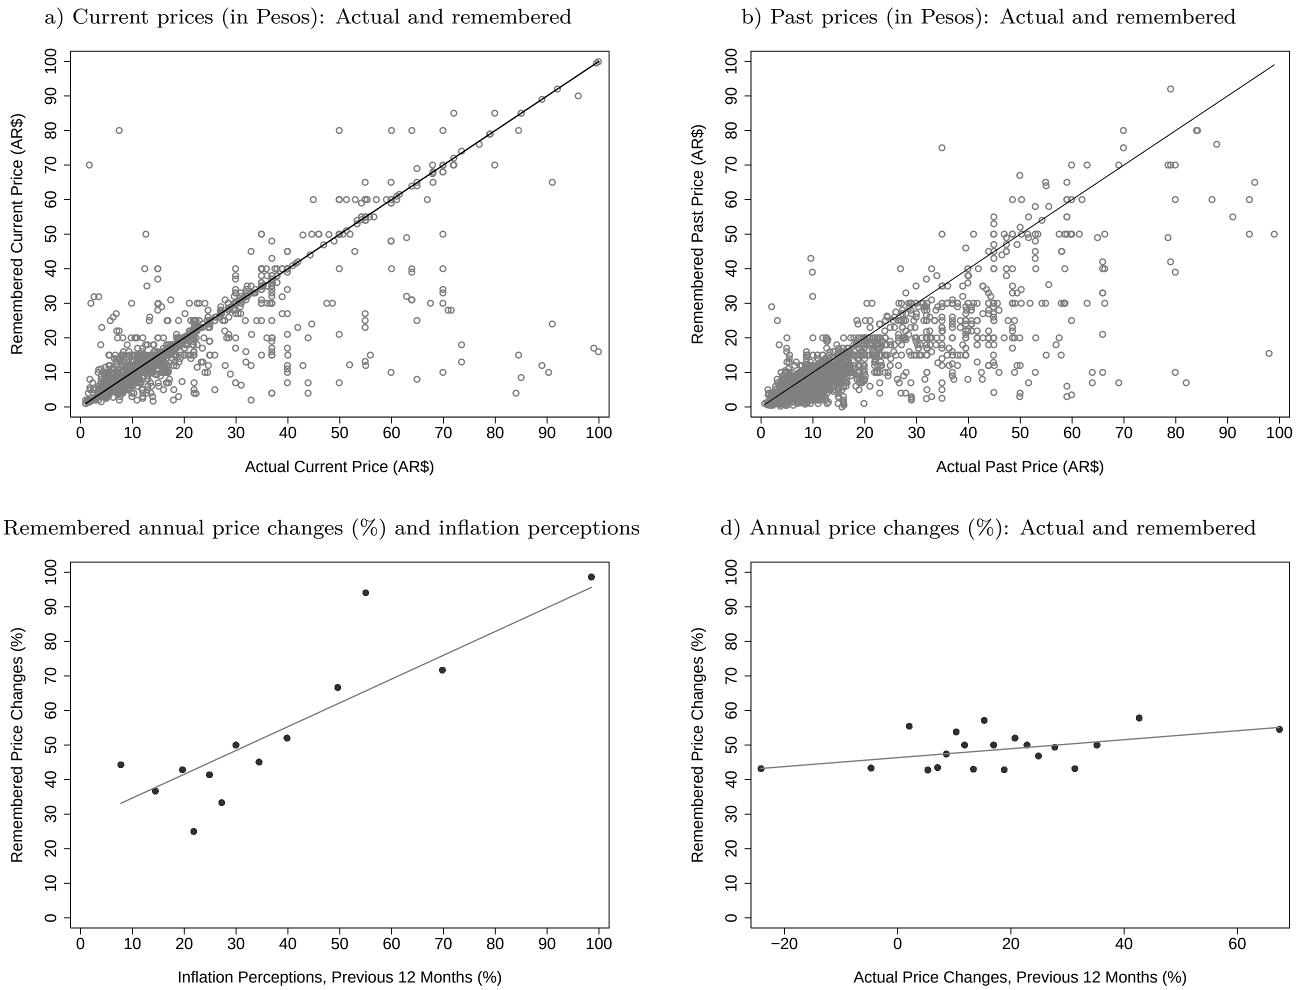


Notes: The total number of observations is 1,140. Panels (c) and (d) represent binned scatterplots. The annual price changes in panels (c) and (d) are implicit; they are obtained from the current and past prices in pesos (AR$) reported by the respondents.

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

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
|---|---|---|---|---|---|---|---|---|
| | $\pi_{\mathrm{i,t+1}}$ | $\pi_{\mathrm{i,t+1}}$ | $\pi_{\mathrm{i,t+1}}$ | $\pi_{\mathrm{i,t+1}}$ | $\pi_{\mathrm{i,t+1}}$ | $\pi_{\mathrm{i,t+1}}^{follow-up}$ | $i_{\mathrm{i,t+1}}$ | $i_{\mathrm{i,t+1}}$ |
| $\beta$ | 0.829*** | 0.777*** | 0.817*** | 0.718*** | 0.845*** | 0.402*** | 0.275*** | 0.267*** |
| | (0.057) | (0.051) | (0.058) | (0.047) | (0.072) | (0.075) | (0.059) | (0.061) |
| *Statistics* | | | | | | | | |
| $\alpha$ | 0.847*** | | 0.283*** | | | | | |
| | (0.033) | | (0.063) | | | | | |
| *Products* | | | | | | | | |
| $\alpha$ | | 0.710*** | 0.449*** | | 0.704*** | 0.325* | 0.549*** | |
| | | (0.036) | (0.050) | | (0.044) | (0.167) | (0.150) | |
| *Hypothetical* | | | | | | | | |
| $\alpha$ | | | | 0.185*** | | | | 0.009 |
| | | | | (0.038) | | | | (0.193) |
| Observations | 1,590 | 1,546 | 1,587 | 1,571 | 513 | 513 | 1,546 | 1,571 |

Notes: The dependent variable in columns (1) to (5) is inflation expectations (for the following 12 months) at the time of the original survey (September 2013), with the sample restricted to a subset of respondents who were re-interviewed two months after the original survey (November 2013). The dependent variable in column (6) is inflation expectations (for the following 12 months) at the time of that follow-up interview. The dependent variable in columns (7) and (8) is the expected interest rate (for the following 12 months) in the original survey. 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 – column (1), 763 in the *Products* treatments column (2), 804 in the *Products+Statistics (1.5%)* combined treatment – column (3) – and 788 in the *Hypothetical (10%)* treatment – columns (4) and (8)). The $\alpha$ and $\beta$ coefficients are obtained from the regression given by Equation 4, section 2.2. The p-value of the difference between the $\alpha$ coefficients in columns (1) (*Statistics*) and (2) (*Products*) is 0.0015; the p-value of the difference between the two $\alpha$ coefficients in column (3) (*Statistics+Products*) is 0.0038; and the p-values of the differences between the sum of the $\alpha$ coefficients in column (3) (*Statistics+Products*) and the $\alpha$ coefficients in columns (1) and (2) 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 ($\alpha$), Robustness Checks, *Statistics (1.5%)* and *Products* Treatments, U.S. Online Experiment

| Treatment: | *Statistics* | | | *Products* | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ |
| $\beta$ | 0.829*** | 0.827*** | 0.822*** | 0.777*** | 0.778*** | 0.775*** |
| | (0.057) | (0.057) | (0.059) | (0.051) | (0.051) | (0.051) |
| $\alpha$ | 0.847*** | 0.918*** | | 0.710*** | 0.690*** | |
| | (0.033) | (0.049) | | (0.036) | (0.042) | |
| $\alpha^2$ | | 0.007 | | | -0.003 | |
| | | (0.007) | | | (0.005) | |
| $\alpha_+$ | | | 0.632*** | | | 0.606*** |
| | | | (0.108) | | | (0.078) |
| $\alpha_-$ | | | 0.859*** | | | 0.736*** |
| | | | (0.037) | | | (0.046) |
| Observations | 1,590 | 1,590 | 1,590 | 1,546 | 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), (2) and (3) – and 763 in the *Products* treatments – columns (4), (5) and (6)). The $\alpha$ and $\beta$ coefficients are obtained from the regression given by Equation 4, section 2.2. $\alpha^2$ represents the squared learning weight parameter. $\alpha_+$ and $\alpha_-$ 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,$\left(\pi^T_{i,t} - \pi^0_{i,t}\right)$. The p-values for the differences between the $\alpha_+$ and $\alpha_-$ parameters are 0.0754 for column (3) (*Statistics*) and 0.1985 for column (6) (*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 ($\alpha$), Argentina Online Experiment

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}$ | $\pi_{i,t+1}^{follow-up}$ | $i_{i,t+1}$ | $\log(e_{i,t+1})$ |
| $\beta$ | 1.326*** | 1.133*** | 0.902*** | 0.909*** | 0.902*** | 0.963*** | 0.754*** | 0.155*** | 0.328*** |
| | (0.158) | (0.119) | (0.042) | (0.043) | (0.042) | (0.041) | (0.086) | (0.035) | (0.088) |
| *Statistics* | | | | | | | | | |
| $\alpha$ | 0.506*** | | | | | | | | |
| | (0.090) | | | | | | | | |
| *Products* | | | | | | | | | |
| $\alpha$ | | 0.459*** | 0.494*** | 0.472*** | | 0.456*** | 0.208** | 0.468*** | 0.435** |
| | | (0.062) | (0.027) | (0.025) | | (0.037) | (0.094) | (0.133) | (0.173) |
| $\alpha^2$ | | | | -0.001 | | | | | |
| | | | | (0.001) | | | | | |
| $\alpha_+$ | | | | | 0.484*** | | | | |
| | | | | | (0.040) | | | | |
| $\alpha_-$ | | | | | 0.497*** | | | | |
| | | | | | (0.037) | | | | |
| Observations | 343 | 530 | 3,653 | 3,653 | 3,653 | 1,320 | 1,320 | 3,373 | 1,660 |
| Sample | I | I | II | II | II | II | II | II | II |

Notes: The dependent variable in columns (1) to (6) 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 survey (August 2013). The dependent variable in column (7) is inflation expectations (for the following 12 months) at the time of the follow-up interview. The dependent variable in column (8) is the expected interest rate (for the following 12 months) in the original survey. The dependent variable in column (9) 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 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 columns (1) is the sum of 182 observations in the *Control* group and 161 in the *Statistics 24%* treatment for the college graduates sample. The total number of observations in column (2) is the sum of 182 observations in the *Control* group and 348 in the *Products* (*19%*, *24%* and *29%*) for the college graduates sample. The total number of observations for columns (3)-(5) 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 (6) and (7) 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 (8) 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 (9) 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 $\alpha$ and $\beta$ coefficients are obtained from the regression given by Equation 4, section 2.2. $\alpha^2$ represents the squared learning weight parameter. $\alpha_+$ and $\alpha_-$ 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, $\left(\pi_{i,t}^T - \pi_{i,t}^0\right)$. Robust standard errors. *significant at the 10% level, ** at the 5% level, *** at the 1% level.