

Online Appendix

Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments

By Alberto Cavallo, Guillermo Cruces and Ricardo Perez-Truglia

A Memory Limitations and Excess Dispersion in Inflation Expectations

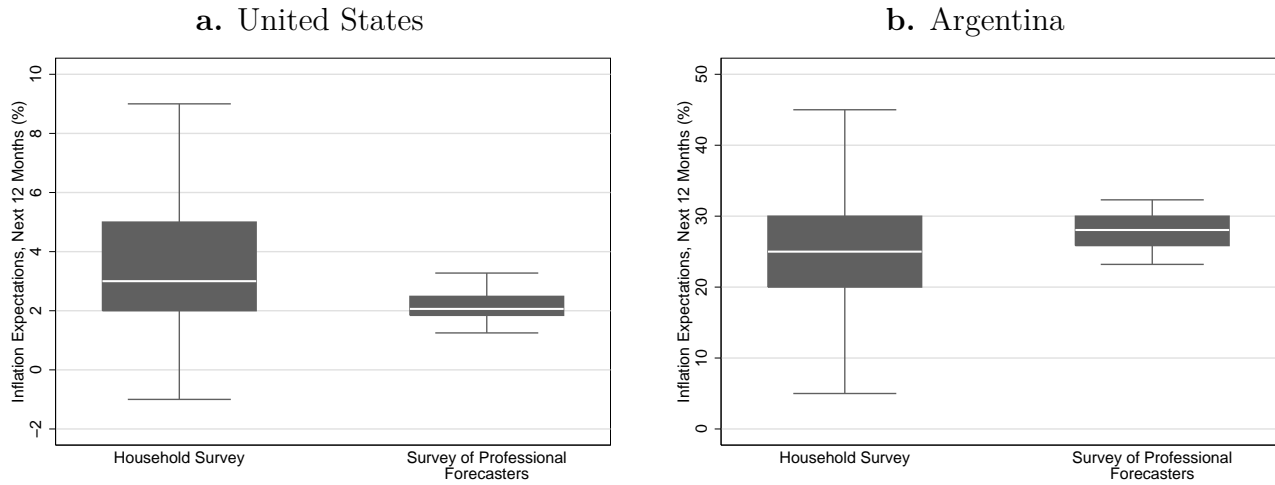
As discussed in the main body of the paper, memory limitations might induce excess dispersion in inflation expectations. In this section, we present some evidence consistent with this hypothesis. Figure A.1 presents the distribution of inflation expectations for 2013, reported at the end of 2012 by individuals in the household surveys and by professional forecasters in surveys of professionals. 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 United States, 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 (i.e., the use of price memories in forming inflation expectations) could explain a small or a large share of excess dispersion in inflation expectations. The evidence suggests that it can explain a large share of this dispersion. Our results indicate that individuals assign a significant weight to the price changes of individual products, and this is further reinforced by our finding of a nearly-orthogonal relationship between remembered price changes and actual price changes.

As a final empirical exercise, we illustrate how, due to the substantial dispersion in the distribution of price changes, both in low- and high-inflation contexts, even small limitations in the ability to recall prices can generate substantial dispersion in perceptions about inflation. Denote $p_{j,t}^a$ the actual price of product $j = 1, \dots, J$, with corresponding price changes for j given by $1 + \pi_{j,t}^a = \frac{p_{j,t}^a}{p_{j,t-1}^a}$. One way of modeling memory limitations is to assume individuals have perfect memory about price changes, but they can only recall prices for a limited number of products, i.e., 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^* .³⁵ 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

³⁵The 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.

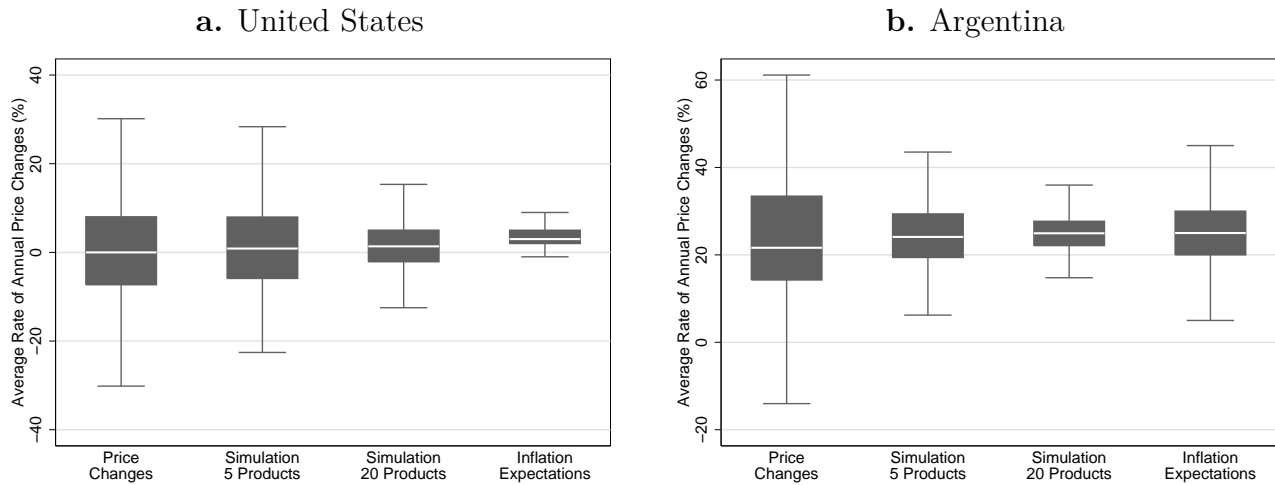
memory and were able to perfectly recall the current and past prices of 20 products (i.e., 40 individual prices) and correctly computed all changes and their averages, the inflation perceptions resulting from these limited samples would still be substantially dispersed. This evidence complements our 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.

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: Panel a: 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). Panel b: WP Public Opinion Survey, December 2012 (household survey, Argentina, N=777) and Latin Focus 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 Online Experiments: Further Details and Results

B.1 Further Details about Data Collection and Descriptive

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.³⁶

Potential recruits were offered to participate in a short online “public opinion survey.” We avoided conditioning the subjects by using this vague description and by refraining from using words such as “economic expectations”, “inflation” and others. We collected data during the month of September 2013. Participants were paid \$0.50 for their participation, which is about average for this type of studies in AMT (the average duration of the questionnaire in our sample was about three minutes). We restricted the sample of participants to U.S. residents only,³⁷ and we included attention checks to ensure participants read the instructions and the questions thoroughly.³⁸ The descriptive statistics in the top panel of Table B.1 indicate that, as it is common with this type of studies, subjects in our sample are younger and more educated than the average of the U.S.

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

³⁶See 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.

³⁷While 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.

³⁸All 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.

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

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

The bottom panel in 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 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.

B.2 Further Details about the Information Treatments

This Section complements the discussion of the U.S. online experiment in the body of the paper (Section 3.1) by presenting some additional details about the information treatments. Figure 1

presents examples of the treatment arms in the U.S. online experiment, and Figure B.1 presents equivalent examples for Argentina. Our information provision setup consisted of displaying tables with the prices and price changes of specific products. In the context of the Argentine experiment (sample II), 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 B.1, and Appendix F for the screen captures of the full questionnaire and for all the specific product tables). 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) and 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 as much constant 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 (Figures B.1.a and B.1.b respectively). The descriptions are identical, the initial prices are very similar, but the price changes of the two olive oils are very different: the brand in the *Products (30%)* table increased its price substantially more than the brand in the *Products (16%)* 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 for the screen captures of the full questionnaire). The experiment had the same structure as the previous ones, and we collected information on 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 B.1, but with dates updated accordingly. See Appendix E 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 unofficial indicators pro-

duced by private consultancy firms, analysts and research centers, the annual inflation rate in the last 12 months was approximately 24%.” The original sentence in Spanish and the corresponding English translation are presented in Figure B.1.³⁹

B.3 Further Results

This section complements the discussion in the body of the paper by presenting some of the main results in more detail, and also by discussing some additional results.

B.3.1 Reduced Form Evidence

Figure 4 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 B.2 (*Products*) and B.3 (*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 treatments, grouped in two one percentage point sets. The main results 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 it to the right.

We can also appreciate the effects of the treatments by testing the impact on average outcomes. In the body of the paper, panel (a.i) in Figure 6 depicted the effect of the *Product* treatments on the average of inflation expectations, and panel (a.ii) in the same Figure compares the impact of each treatment level for the *Products* treatment arm on the standardized confidence variable. Figure B.4 reproduces the equivalent results for different levels of the combined *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 B.4 confirms that the impact of the treatments with specific products modified the average reported expectations in a systematic manner, with the impact increasing in the value of the signal. Regarding the effects on confidence, we cannot reject the null hypothesis that all these sub-treatments had the same effect on confidence (p-value of 0.16). The coefficients for the -2% and 0% signals still have non-significant effects on confidence, but we can reject the null hypothesis that the coefficients for -2%, -1% and 0% are jointly insignificant (p-value of 0.02). As in the discussion of Figure 6 (a.ii), individuals might be less prone to incorporate information about

³⁹After the intervention of the national statistical agency in 2007 and the adulteration of official inflation estimates, the government started prosecuting private sector firms and consumer associations who published their own measures of inflation as an alternative to the official statistics. For this reason, members of Argentina’s Congress (who had immunity from prosecution) started compiling in 2012 these private sector estimates confidentially and reported the mean every month as the “IPC Congreso”. Our survey coincided with the April 2013 release of this indicator, with an annual inflation rate of 23.67%. See Cavallo (2013) and Cavallo, Cruces and Pérez-Truglia (2016) for more details.

price decreases than about price increases, although in this case even the negative signals seem to have a significant effect on confidence on inflation expectations, similar to that of the positive signals. Overall, then, 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 evidence in panel (b) of Figure B.4 allows for an additional test. Since we have a situation where the treatment provides a signal for aggregated inflation (1.5%) and information about price changes for concrete products, the two signals disagree for some of the sub-treatments in the the *Statistics (1.5%)+Products* treatment arm. We can test whether when the two signals coincide consumers have more confidence in their forecast. When the product price change is between 1% and 2%, we can consider that the signals “agree”. In our Bayesian model, the gain in confidence should be the same no matter whether two signals drawn from the same distributions are similar or very different. The evidence discussed in the previous paragraph is consistent with this prediction: we cannot reject the null that all the 10 coefficients are equal (p-value of 0.16), and also we cannot reject the null hypothesis that the “agreeing” sub-treatments (1% and 2%) have the same effect on confidence compared to all others (p-value of 0.62). If anything, as in the *Products* treatment discussed in the body of the paper, there is a suggestive difference when comparing the positive signals (1%-7%) against the non-positive signals (-2%-0%), but that difference is most likely due to asymmetry than to agreeing with the prior beliefs.

We also include in this Appendix the complete pattern of distribution of inflation expectations for the different treatments in the Argentina online experiment. Figure B.5 presents the results for all the treatments in the Argentina college graduates sample (I), and Figure B.6 depicts the results for the Argentina opinion poll sample (II), with two or three *Products* treatment levels per panel. The results presented in the two Figures confirm the main paper's result that 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 the treatments (price changes between 22% and 26%) for sample II reduced the dispersion of expectations more than they affected the mean.

B.3.2 Learning Model

We also present here additional evidence and robustness checks on our estimates of the learning model. We first analyze the potential implications of sample selection in our survey for our results. The discussion of Table B.1 in the previous section indicated substantial differences between our online experiment samples and the general population of Argentina and the United States. Our first robustness check is to reproduce the paper's main results from Table 1 using sampling weights. We constructed these weights to make the online survey data representative of the whole country in terms of age, gender balance and education level for both Argentina and the United States.

They are based on population data for both countries, and adjusted for the combined proportion in the population of males and females from three age groups and three groups sorted by education levels.⁴⁰

The discussion of the evidence presented in Panel (a) of Figure 7 indicated low heterogeneity in learning rates along socio-demographic categories for the United States, and this is confirmed in the comparison of the results from unweighted (Panel a, Table 1) and weighted (Panel a, Table B.2) regressions. The coefficients for the pass-through and the learning rates are in very similar ranges in the two tables.

The heterogeneity of learning rates with respect to demographic characteristics is somewhat more significant in Argentina (Panel b, Figure 7). However, the weighted and unweighted results are nonetheless similar in Argentina both for samples I (college graduates) and II (opinion poll, general population). One notable difference is that the learning rate using the follow-up survey decreases from 0.208 in the unweighted results to 0.092 in the weighted results (column 4, Panel b, in Tables 1 and B.2 respectively). However, we cannot reject the null hypothesis that these two coefficients are equal at standard significance levels. In sum, weighting the observations does not affect the overall pattern of results.

We also conduct further tests of the Bayesian model described in section 2.3. We first test for non-linearities or asymmetries in the reaction to the information provided (e.g., if individuals learn more from signals that are closer to their prior belief). Our learning model predicts that an individual's adjustment to the new information is a linear function of the distance between the new information and her prior belief. We can test whether this prediction is accurate by estimating the basic model including an additional quadratic term:

$$\pi_{i,t+1} = \gamma_0 + \gamma_1 \pi_{i,t}^0 + \gamma_2 (\pi_{i,t}^T - \pi_{i,t}^0) + \gamma_3 (\pi_{i,t}^T - \pi_{i,t}^0)^2 + \varepsilon_{i,t+1}$$

and testing whether $\hat{\gamma}_3 = 0$. Similarly, we can test the possibility that individuals react differently to signals above their prior belief than to signals below their prior, by estimating the following model:

$$\pi_{i,t+1} = \gamma_0 + \gamma_1 \pi_{i,t}^0 + \gamma_+ \cdot 1 \{ \pi_{i,t}^T > \pi_{i,t}^0 \} \cdot (\pi_{i,t}^T - \pi_{i,t}^0) + \gamma_- \cdot 1 \{ \pi_{i,t}^T < \pi_{i,t}^0 \} (\pi_{i,t}^T - \pi_{i,t}^0) + \varepsilon_{i,t+1}$$

and then testing whether $\hat{\gamma}_- = \hat{\gamma}_+$.

The results from these additional tests are presented in Table B.3 for the U.S. Online Experiment and in Table B.4 for Argentina's sample II. For the U.S., the alternative specification with a quadratic term is provided in columns (1) and (3) of Table B.3 for the *Statistics (1.5%)* and

⁴⁰For Argentina's sample I (college graduates), which is not representative of the whole population, we adjust for three age groups of college graduates and for the proportion of college graduates with a postgraduate degree in the population, which are over-represented in our sample.

Products treatments, respectively. The results indicate that the linear terms for α and β are very similar to the main results without the quadratic term in Panel (a) in Table 1, while the coefficients for the quadratic terms in columns (1) and (3) are not statistically significant and virtually equal to zero (0.007 and -0.003, respectively). Columns (2) and (4) present the results yielded by a specification that allows differential learning for positive and negative differences between the signal and the prior belief, with a coefficient α of 0.632 (*Statistics*) and 0.606 (*Products*) for those with $\pi_{i,t}^T - \pi_{i,t}^0 \geq 0$, and of 0.859 and 0.736 for those with $\pi_{i,t}^T - \pi_{i,t}^0 < 0$. The difference between the two pairs of coefficients is statistically significant for the *Statistics* treatment (p-value of 0.08) but not for the *Products* treatment (p-value of 0.22). Thus, there is some weak evidence of a mild asymmetry in our U.S. sample, indicating that individuals seem more prone to revise their expectations downwards rather than upwards.

These alternative specifications for Argentina (sample II) are presented in Table B.4. The linear terms for α and β with a quadratic term presented in column (2) are very similar to the benchmark (linear only) results presented in column (1), while the coefficient for the quadratic term is not statistically significant and virtually zero (-0.001). Column (3) in Table B.4, in turn, presents the results of an alternative specification that contemplates differential learning for upward and downward corrections of the prior beliefs. The estimated coefficient α is 0.484 for those with $\pi_{i,t}^T - \pi_{i,t}^0 \geq 0$ and of 0.497 for those with $\pi_{i,t}^T - \pi_{i,t}^0 < 0$, and their difference is not statistically significant. This evidence suggests that learning was symmetric in our Argentina experiment, as predicted by the Bayesian model. This result contrasts with the evidence in the U.S. sample, where we found some limited but statistically significant evidence of a mild asymmetry. Overall, this evidence also suggests that the Bayesian model fits the data very well.

Finally, in column (4) of Table B.4, we report the results from the estimation of learning rates using the log of the expected nominal exchange rate of the Argentine Peso with respect to the U.S. Dollar on the free currency market instead of inflation or interest rate expectations. This is a key macroeconomic variable in Argentina: due to a history of high inflation, a substantial fraction of savings are held in U.S. dollars, so most individuals are aware of the market value of this exchange rate and have interest in its future evolution. The α coefficient from this estimation, presented in column (8), is 0.435, that is, very close to the figure for the nominal interest rate (Table 1, panel b, column 5 – 0.468) and for inflation expectations (Table B.4, column 1 – 0.494). This result further confirms the notion that individuals incorporate the information about prices when forming their perceptions of all relevant nominal variables in the economy.

B.4 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 emerge because individuals believe that inflation is bad for growth, while a positive correlation could imply that individuals believe in some form of the Phillips curve. Alternatively, that correlation could be entirely spurious, reflecting the fact that more pessimistic individuals expect both higher inflation and lower growth. Holding this pessimism constant, that fact that an individual is induced to believe that inflation is going to be higher in the future should not affect her expectations about growth. As a result, using growth and similar outcomes as dependent variables to estimate α would lead to wildly inaccurate conclusions. Nevertheless, we can still perform a qualitative version of this falsification exercise. For each of these outcomes, we can estimate two versions of the following regression:

$$y_{i,t+1} = \alpha + \delta\pi_{i,t+1} + \varepsilon_{i,t+1} \tag{B.1}$$

The first version, labeled as the “experimental correlation,” uses the learning equation (6) as the first stage for $\pi_{i,t+1}$ in an 2SLS estimation of (B.1).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 (B.1) 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, the comparison of the two correlations (the two estimates of δ) can still be informative. If the non-experimental correlations were significantly different from zero for most outcome variables but the experimental correlations were always zero, this would be a strong indication that the learning from the treatments is spurious. This would provide a qualitative rather than a quantitative test of spurious vs. genuine learning.

Panel (a) in Figure B.7 presents these correlations for a series of additional standardized outcomes for our U.S. online experiment.⁴² All the outcomes were constructed such that 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. The observational correlations for the outcomes presented in Figure B.7 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

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

⁴²The categorical dependent variables presented in Figure B.7 (all but the nominal interest rate, the propensity to consume and the perceived interest rate) were rescaled and standardized according to the Probability-OLS procedure described in Van Praag and Ferrer-i-Carbonell, “Happiness Quantified: A Satisfaction Calculus Approach,” Oxford: Oxford University Press, 2007.

to a combination of two factors: i. Some spurious learning; ii. Omitted-variable biases in the observational correlations.

Finally, as in the U.S. online experiment, the Argentina online experiment 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. These results for the opinion poll sample (II) are summarized in Panel (b) in Figure B.7. 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 B.1: Example of Information Treatments (English Translation), 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%


c) *Statistics (24%)*

De acuerdo a un promedio de los indicadores no oficiales realizados por consultoras privadas, analistas y centros de estudios, la tasa anual de inflación con respecto a los últimos 12 meses fue aproximadamente de 24%.

Translation: *According to an average of unofficial indicators produced by private consultancy firms, analysts and research centers, the annual inflation rate in the last 12 months was approximately 24%.*

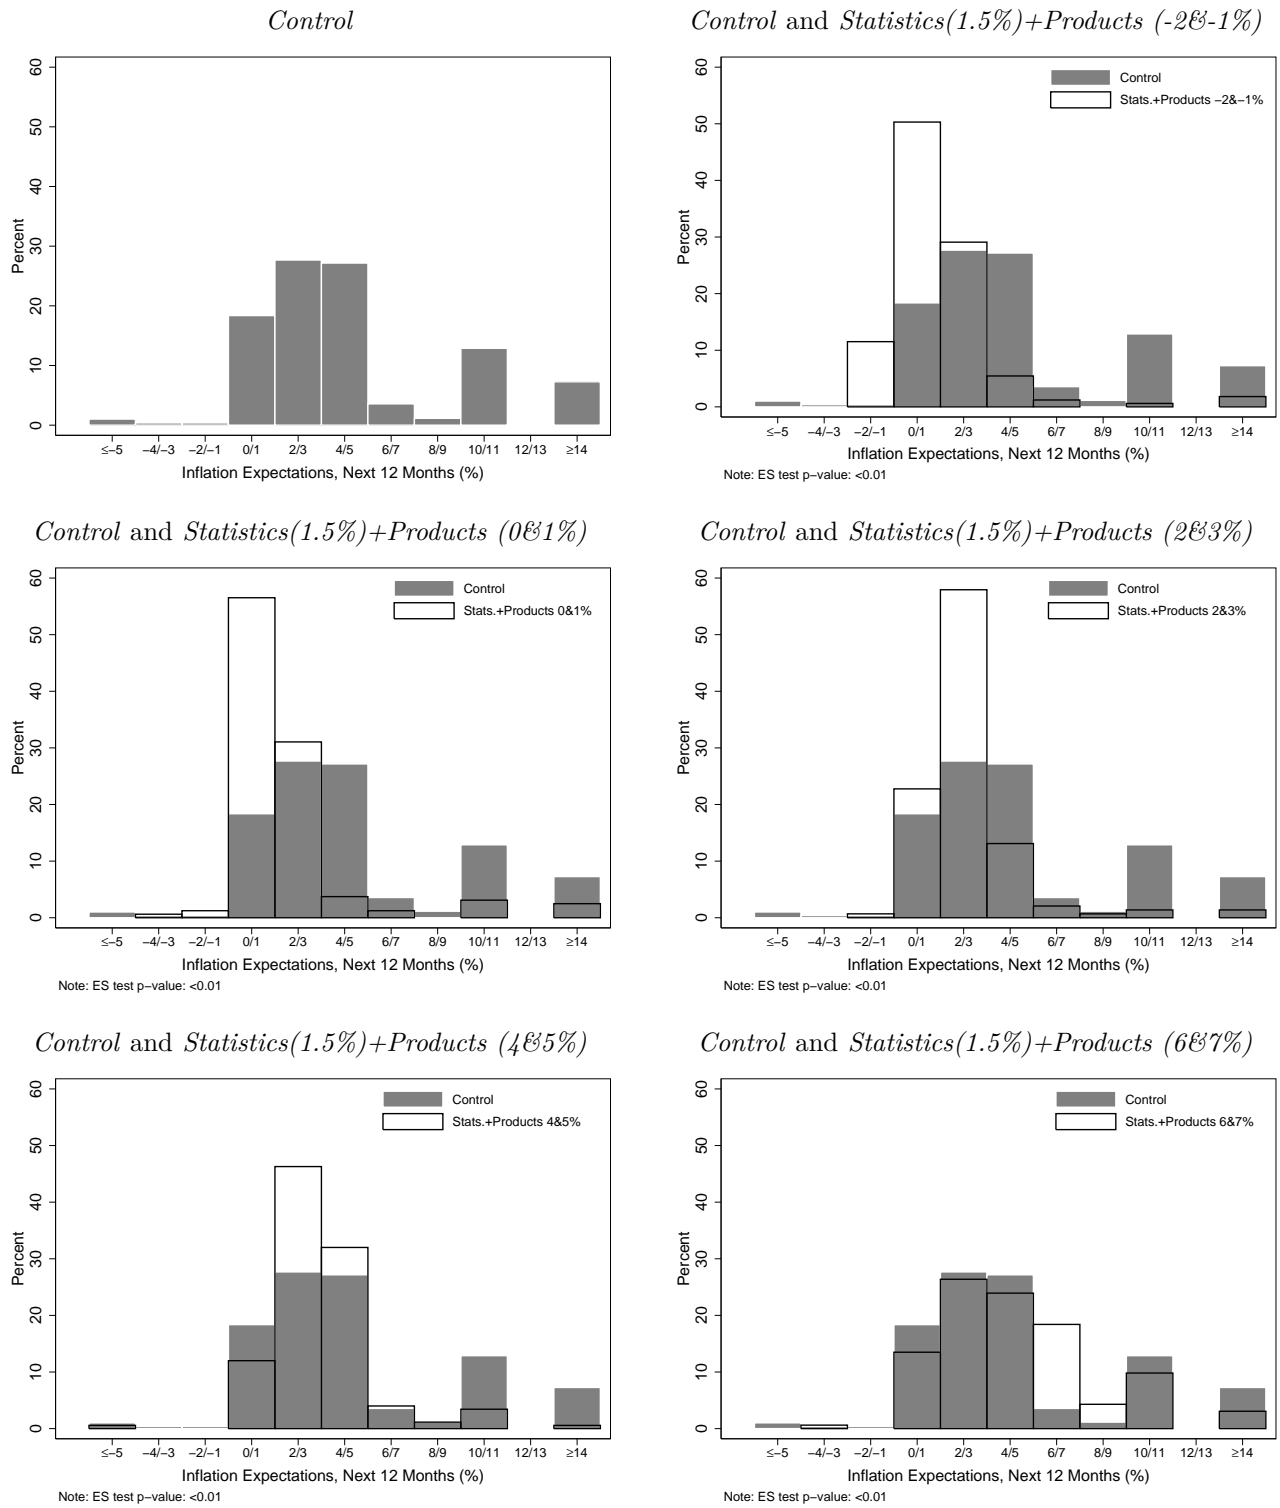
Notes: Prices obtained from online scrapped supermarket prices, from one of Argentina’s largest supermarket chains. The examples in Panels (a) and (b) were used in the Argentina Online Experiment sample II (opinion poll), whereas the treatment in Panel (c) was included in the sample I (college graduates) experiment. The *Products* treatments were preceded by the following text: “Before answering, please look at the table below. For each listed product, the table shows the price on May 1st, 2012 and on May 1st, 2013 (that is, one year later). These prices were taken from the same branch of the main supermarket chain in Argentina.” with the note to the table: “The four products that appear in this table were randomly selected from a database containing hundreds of products. They all belong to well-known brands in Argentina.” The *Statistics (24%)* treatment was preceded by the following text: “According to an average of unofficial indicators produced by private consultancy firms, analysts and research centers, the annual inflation rate in the last 12 months was approximately 24%.”

Figure B.2: Inflation Expectations by Level of *Products* Treatment, U.S. Online Experiment



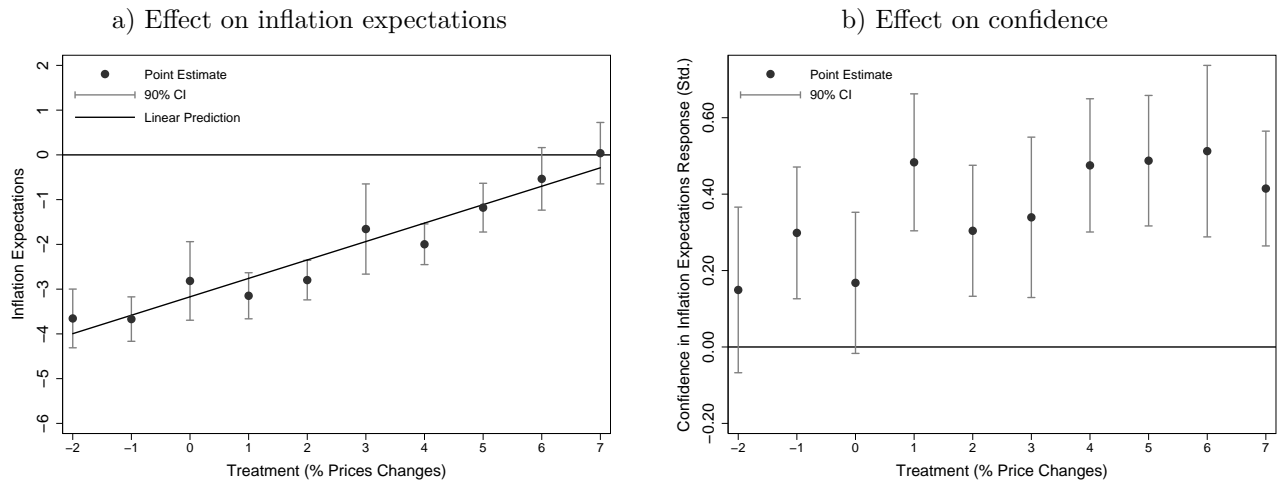
Notes: The observations are from the U.S. Online Experiment, with 783 from the *Control* group and 763 from the *Products+Statistics (1.5%)* treatment (10 tables with average price changes from -2% to 7% in 1 percentage point increments within this treatment). ES is the Epps–Singleton characteristic function test of equality of two distributions.

Figure B.3: Inflation Expectations by Levels of *Products* and *Statistics* (1.5%)+*Products* Treatments, U.S. Online Experiment



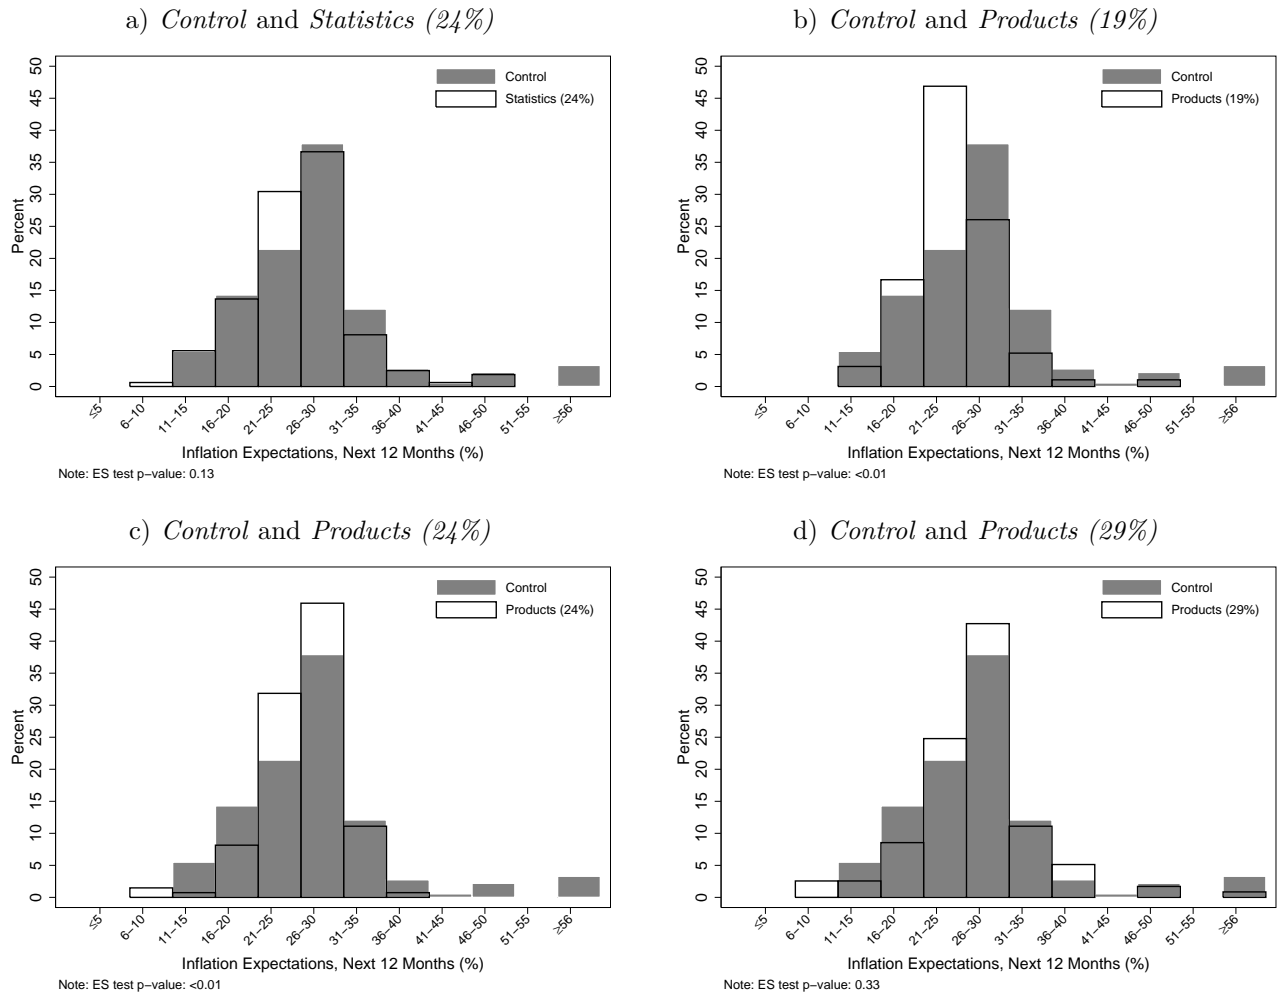
Notes: The observations are from the U.S. Online Experiment, with 783 from the *Control* group and 804 from the *Products+Statistics* (1.5%) combined treatment (10 tables with average price changes from -2 to 7% in 1 percentage point increments within this treatment). ES is the Epps-Singleton characteristic function test of equality of two distributions.

Figure B.4: Treatment Effects on Inflation Expectations and Confidence about Own Expectations by Levels of *Products* and *Statistics (1.5%)+Products* Treatments, U.S. Online Experiment



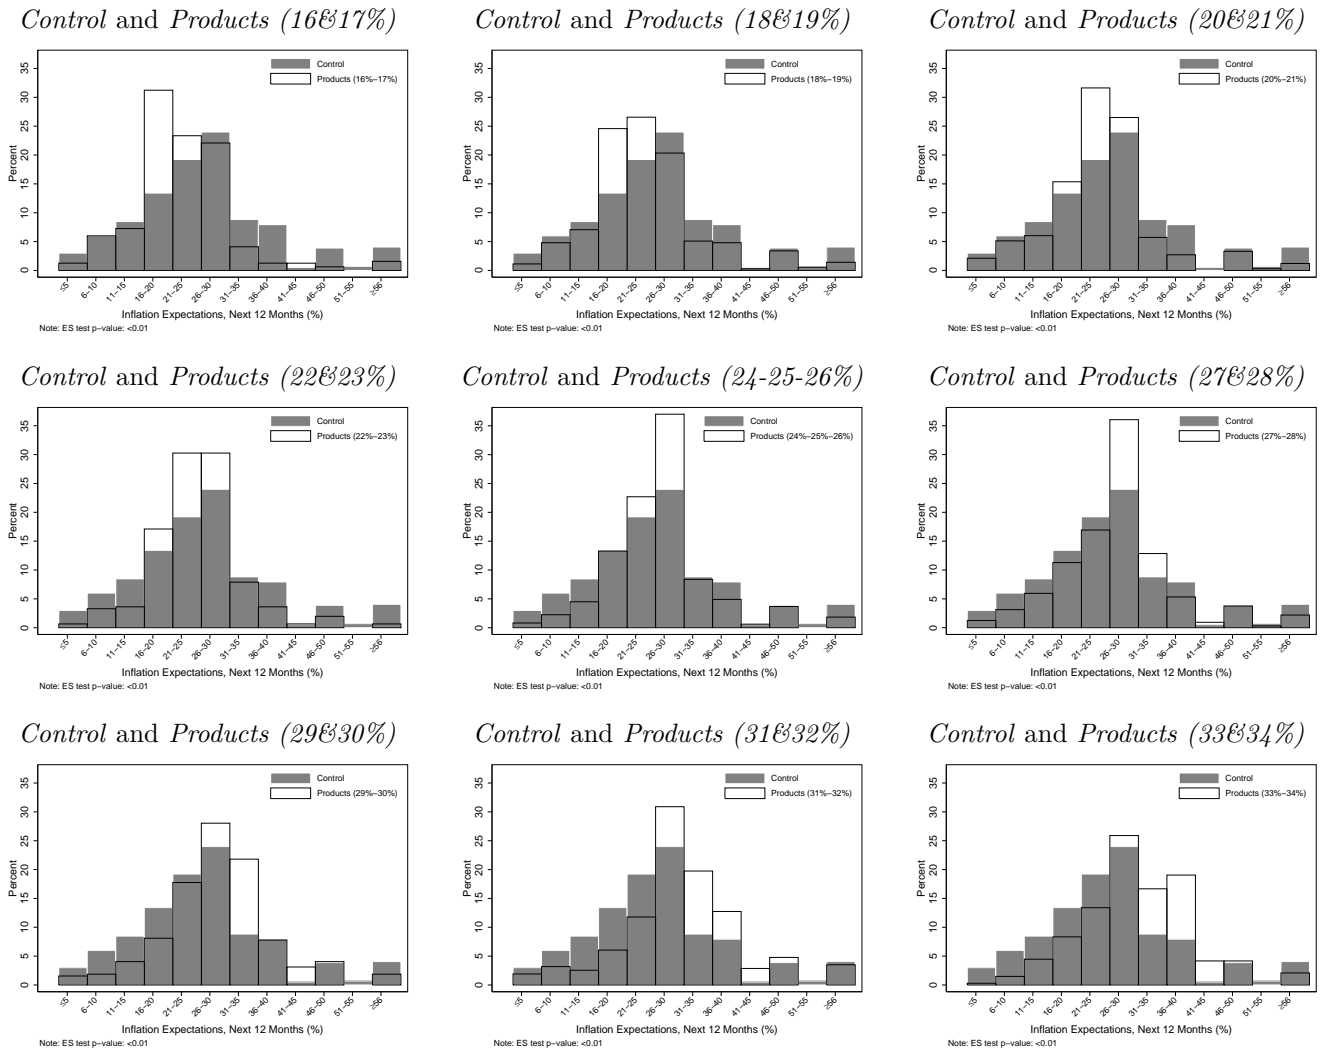
Notes: The total number of observations is 1,732 (789 from the control group and 804 from the 10 variations of the combined specific prices and statistics treatment). Each bar represents the point estimate of the effect of the specific price treatment compared to the control group. Robust standard errors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

Figure B.5: Inflation Expectations by Informational Treatments, Argentina Online Experiment Sample I (College Graduates)



Notes: The Figure presents results for the Argentina college graduates online experiment sample (sample I). The observations correspond to 182 from the *Control* group, 161 from the *Statistics (24%)* treatment arm, and 96, 135 and 117 for the *Products 19%*, *Products 14%* and *Products 29%* treatment arms, respectively. 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 B.6: Inflation Expectations, *Control Products* Groups, Argentina Online Experiment



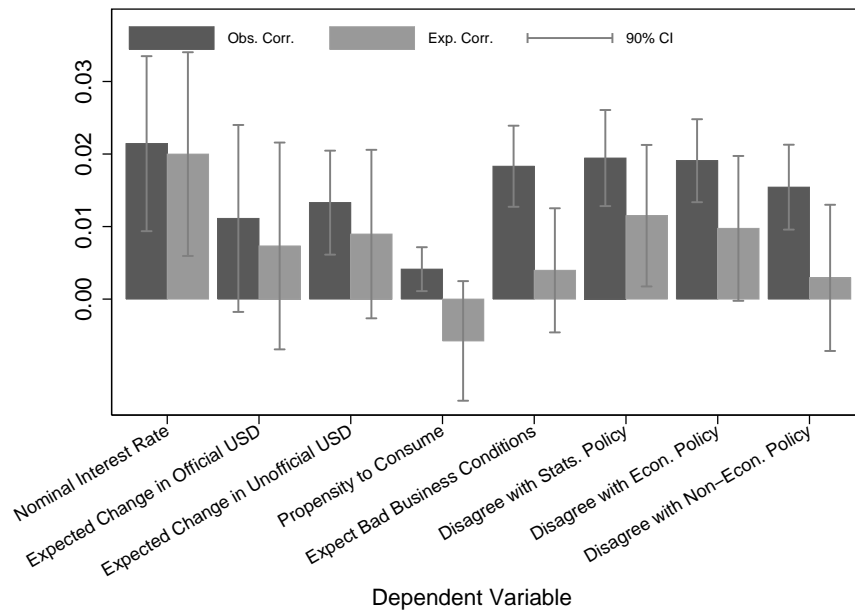
Notes: The source is the Argentina online experiment sample II (opinion poll). The total number of observations is 3,653, with 567 from the control group and 141-177 from each of the 19 treatment groups. ES is the Epps–Singleton characteristic function test of equality of two distributions.

Figure B.7: Observational and Experimental Correlations between Inflation Expectations and Other Economic Variables, U.S. and Argentina (Sample II) Online Experiments

a. United States



b. Argentina



Notes: The total number of observations for Panel (a) is 3,157 (control group and all treatments except *Hypothetical (10%)*). For Panel (b), the total number of observations is 3,653 (Argentina Online Experiment Sample II). 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. For the U.S. experiment, 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.

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%	3,653
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).

Table B.2: Estimates of Learning Rates, Online Experiments, Weighted Estimates

a. United States					
	(1)	(2)	(3)	(4)	(5)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}^{follow-up}$	$i_{i,t+1}$
β	0.776*** (0.040)	0.830*** (0.066)	0.843*** (0.078)	0.519*** (0.066)	0.377*** (0.054)
α -Products	0.593*** (0.046)	0.434*** (0.068)	0.661*** (0.042)	0.487*** (0.180)	0.556*** (0.127)
α -Statistics	0.767*** (0.051)	0.278*** (0.078)	0.696*** (0.094)	0.330** (0.158)	0.318* (0.187)
α -Hypothetical	0.184*** (0.032)		0.214*** (0.055)	-0.071 (0.108)	0.138 (0.103)
Observations	3,141	1,587	1,073	1,073	3,141
Simultaneous treatments	No	Yes	No	No	No
b. Argentina					
	(1)	(2)	(3)	(4)	(5)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}^{follow-up}$	$i_{i,t+1}$
β	1.120*** (0.146)	0.948*** (0.049)	1.018*** (0.050)	0.596*** (0.090)	0.209*** (0.069)
α -Products	0.438*** (0.073)	0.548*** (0.050)	0.461*** (0.091)	0.092 (0.186)	0.665*** (0.183)
α -Statistics	0.449*** (0.111)				
Observations	691	3,653	1,320	1,320	3,373
Sample (experts, online)	I	II	II	II	II

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

Table B.3: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), Robustness Checks, *Statistics (1.5%)* and *Products* Treatments, U.S. Online Experiment

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

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

Table B.4: Learning Model: Weight Given to the Information Provided in the Experiment Relative to Prior Beliefs (α), Argentina Online Experiment Sample II

	(1)	(2)	(3)	(4)
	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\log(e_{i,t+1})$
β	0.902*** (0.042)	0.909*** (0.043)	0.902*** (0.042)	0.328*** (0.088)
<i>Products</i>				
α	0.494*** (0.027)	0.472*** (0.025)		0.435** (0.173)
α^2		-0.001 (0.001)		
α_+			0.484*** (0.040)	
α_-			0.497*** (0.037)	
Observations	3,653	3,653	3,653	1,660
Sample (experts, online)	II	II	II	II

Notes: * significant at the 10% level, ** at the 5% level, *** at the 1% level. Heteroskedasticity-robust standard errors in parenthesis. The α and β coefficients are obtained from the regression given by equation 6, section 2.3, and described in the notes to Table 1. The dependent variable in columns (1) to (3) is inflation expectations (for the following 12 months) at the time of the original survey (March 2013 for sample II). The dependent variable in column (4) is the log of the expected nominal exchange rate of the Argentine Peso with respect to the U.S. Dollar (for the following 12 months) in the original sample II survey. The total number of observations for columns (1)-(3) is 3,653, with 567 from the control group and 141-177 in each of the 19 *Products* treatment groups for the WP Public Opinion Survey. The 1,660 observations in column (4) represent the half of respondents of the WP Public Opinion Survey who were randomly assigned to be asked about the nominal exchange rate and provided a valid answer to this question. The α and β coefficients are obtained from the regression given by equation 6, section 2.3. α^2 represents the squared learning weight parameter. α_+ and α_- represent the learning weight parameters differentiated for those with positive and negative differences (respectively) between the reported value of the difference between the informational signal provided and the own reported value of past inflation perception, $(\pi_{i,t}^T - \pi_{i,t}^0)$.

C Argentina Supermarket Experiment: Further Details and Results

C.1 Further Details about the Supermarket Experiment

This Section complements the discussion of the supermarket experiment in the body of the paper by presenting additional details about the implementation of the survey. The study 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 interviewers carried a handheld scanner, with which they scanned the respondents’ receipt from the supermarket purchase. The interviewers reported high levels of interest and curiosity from the respondents, especially about the use of the handheld scanners.

The following is an extract from the enumerators instruction manuals, translated from Spanish. Verbal statement to engage interviewees: “Hi, we are from the Universidad Nacional de La Plata. Are you willing to participate in a study on economic expectations? It will only take 5 minutes”. *To those who accept, please explain the following:* “This study attempts to relate individual shopping patterns with their economic perceptions. For this purpose, we need you to let us scan your shopping receipt. This information, the list of products, will allow us to develop the empirical analysis for our study. The receipt does neither contain your name nor any sensitive information. The survey is completely anonymous. Once 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). 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 receipts for scanning, the respondents were asked 12 questions to gather evidence on inflation perceptions and memories of price changes, among other outcomes of interest. As in the research design of our online experiments, we capture the subjects’ prior belief about inflation by asking them about their perceptions of the rate of inflation over the past year. This question was followed by some randomized treatments, and then by a final question about inflation expectations. Appendix G presents the original survey instrument, the three specific product tables, and the enumerators instruction manual.

C.2 Further Results

The results from the supermarket experiment presented in the body of the paper were based on actual and remembered price changes for products the respondents had just purchased. The

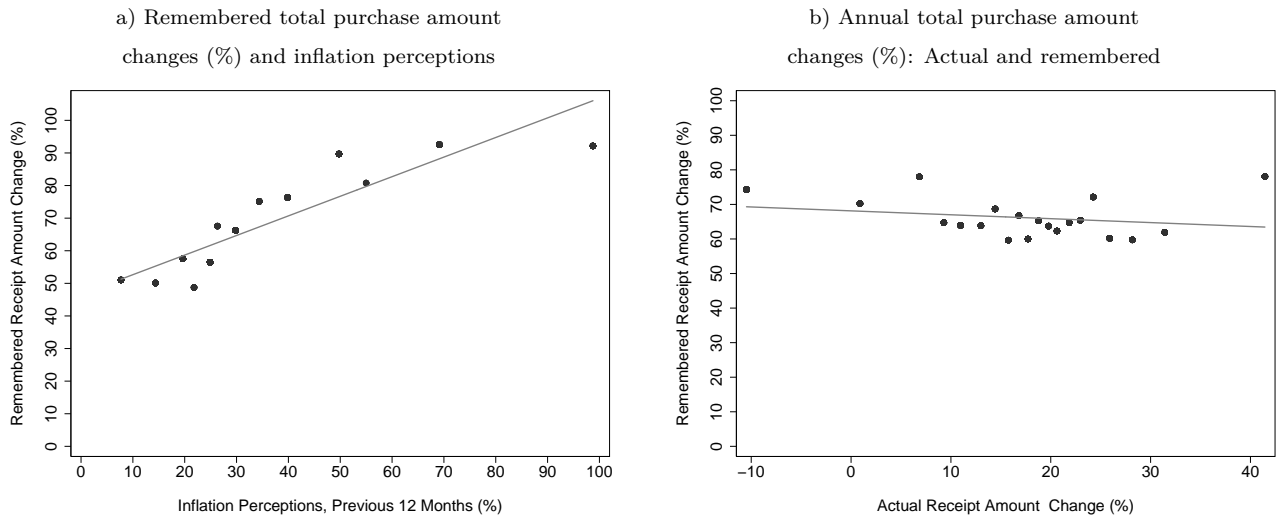
results indicate that individuals seem to have a poor memory about price changes for individual products. However, individuals may have a better recollection of the price of bundles of products, for instance, the price of the basket of products they had just purchased. To test this hypothesis, in our supermarket experiment, 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 her estimate of the total amount she would have had to pay for the same goods 12 months earlier.

As a further robustness check of the results in the body of the paper, we compare the individual's estimate of the change in her purchase's total amount and the actual change in the total cost according to our price database. Figure C.1 is based on this comparison. It depicts the relationship between the estimate of the change in the receipt's total amount and inflation expectations (Panel a), as well as the relationship between this estimate and the actual change (Panel b). The results are very similar to those we obtain with the changes in specific product prices: there is a positive relationship between the subjects' estimates and their inflation perceptions, but virtually no correlation between the actual and the estimate of the receipt's total amount change. The similarity of these results indicates that respondents do not seem to fare any better when asked about total purchase amounts instead of specific products.

The supermarket experiment also included an informational treatment with tables of products with three levels of average price changes. Panels (a) and (b) in Figure C.2 present the distributions of inflation expectations in pairwise comparisons between the *Products* treatments. While there is no statistically significant difference between the distributions of the 19% and the 24% treatments (the ES test does not reject the null of equality of distributions – p-value of 0.24), the *Products (19%)* and *Products (29%)* treatments are statistically different: average inflation expectations are clearly higher when the subjects were shown tables with the highest average price changes. This evidence merely confirms the findings from the online experiments that individuals incorporate objective information about prices of specific products.

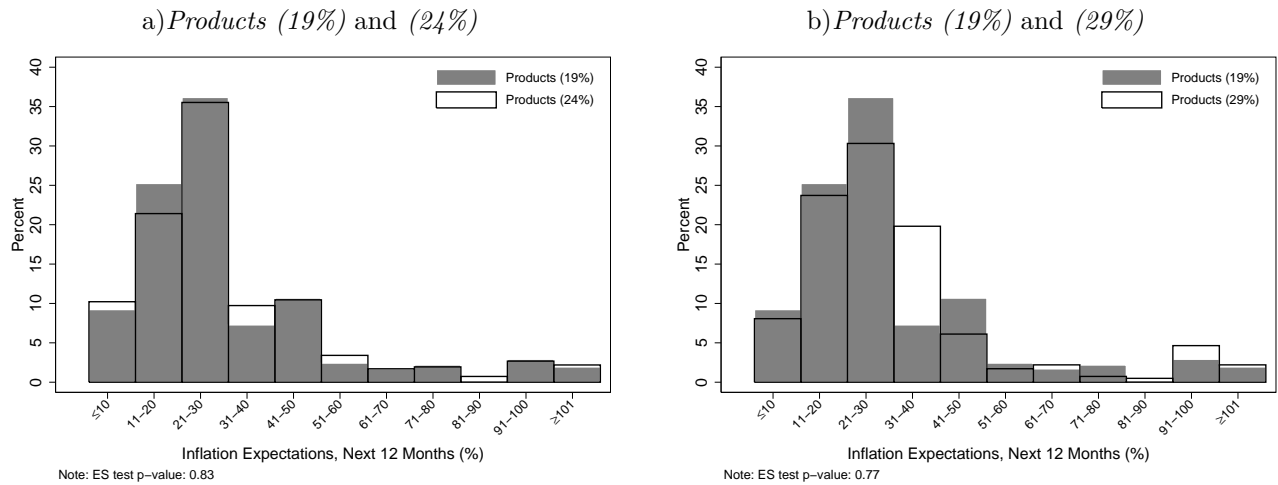
While the rate of learning from remembered price changes of specific products could also be depicted by means of the Bayesian learning model used before, we must note that, in contrast to the other informational treatments, we did not randomize the remembered price changes directly, but instead we randomized the salience for a group of products. As a result, we cannot compare the α from randomizing salience with that from randomizing the information directly.

Figure C.1: Robustness: Implicit Price Changes from Total Purchase Amount and Inflation Expectations, Supermarket Experiment, Argentina



Notes: The total number of observations is 1,140. Panels (a) and (b) represent binned scatterplots, where the number of observations are almost identical across bins. The percentage changes in both panels are implicit – they are obtained from the total purchase amounts in Argentine 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.

Figure C.2: Inflation Expectations by *Product* Treatment Levels, Argentine Supermarket Experiment



Notes: Source: Argentina Supermarket Experiment. The total number of observations is 1,232 for panels (a) and (b) (412 from the *Products (19%)* group, 411 from the *Products (24%)* group and 409 from the *Products (29%)* group). ES is the Epps–Singleton characteristic function test of equality of two distributions.