APPENDICES: NOT FOR PUBLICATION

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

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

As a final empirical exercise, we illustrate how – due to the substantial dispersion in the distribution of price changes, both in low- and high-inflation contexts – even small limitations in the ability to recall prices can generate substantial dispersion in perceptions about inflation. Denote $p^a_{j,t}$ the actual price of product $j = 1, ..., J$, with corresponding prices changes for j given by $1 + \pi^a_{j,t} = \frac{p^a_{j,t}}{p^a_{j,t-1}}$. One way of modeling memory limitations is to assume individuals have perfect memory about price changes, but they can only recall prices for a limited number of products – a subset $J*$. To estimate the aggregate inflation rate, individuals simply compute the average of price changes for their own basket of $J*$ products. Using our data on actual price changes for supermarket products, we can simulate how these perceptions vary for different values of $J*$.[37] Figure A.2 shows the distribution of annual price changes for $J* = 5$ and $J* = 20$, as well as the distribution of individual inflation expectations for the same time period for the U.S. (panel (a)) and Argentina (panel (b)). This Figure illustrates that even if individuals exhibited a remarkable memory and were able to perfectly recall the current and past prices of 20 products (i.e., 40 individual prices) and correctly compute all changes and their averages, the inflation perceptions resulting from these limited samples would still be substantially dispersed. This evidence complements our finding about the noisiness of individuals' memories about specific prices. Taken together, these two pieces of evidence reinforce the case for a link between memory limitations and the heterogeneity of inflation expectations.

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

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

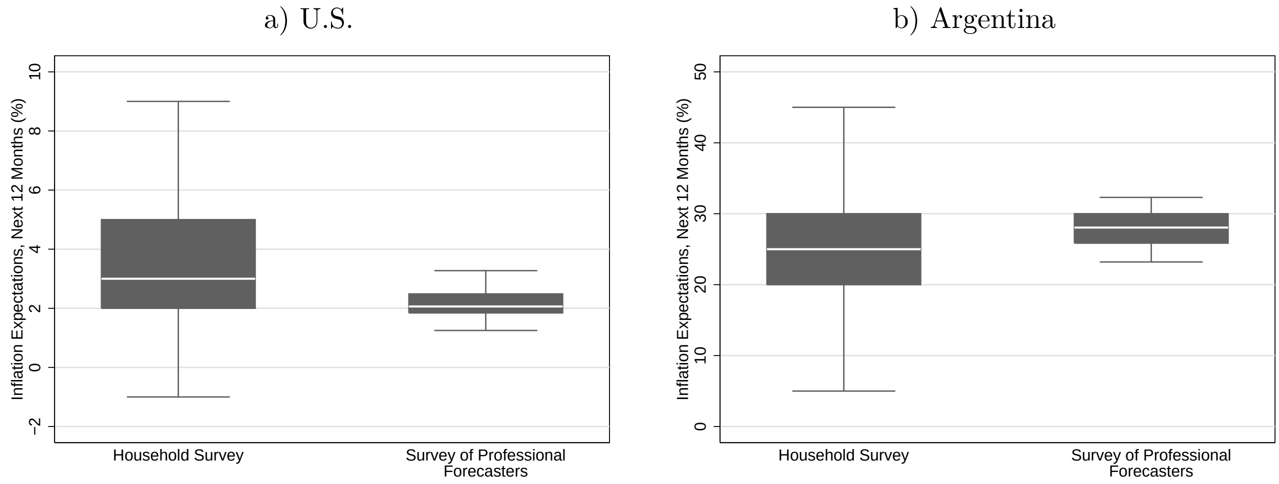

Notes: Expected inflation for the period January 1-December 31 2013, reported in December 2012. Sources: University of Michigan's Survey of Consumers, December 2012 (household survey, U.S., N=502), Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters, fourth quarter of 2012 (professional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

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

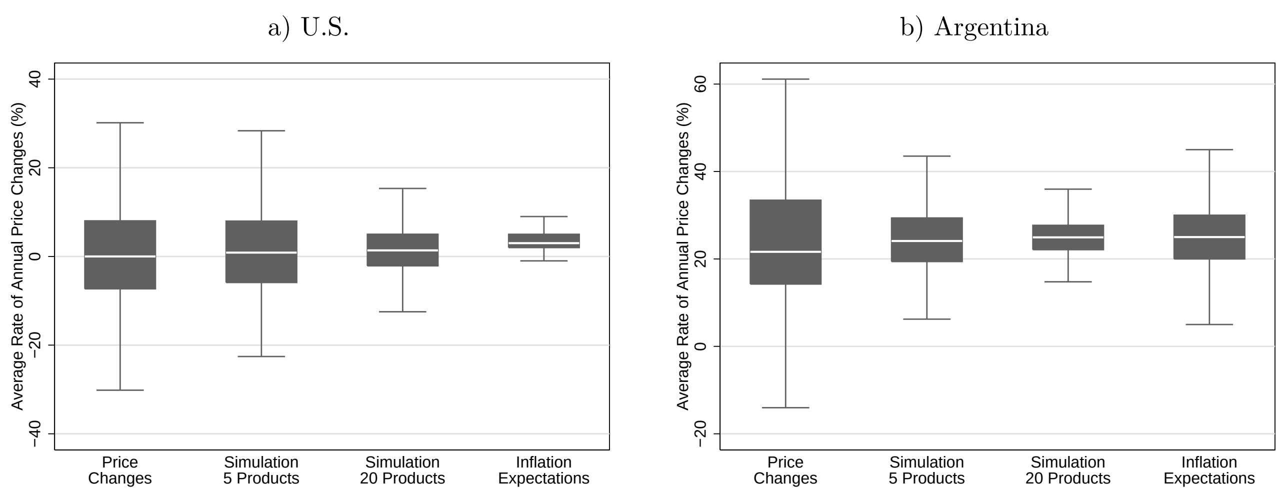

Notes: The price changes refer to the period January 1 2012 to January 1 2013 for both countries. The first box in each panel represents the actual distribution of price changes for the products in each database (N=10,518 and N=9,276 for the U.S. and Argentina, respectively). The following two boxes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

B Descriptive Statistics and Representativeness of the Subject Pools

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

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

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

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

C.1 Subject Pool and Descriptive Statistics

The subject pool for the U.S. online experiment was recruited from Amazon's Mechanical Turk (AMT) online marketplace. We followed several references that describe the best practices for recruiting individuals for online surveys and experiments using AMT, and adopted some of these recommendations to ensure high quality responses.[38]

Potential recruits were offered to participate in a short online "public opinion survey" – we avoided conditioning the subjects by using this vague description and by refraining form using words such as "economic expectations", inflation and others. We collected data during the month of September 2013. Participants were paid \$0.50 for their participation, which is about average for this type of studies in AMT (the average duration of the questionnaire in our sample was about three minutes). We restricted the sample of participants to U.S. residents only,[39] and we included attention checks to ensure participants read the instructions and the questions thoroughly.[40] The descriptive statistics in Table B.1 indicate that, as it is common with this type of studies, subjects in our sample are younger and more educated than the average of the U.S.

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

[38]See for instance:

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

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

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

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

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

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

C.2 Further Results

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

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

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

C.3 Additional Test of Spurious Learning

A key assumption for the test between spurious and genuine learning is that the observational correlation between $\pi_{i,t+1}$ and the outcome variable ($i_{i,t+1}$) reflects a causal effect running from the first to the latter. For other outcomes, denoted $y_{i,t+1}$, the observational correlation with $\pi_{i,t+1}$ may suffer from substantial omitted variable bias. For example, a negative correlation between inflation expectations and expected growth rate could be due to individuals believing that inflation is bad for growth, while a positive correlation could imply that individuals believe in some form of the Phillips curve. Alternatively, that correlation could be entirely spurious, reflecting the fact that more pessimistic individuals expect both higher inflation and lower growth. Holding this pessimism constant, that fact than an individual is induced to believe that inflation is going to be higher in the future should not affect her expectations about growth. As a result, using growth and similar outcomes as dependent variables to estimate α would lead to wildly inaccurate conclusions. Nevertheless, we can still perform a qualitative version of this falsification exercise. For each of these outcomes, we can estimate two versions of the following regression:

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

The first version, labeled as the "experimental correlation," uses the learning equation (4) as the first stage for $\pi_{i,t+1}$ in an 2SLS estimation of (7).[41] Intuitively, this "experimental correlation" provides a measure of how much the outcome $y_{i,t+1}$ changes for every 1 percentage point increase in $\pi_{i,t+1}$ due to provision of information. Ideally, we would like to compare this experimental correlation to the true causal effect of inflation expectations on $y_{i,t+1}$ (i.e., the true δ). We denote the "non-experimental correlation" to the OLS estimate of δ from equation (7) based on subjects in the control group. Even though this non-experimental correlation may be biased with respect to the true δ because of the potential omitted variable biases discussed above, there comparison of the two correlations (the two estimates of δ) can still be informative. If the non-experimental correlations were significantly different from zero for most outcome variables but the experimental correlations were always zero, this would be a strong indication that the learning from the treatments is spurious rather than genuine. As a result, this would provide a quantitatively rather than a quantitatively test of spurious vs. genuine learning.

Figure C.5 presents these correlations for a series of additional standardized outcomes.[42] All the outcomes were constructed such as the expected correlation with inflation is positive (e.g., higher inflation should be correlated to higher interest rate). To increase the statistical power of these regressions, we pooled the three factual information treatments – the experimental correlations are statistically the same for these three treatments (see the Appendix for an illustration with the nominal interest rate). As expected, the observational correlations for the outcomes presented in Figure C.5 are all positive and significant at standard confidence levels. The experimental correlations are also positive in general, suggesting that a substantial portion of the learning was genuine. The experimental correlations, however, are lower – on absolute value – than the observational correlations. This is probably due to a combination of two factors: i. Some spurious learning; ii.

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

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

Omitted-variable biases in the observational correlations. The results are thus consistent with the result presented in body of the paper that there is some spurious learning but a majority of the learning is still genuine.

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

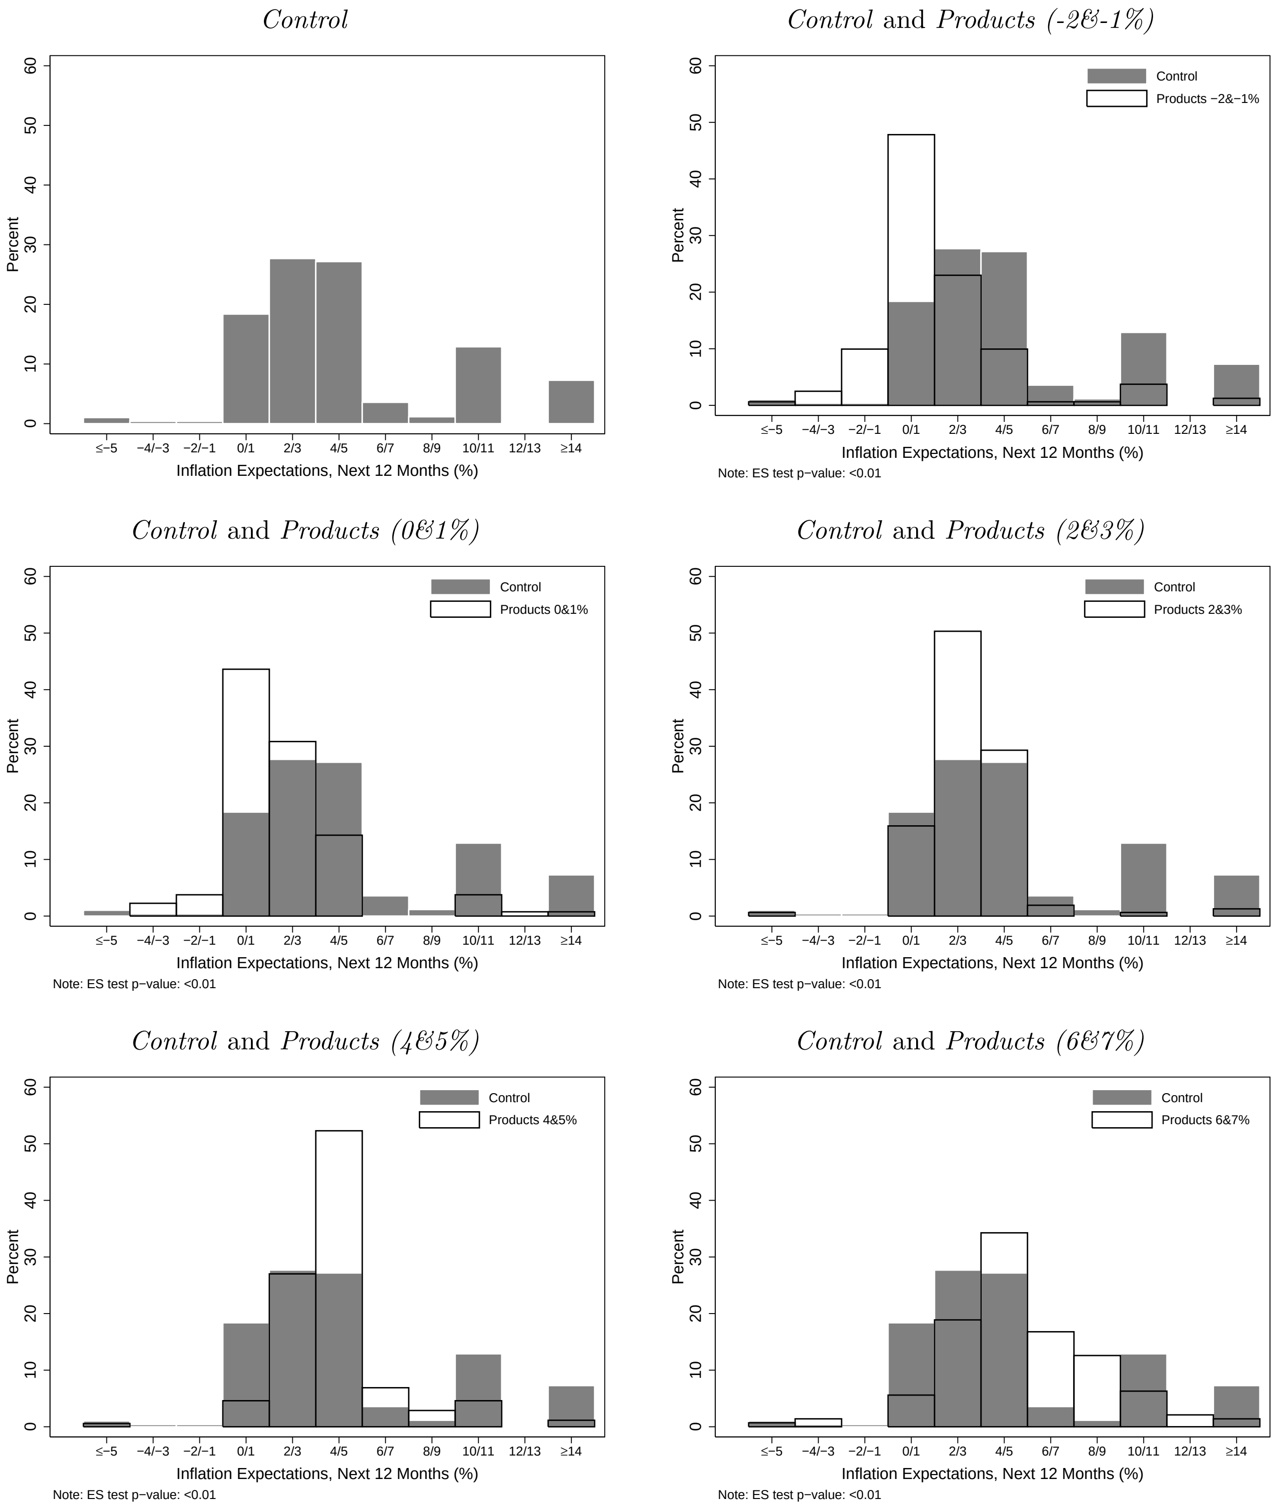

Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. ES is the Epps–Singleton characteristic function test of equality of two distributions.

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

Control

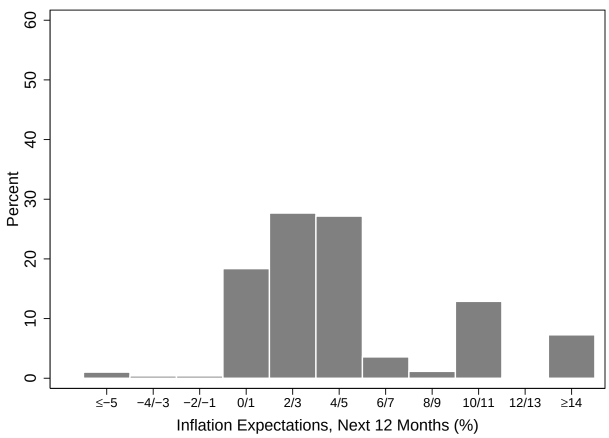

Control and *Statistics(1.5%)+Products (-2&-1%)*

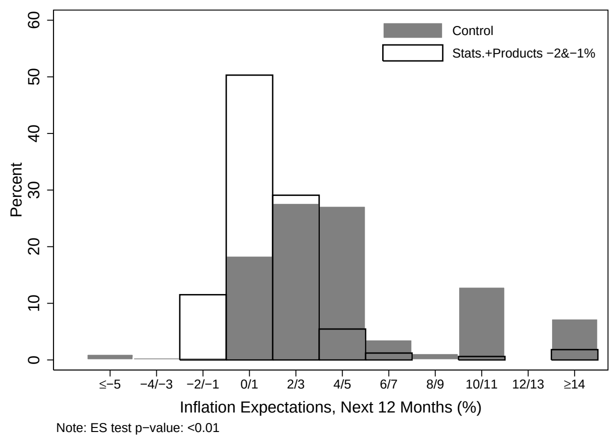

Control and *Statistics(1.5%)+Products (0&1%)*

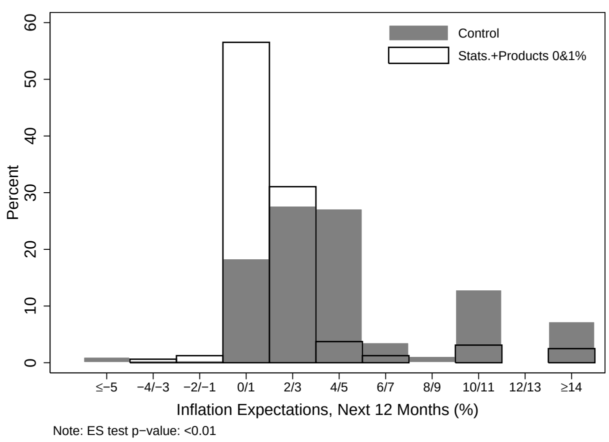

Control and *Statistics(1.5%)+Products (2&3%)*

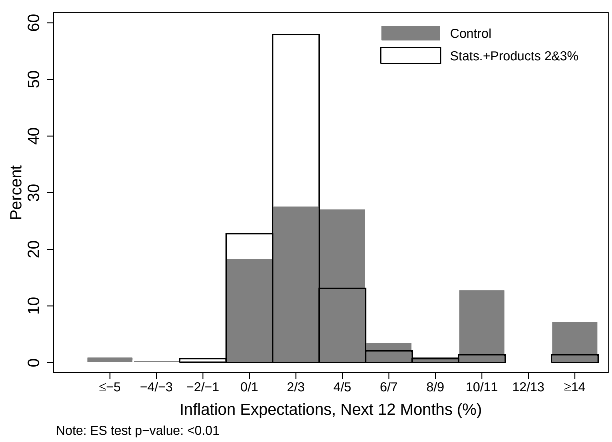

Control and *Statistics(1.5%)+Products (4&5%)*

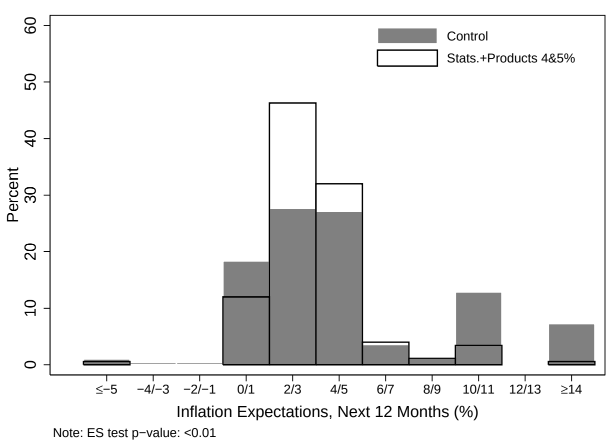

Control and *Statistics(1.5%)+Products (6&7%)*

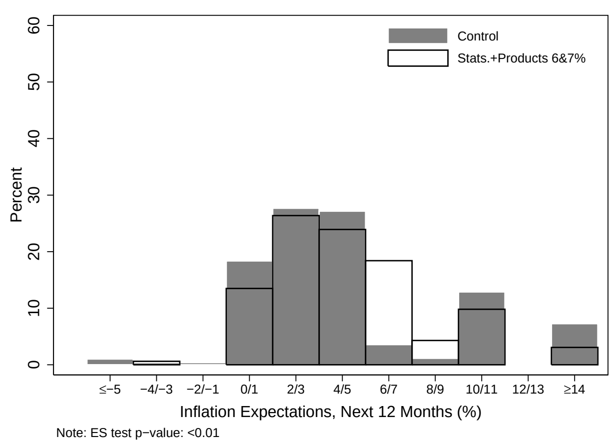

Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. ES is the Epps–Singleton characteristic function test of equality of two distributions.

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

a) Effect on inflation expectations

b) Effect on confidence

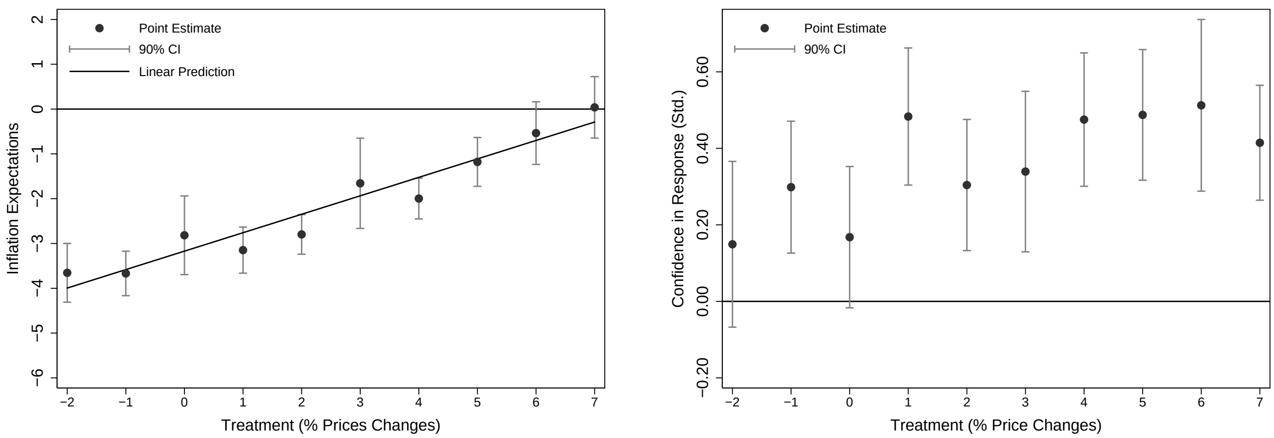

Notes: The total number of observations is 1,732 (789 in the control group and 804 in the 10 variations of the combined specific prices and official statistics treatment). Each bar represents the point estimate of the effect of the specific price treatment compared to the control group. Robust standard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

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

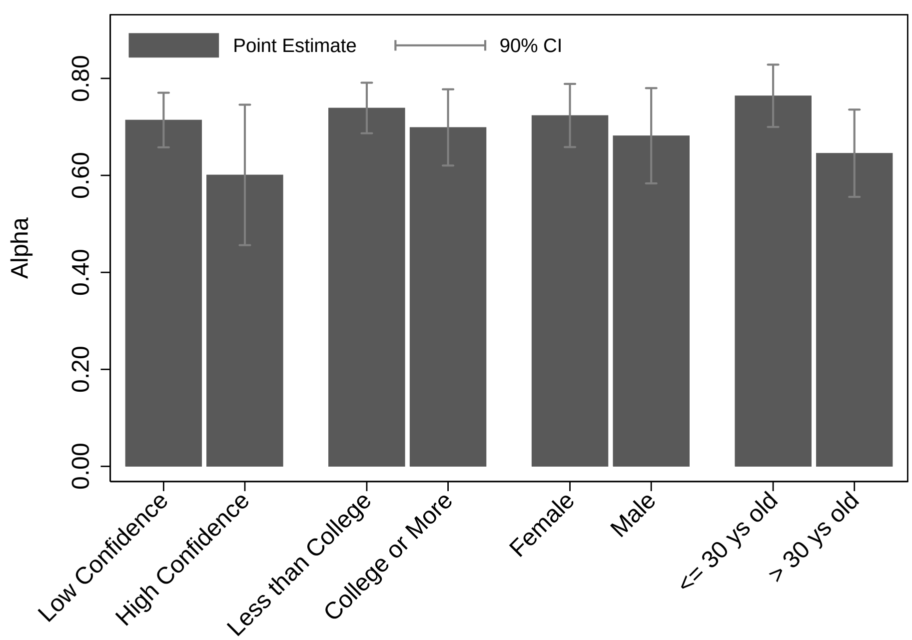

Notes: Results for the specific prices-tables treatment only. The total number of observations is 1,552 (789 in the control group and 763 in the 10 variations of the *Products* treatment). Robust standard errors reported.

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

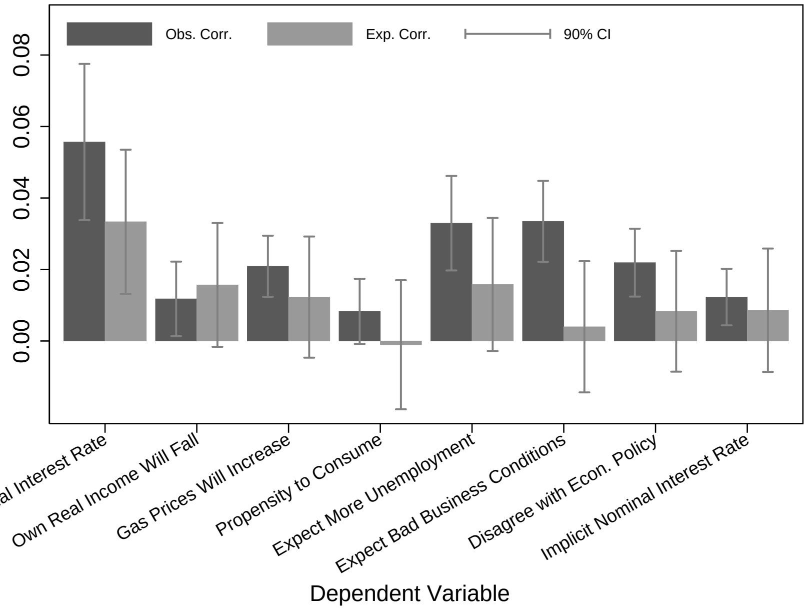

Notes: The total number of observations is 3,157 (control group and all treatments except *Hypothetical (10%)*). The observational correlations correspond to the coefficient of inflation expectations in OLS regressions of the dependent variables on inflation expectations for the *Control* group. The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

D Argentina Online Experiment: Further Details and Results

D.1 Samples

The Argentina online experiment results are drawn from two different sets of respondents. The first group is comprised by a sample of economics, accountancy, business and political science graduates. This sample, with a total of 691 observations, was assigned to a control group, or to *Statistics (24%)*[43] and *Products* treatment arms, the latter with three sub-treatments with tables with average price changes of 19%, 24% and 29%. This experiment was implemented between May and June 2013 using only graduates in economics, management, accountancy, finance, international relations and political science from Argentina. We approached these subjects through mailings of graduates from the Universidad Nacional de La Plata (UNLP), Universidad Torcuato Di Tella (UTDT), and through a professional association, the Consejo de Profesionales en Ciencias Económicas of the Buenos Aires province (CPBA). About half of theindividuals contacted responded to the survey resulting in a total sample of 691 respondents. Of those, 277 were accountants, 135 had a BA or MA in Economics, 89 a BA in Management, 57 an MBA or an MA in Finance, and the rest were Political Scientist and Bachelors in International Relations. All of these individuals had at least basic Economics training as part of their degrees.

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

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

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

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

D.2 Construction of the Informational Treatments

As in the U.S. experiment, our *Products* information provision setup consisted of displaying tables with the prices and price changes of specific products after eliciting past inflation perceptions and right before asking about respondents' inflation expectations. In the context of the Argentine experiment, we displayed a series of 19 different tables with four products each, with average price changes over the previous year (March 1 2012 to March 1 2013) ranging from 16 to 34% in one percentage point increments (see Appendix E.3 for the screen captures of the full questionnaire and for all the specific product tables).[44] The source for these tables is a database of scrapped online data from the largest supermarket chain in Argentina, and the products correspond to a subsample of four common products: olive oil, pasta, wine, and shampoos/conditioners. As in the U.S. experiment, no suggestion was made that the prices or the price changes shown in the table were representative, and that there was no deception. The text only stated that the products were selected randomly, without specifying any details about the sampling procedure.

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

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

increased its price substantially more than the brand in the 30% table. The 750ml bottles of wine in the two tables also have a similar initial price, but the price increase of the Malbec in the 30% table was much larger than that of the Syrah. The tables were introduced with the following text: "Before replying, please take a look at the following table. For each of the listed products, the table presents the price on March 1, 2012 and March 1, 2013 (that is, one year later). These prices were taken from the same branch from the main supermarket chain in Argentina". It should be noted that no suggestion was made that the prices or the price changes shown in the table were representative, and that there was no deception. The text only stated that the products were selected randomly, without specifying any details about the sampling procedure.

We implemented a shorter version of the questionnaire-experiment for the sample of college graduates (see Appendix E.3 for the screen captures of the full questionnaire). The experiment had the same structure as the previous ones, and a subset of the outcomes from the larger sample Argentina experiment described above. In terms of treatments, we included three tables with specific prices (with the same format as in Figure D.1, but with dates updated accordingly – see Appendix E.3 for all the original tables included in the experiment), with average price changes of 19%, 24% and 29%. We also included a fourth treatment branch, where instead of a table, we included the following statement: "According to an average of non-official indicators produced by private firms, analysts and research centers, the annual rate of inflation with respect to the last 12 months was approximately 24%".[45]

D.3 Further Analysis

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

The Argentina opinion poll sample also allowed for a more detailed analysis of heterogeneous effects in learning. The coefficients of the learning model in Table 3 may also have different parameter values for different groups. Figure D.3 presents some evidence for differences in α between relevant groups in the population. The first two columns in the Figure present the coefficients for those with high and low levels of confidence in their inflation perceptions. In contrast

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

to the results for the U.S., we find significant differences between the two groups: individuals who reported lower levels of confidence on their own perceptions of inflation placed a significantly higher weight on the information we provided (about 0.61 compared to about 0.41). There are also similar and significant differences by education level and by age: respondents with less than a college degree and those under 40 years old place a higher weight on the information provided as part of the treatment. Females (with respect to males) also seem to learn more from the informational treatments, although this difference is not statistically significant.

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

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

a) *Products (16%)*

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

b) *Products (30%)*

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

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

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

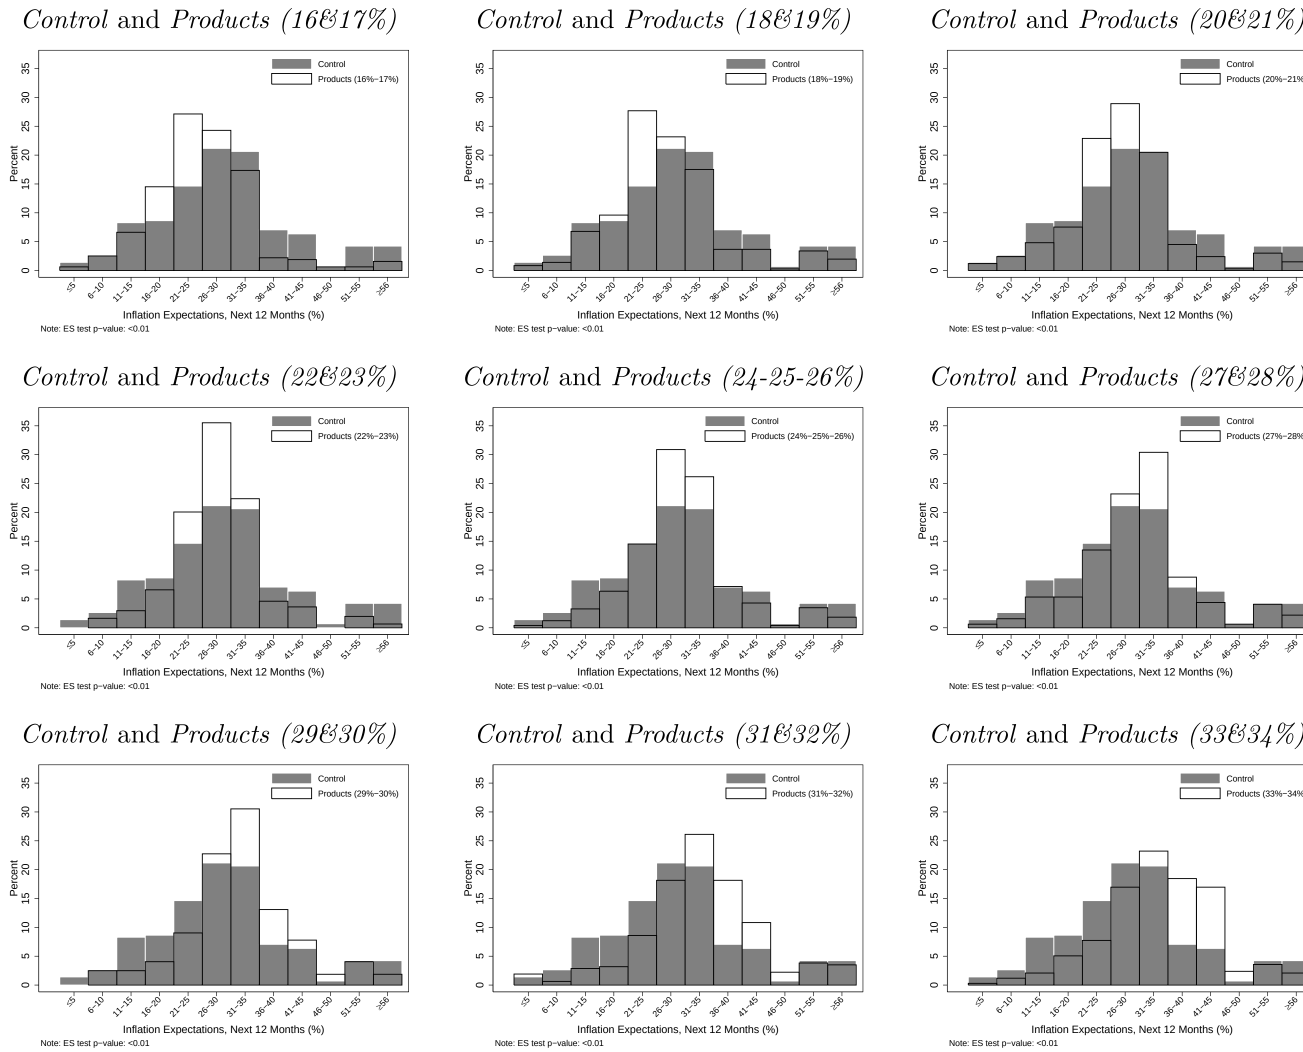

Notes: The total number of observations is 3,686, with 568 in the control group and 146-181 in each of the 19 treatment groups. ES is the Epps–Singleton characteristic function test of equality of two distributions.

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

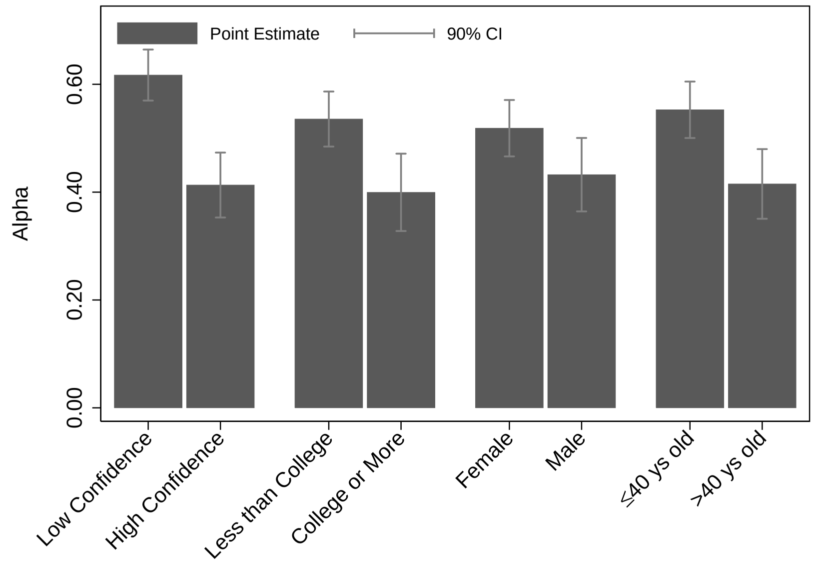

Notes: The total number of observations is 3,686. Robust standard errors reported.

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

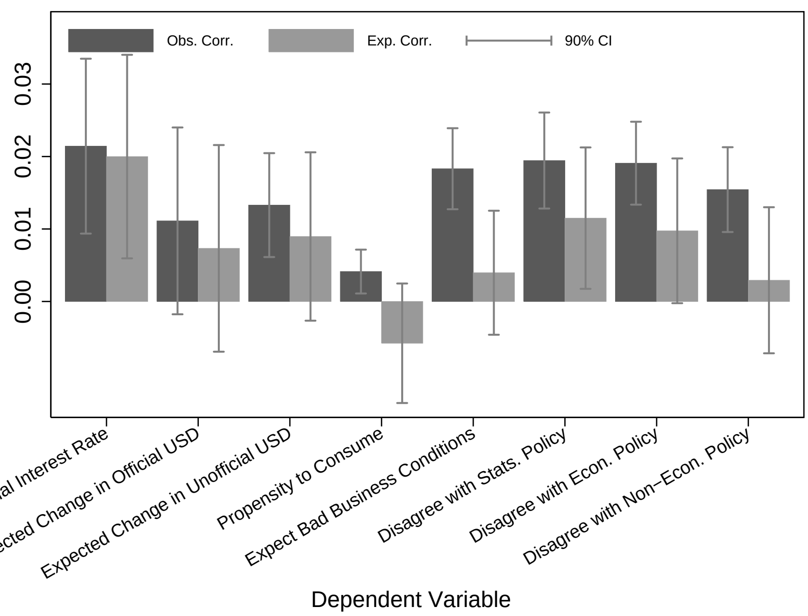

Notes: The total number of observations is 3,686. The observational correlations correspond to the coefficient of inflation expectations in OLS regressions of the dependent variables on inflation expectations for the *Control* group. The experimental correlations correspond to IV versions of the same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

E Argentine Supermarket Experiment

E.1 Data Collection

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

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

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

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

Figures E.1 and E.2 present robustness checks of the results in the main body of the paper. The previous results where based on actual and remembered price changes for products the respondents had just purchased. The survey, however, also recorded the total amount spent, and asked the respondents about their estimate of the total they would have had to pay for the same goods 12 months earlier.

The results presented in this Appendix are not based on these remembered price changes. Instead, they compare the distribution of inflation expectations (Figure E.1) for individuals for high and low remembered and actual changes in their purchase receipts total amount. Figure E.2 in turn depicts the relationship between the price changes in the receipt and inflation expectations (panel (a)), as well as the relationship between price changes in the receipt (actual and remembered).

E.3 Estimating Learning Rates

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

Table E.1 presents the estimates from the learning model described in section 2.2 for our supermarket study. The first randomly assigned information for which we compute the learning model is the average remembered price change for the four products that the respondent was asked about.[46] The α coefficient is about 0.11 and strongly significant. This weight is substantially lower than the one obtained from the online experiments (about 0.5 for Argentina), but this was expected due to the reasons listed above due to the reasons listed above. This implies that individuals form their inflation expectations, in part, based on information that is mostly noise (i.e., it is not correlated with actual price changes – see Figure 8, panels (c) and (d)), as we established previously. To stress this point, in column (2), instead of using remembered price changes, we use the actual price changes in the list of randomly selected products. As expected, the estimated α is close to zero and statistically insignificant. In column (3), we present the estimates from the replication of the *Products* treatment with the three levels discussed in the previous paragraph.

[46] Given the biases documented above in terms of the average price changes reported by respondents, in these regressions we use a "corrected" value using a deflation factor of 30%. In any case, the results are similar under alternative specifications.

The α coefficient, which represents the weight given by respondents to the price information we provided, is similar in value to the α for (salient) remembered prices (although it is statistically insignificant. The last column (4) in the table pools all these alternative treatments, and the results are very similar.

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

a) Low and high remembered total purchase amount change

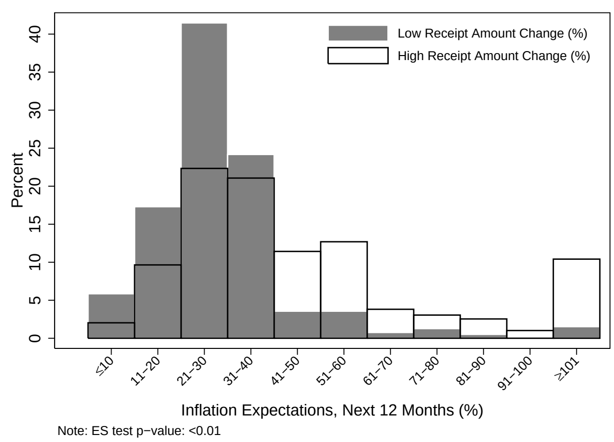

b) Low and high actual total purchase amount change

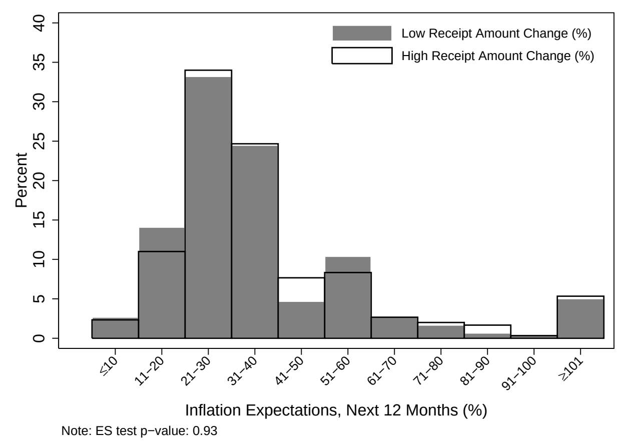

Notes: The total number of observations is 375 (lowest third of total purchased amount changes, panel (a)) and 372 (top third of of total purchased amount changes, panel (b)). ES is the Epps–Singleton characteristic function test of equality of two distributions.

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

a) Remembered total purchase amount changes (%) and inflation perceptions

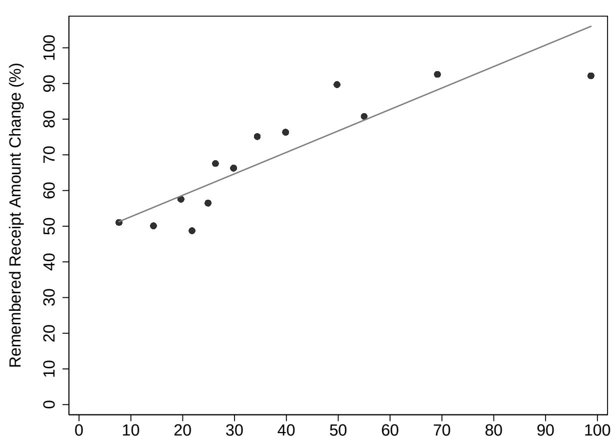

b) Annual total purchase amount changes (%): Actual and remembered

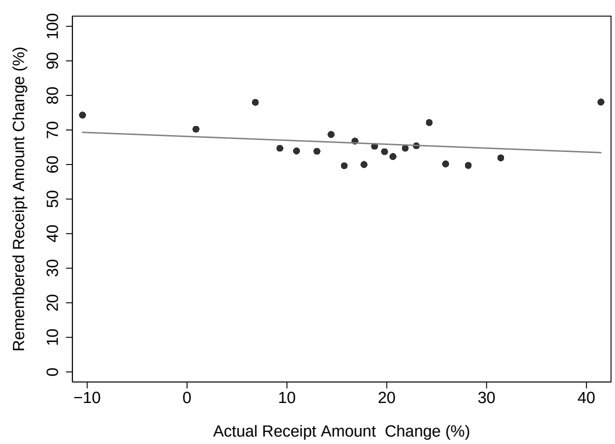

Notes: The total number of observations is 1,140. Panels (a) and (b) represent binned scatterplots. The percentage changes in both panels are implicit – they are obtained from the total purchase amounts in pesos (AR$) from the scanned receipt and from the estimate of the total for the same purchase a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

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

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

Notes: The total number of observations correspond to 1,070 participants of the Argentina Supermarket Experiment with valid responses for inflation expectations and remembered price changes, and for which it was possible to establish the actual price changes from the scanned purchase receipts (actual price changes). The α and β coefficients are obtained from the regression given by Equation 4, section 2.2. The p-value of the difference between the α coefficient for *Remembered Price Changes* and *Actual Prices Changes* is 0.0102. Robust standard errors. *significant at the 10% level, ** at the 5% level, *** at the 1% level.