Federal Reserve Bank of New York
Staff Reports

The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment

Olivier Armantier
Scott Nelson
Giorgio Topa
Wilbert van der Klaauw
Basit Zafar

Staff Report No. 543
January 2012
Revised January 2013

The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment
Olivier Armantier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar
Federal Reserve Bank of New York Staff Reports, no. 543
January 2012; revised January 2013
JEL classification: C81, D80, E31

Abstract

Understanding the formation of consumer inflation expectations is considered crucial for managing monetary policy. Using a unique "information" experiment embedded in a survey, this paper investigates how consumers' inflation expectations respond to new information. We elicit respondents' expectations for future inflation before and after providing a random subset of respondents with factual information that may affect their expectations. This design creates unique panel data that allow us to identify causal effects of new information. We find, first, that baseline inflation expectations are right-skewed, and that consumers in the high-expectation right tail are relatively underinformed about objective, inflation-relevant facts. We next find that providing consumers with new information causes them to update their expectations, such that the expectations distribution converges toward its center. Furthermore, respondents who update do so in sensible ways: revisions are proportional to the strength of the information signal, and inversely proportional to the precision of baseline inflation expectations. Our findings indicate that heterogeneous consumer expectations are a result of both different information sets, as well as different information-processing rules. Overall, our results are consistent with a Bayesian learning model. We discuss implications of these results for monetary policy and for macroeconomic modeling.

Key words: inflation expectations, information, heterogeneous expectations, updating, Bayesian learning

---

Armantier, Topa, van der Klaauw, Zafar: Federal Reserve Bank of New York. Address correspondence to Basit Zafar (e-mail: basit.zafar@ny.frb.org). Nelson: Massachusetts Institute of Technology (e-mail: nelsons@mit.edu). The authors thank Carlos Madeira, Ricardo Reis, and participants at presentations at the Federal Reserve Bank of New York seminar series, the State Bank of Pakistan, and the CIPREE Conference on Formation and Revision of Subjective Expectations for comments. All errors that remain are the authors'. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1. Introduction

“A fuller understanding of the public's learning rules would improve the central bank's capacity to assess its own credibility, to evaluate the implications of its policy decisions and communications strategy, and perhaps to forecast inflation.”

(Ben Bernanke, 2007)

Many economic decisions – consumption, saving, wage bargaining, investing – are believed to be influenced by expectations about inflation. Inflation expectations have now become central to macro-economic models and monetary policy (Gali, 2008; Sims, 2009), and managing consumers' inflation expectations has become one of the main goals of policy makers.[1] Indeed, national surveys of public inflation expectations are now conducted in multiple countries.[2] However, managing inflation expectations requires not just monitoring expectations, but also understanding how these expectations are formed.

Studies based on survey data have shown substantial divergence among individuals' beliefs about future inflation (Mankiw, Reis, and Wolfers, 2003), which the recent literature attempts to explain as a result of different expectation-formation processes: e.g., some form of bounded rationality (Sargent, 1993; Evans and Honkapohja, 2001; Mankiw and Reis, 2002), adaptive learning (Orphanides and Williams, 2006), switching between different prediction rules (Branch, 2007), time-dependent rules under which expectations are updated only at fixed intervals (Carroll, 2003), or learning from experience (Malmendier and Nagel, 2010; Madeira and Zafar, 2011). However, while this literature has found some aggregate data patterns consistent with these models, there nevertheless remains little direct empirical evidence on how individual consumers form their inflation expectations. This paper helps fill that gap.

We conduct an experiment in which we randomly provide a subset of survey respondents with information (which we refer to as “treatment information”) about either past-year average food price inflation, or professional economists’ median forecast of next-year overall inflation. Before this subset of respondents receives this information, and again after this subset receives this

[1] Bernanke, 2004, argues that "an essential prerequisite for controlling inflation is controlling inflation expectations".

[2] These include the Reuters/University of Michigan Survey of Consumers, the Livingston Survey, the Conference Board’s Consumer Confidence Survey and the Survey of Professional Forecasters in the US. Other central banks that survey consumers about their inflation expectations include the Bank of England, the European Central Bank, the Bank of Japan, the Reserve Bank of India, and the Sveriges Riksbank.

information, we ask *all* respondents for their expectations of future inflation. This experimental design thus creates a unique panel dataset which allows us to observe how this new information induces respondents to update their inflation expectations.

We also ask all respondents for their priors about the randomly-provided information (henceforth, referred to as "information priors"), and test whether respondents' ex-ante informedness can help explain (1) baseline heterogeneity of expectations for future inflation, and (2) updating of these expectations after the information is revealed. We expect respondents who are less informed about either of the information treatments ex ante (that is, those who exhibit larger gaps on average between their information priors and the treatment information) to have more extreme expectations for future inflation. For these respondents, the information treatment may contain more valuable information that causes expectations updating and may thus result in larger expectation revisions. The patterns in who updates – and how much they do so – shed light on how expectations are formed and how consumers react to (possibly) inflation-relevant, new information.

Furthermore, following the recent literature on the importance of inflation survey question-wording (Bruine de Bruin et al., 2010a), when we elicit inflation expectations we randomly ask half of our respondents for their expectations of overall inflation (the "rate of inflation") while asking the other half of respondents for their expectations of their own-basket inflation (the "prices of things you usually spend money on", hereafter referred to as "prices you pay").[3] We test here whether consumers find different types of information (from our two information treatments) more or less appurtenant to these two "types" of inflation expectations. This is relevant for a better understanding of the heterogeneity and content of consumer inflation expectations, and for determining the reliability of different survey question wordings for elicitation of consumer inflation expectations.

Compared to existing studies, the approach used in this paper differs in that we (1) can remain agnostic about each respondent's information set, (2) explain the heterogeneity of expectations without imposing any particular learning rule or information-processing rule for consumers, and (3) directly infer the *causal* effects of different types of inflation-relevant information on individual consumers. Previous studies have mostly overlooked the panel dimension of survey expectations (see Keane and Runkle, 1990, for an exception), and instead have studied the

[3] Thus, we have four information treatment cells (given two information treatments and two expectation-type questions), and four corresponding control treatment cells.

aggregate evolution of beliefs in repeated cross-sections; this complicates the interpretation of previous work on learning in expectation updating.

Our first result concerns respondents' *perception gaps* – the gaps between the treatment information and respondents' information priors. While a substantial number of our respondents have information priors that are closely in line with the treatment information, the distribution of information priors is highly skewed, and 37.5% of respondents have perception gaps of three percentage points or more. Thus, average perceptions gaps are substantial: -56% (i.e., over-estimation by 56%) for the median forecast of professional economists, and -139% for past changes in food prices. This result closely matches that of Bryan & Venkatu (2001a), who find that respondents' "perceptions" of past-year price changes are on average too high. [4]

Exploiting our information treatment, next we find that new information indeed causes respondents to update their inflation expectations, and to do so in a sensible manner. On average, we find that respondents: (1) revise their inflation expectations down if their perception gaps were over-estimates (and vice versa for under-estimates), (2) revise their expectations more when their perception gaps are larger, and (3) are more receptive to the information when the uncertainty in their baseline inflation expectations is greater (as would be the case under Bayesian updating). Thus the information treatment causes the distribution of inflation expectations to converge towards its center. Looking beyond average effects, there is also substantial heterogeneity in how information is processed by the various demographic groups. In one striking finding, we find that most substantial updating behavior in our sample is driven by female respondents, even after controlling for female respondents' higher average perception gaps. This is partly a consequence of females' less precise baseline inflation expectations relative to male respondents, and is likely a result of differences in females' expectation formation rules.

Both of our information treatments — past-year food price inflation and a professional forecast of next-year overall inflation — affect respondents' expectations. This provides empirical support for models in which consumers derive their forecasts from news reports of the forecasts of professional economists (Carroll, 2003), and models in which consumers base their inflation

[4] This finding is generally consistent with a literature that shows individuals can be uninformed when making decisions of economic significance: low-income families are unaware of basic features of the Earned Income Tax Credit (Chetty and Saez, 2009); students have incorrect perceptions of returns to schooling (Jensen, 2010; Wiswall and Zafar, 2011); most households are unaware of their marginal price for electricity and water (Brown et al., 1975; Carter and Milon, 2005).

expectations on news releases about previous-period inflation (Garner, 1982; Hey, 1994; Lanne, Luoma, and Luoto, 2009).

The effect of our two information treatments is, however, heterogeneous. Our analysis shows that information about food prices causes consumers to update expectations more for their own-basket inflation rate, and less for "rate of inflation," whereas we find that information about forecasts of *overall* inflation causes consumers to update expectations for the "rate of inflation" primarily. The greater response of own-basket inflation expectations to food price information (relative to the rate of inflation) may be a consequence of either (1) food price changes having more relevance for consumers' own-basket inflation rate than for the overall Consumer Price Index,[5] or (2) consumers having biased perceptions about the share of food expenditures in their budget, or about the co-movement between food prices and prices of other items in their basket, or (3) consumers focusing on specific salient price changes when considering the change in "prices" rather than "inflation" (Bruine de Bruin et al., 2012). Looking across demographic groups, we see further differences between our two information treatments' effects. For example, low-income respondents revise significantly their "prices you pay" forecast when provided with information about past food prices, while high-income respondents do not, even though both groups receive equally informative signals.[6] This result is consistent with food purchases occupying a larger share of lower income respondents' consumption baskets. Overall, the heterogeneous response for the two question texts and information treatments underscores how certain types of information may be more or less relevant for different measures of inflation, and has implications for the modeling of inflation expectations as well as survey research that collects such data.

We also find important heterogeneity in respondents' informedness about our treatment information in general, and in respondents' subsequent responsiveness to this information. Certain demographic groups – female, older, and less-educated respondents, as well as those with less financial literacy – generally have higher perception gaps about objective inflation information ex ante. It is well-documented that these same demographic groups tend to have higher expectations

[5] The CPI is a plutocratic index, i.e., it weights individual consumers' consumption baskets by their expenditure share (relative to total expenditure). Thus, people who consume more have larger weights in the CPI. Therefore, the simple average of food shares in individual's consumption baskets may not be the same as the food share in the CPI. In particular, if food shares are lower for higher income individuals (McGranahan and Paulson, 2006), the food share in the CPI would be lower than the simple average of food shares in individuals' consumption baskets).

[6] That is, perception gaps among high-income individuals – as revealed to respondents through the information treatment – are not significantly different from those of low-income individuals.

for *future* inflation than their counterparts do,[7] contributing to the strong right-skew of the inflation expectations distribution.[8] Thus we offer an alternative novel explanation for the systematically high inflation expectations of these demographic groups, by identifying a relative gap in their own information sets about objective inflation measures. Furthermore, we find that some of these demographic groups – and, as mentioned above, females in particular – are more responsive to new information. This is an encouraging result: it suggests that policy-makers could partially control the high-expectation right-tail of the inflation expectations distribution through public information campaigns in the spirit of our information treatments.

On the other hand, not all respondents with high perception gaps follow this pattern, and a number of respondents do not revise their expectations at all. While our results shed some light on these non-revisers – they tend to have smaller perception gaps (and hence the treatments have less informational content for them), and are more likely to be male – we are not able to completely predict non-revisions based on demographics or perception gaps. Thus we must conclude that some respondents either do not find the provided information relevant for inflation expectations, or do not find the information credible. This suggests any public information campaigns to help anchor consumer inflation expectations need to be carefully designed and multi-pronged.

Our finding of heterogeneous perception gaps is indicative of differences in consumers' information sets, while the heterogeneous response to information (conditional on the perception gaps) suggests that consumers use heterogeneous updating rules. Therefore, our results suggest that it could be fruitful to model consumers' heterogeneous expectations as a result of both different information sets as well as different information-processing rules. A further insight for future modeling is that, since our information experiment provides information that is readily available, such information should have no systematic effect on individuals' forecasts in a purely rational expectations framework (as defined in Muth, 1961). Our results, therefore, are consistent with the average consumer being boundedly rational. Moreover, adaptive learning would rule out consumers responding to certain kinds of information, such as experts' forecasts of future inflation. In the end, we conclude that our results are most consistent with a Bayesian updating rule for consumer inflation expectations.

---

[7] See Jonung (1981), Bryan and Venkatu (2001a, 2001b), Lombardelli and Saleheen (2003), Blanchflower and Coille (2009), and Bruine de Bruin et al. (2010a). We also replicate these patterns in expectations data here, though results are not always significant at standard levels.

[8] Indeed, whereas *median* consumer inflation expectation survey responses generally track official estimates of realized inflation and sometimes even outperform professional forecasters (Thomas, 1999; Ang, Bekaert & Wei, 2007), *average* consumer inflation expectations are systematically higher than realized inflation (Bryan and Venkatu, 2001a, 2001b)

This paper is organized as follows. The survey design and data collection methodology are described in Section 2. Section 3 conducts the empirical analysis. Section 4 presents the main results, and outlines the implications for modeling of consumer expectations. We conclude in Section 5 with a discussion of the policy implications of our study and its limitations.

2. Data

Our data are from an original survey that is part of an ongoing effort by the Federal Reserve Bank of New York, with support from academic economists and psychologists at Carnegie Mellon University.[9]

The survey was conducted over the internet with RAND's American Life Panel (ALP). Our target population consists of individuals 18 or older who participated in the Reuters/University of Michigan Survey of Consumers Survey between November 2006 and July 2010 and subsequently agreed to participate in the ALP.[10] Out of a total sample of 771 individuals invited to participate in the survey, 735 did so, implying a response rate of 95.3%. The survey was fielded between January 3rd, 2011 and February 9, 2011. Respondents received $20 for each completed survey.

2.1 Survey Design

The survey consisted of two sets of questions. The first set of questions, analyzed in Armantier et al. (2011), examines the link between self-reported beliefs and economic behavior. The second set of questions—the focus of this paper—investigates how individuals revise their inflation expectations after being exposed to new information.

For the sake of concreteness, we introduce notation here that we will use throughout the paper to refer to our survey-measured quantities of interest. As diagrammed in Figure 1, these quantities were measured in four different survey stages:

1. Baseline Inflation Expectations: In the first stage, respondents were randomly assigned to one of two questions that elicited their baseline expectations for future inflation, either for

[9] The general goal of this initiative is to develop better tools to measure consumers' inflation expectations, to study the link between expectations and behavior (Armantier et al., 2011), and to better understand how the public forms and updates expectations about future inflation (Bruine de Bruin et al., 2010b).

[10] The Michigan survey is a monthly telephone survey with 500 respondents, consisting of a representative list assisted random-digit-dial sample of 300, and 200 respondents who were re-interviewed from the random-digit-dial sample surveyed six months earlier. Our target population is further restricted to active ALP members, defined as those who either participated in at least one ALP survey within the preceding year, or were recruited into the ALP within the past year.

their own consumption basket (the "prices you pay") or for the economy overall (the "rate of inflation"). We refer to these measured quantities as "baseline inflation expectations," and represent these quantities for each consumer i as either $\pi_{i,PP}$ ('PP' is for "Prices you Pay") or $\pi_{i,RI}$ ('RI' is for "Rate of Inflation").

2. Information Priors: In the second stage, respondents were randomly assigned to one of two questions that measured their ex-ante informedness about (possibly) inflation-relevant information. The two questions, detailed below, asked respondents for either their belief of the average change in food and beverage prices over the last year, or their belief about professional economists' median forecast for next-year inflation. We refer to respondents' beliefs as their "information priors", which we denote for each consumer i as either $\omega_{i,Food}$ or $\omega_{i,SPF}$ (where 'SPF' is for "Survey of Professional Forecasters").
3. Treatment Information: With a probability of 75%, respondents were provided with the true values (defined as values published in two publicly available data series) of the quantities about which they had been asked: again, either the average change in food and beverage prices over the last year, or professional economists' median forecast for one-year ahead inflation. We refer to these quantities as "treatment information," represented as ω^*_{Food} and ω^*_{SPF} (with no i subscript).
4. Final Inflation Expectations: In the final stage, inflation expectations were re-elicited from all respondents, with each respondent being asked the same inflation question they were asked in the first stage. We refer to these expectations as "final inflation expectations," represented as either $\pi'_{i,PP}$ or $\pi'_{i,RI}$.

2.1.1 Stage 1

In the first stage, respondents were asked to report their baseline inflation expectations using one of two randomly assigned questions. The two question-texts that elicit inflation expectations are:

1) "Prices You Pay" (PP) which asks for "*your expectations for the prices of things you usually spend money on going into the future*";
2) "Rate of Inflation" (RI) which asks for "*your expectations for the rate of inflation/deflation going into the future.*"

Both questions were asked for two different horizons: (1) a point forecast "over the next 12 months", which corresponds to the period January 2011 and January 2012; and (2) a point forecast "over the one-year period between January 2013 and January 2014," which, at the time of the survey, was three-year ahead one-year inflation. Since most of the analysis in the paper relates to the next-twelve-months point forecast, this is the quantity we represent by either $\pi_{i,PP}$ or $\pi_{i,RI}$.

In addition, we also asked respondents for their density forecast over the next 12 months: here respondents assigned probabilities to possible future inflation outcomes such as "the rate of deflation will be between 0% and 2%" or "the rate of inflation will be 12% or more". These choices were mutually exclusive and collectively exhaustive, and respondents could verify that their answers summed to 100% probability. Following the approach developed by Engelberg, Manski and Williams (2009), a generalized beta distribution is fitted to each respondent's stated probabilistic beliefs (see also Bruine de Bruin et al., 2011a). We then generate the mean and variance of the respondent's beta distribution, which is used in the empirical analysis.

The PP question text is similar to the "prices in general" question text studied by Bruine de Bruin et al. (2011a). That "prices in general" question text is the version used in the University of Michigan's Survey of Consumers, which produces the often-cited monthly measure of consumer inflation expectations. While the Michigan Survey's question asks respondents "*During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?*", we ask respondents about "*the prices of things you usually spend money on*". This change to the question wording was prompted by research showing that the Michigan Survey's question-text induces mixed interpretations, with some respondents thinking about specific prices they pay and others thinking about the overall rate of inflation (Bruine de Bruin et al., 2011b). The PP question is designed to be less likely to have mixed interpretations: the PP question is meant to cue respondents to think about prices of specific purchases in their consumption basket, while the RI question is meant to focus respondents on general price levels, or the overall cost of living (Bruine de Bruin et al., 2012.

2.1.2 Stage 2

The second stage consisted of eliciting information priors:

1. The "Food" treatment: asked *"Over the last twelve months, by how much do you think the average prices of food and beverages in the US have changed?"*

2. The "Survey of Professional Forecasters (SPF) Forecast" Treatment: asked "*A group of professional economists report their expectations of future inflation on a regular basis. What do you think these professional economists predicted inflation to be over the next twelve months?*"

In both cases, respondents are asked for a point estimate of year-over-year percentage change. As described above, these are the quantities $\omega_{i,Food}$ or $\omega_{i,SPF}$ for consumer i. Between stages 1 and 2, respondents participated in a battery of experimental questions related to inflation and investment (discussed in Armantier et al., 2011), and also answered several questions about consumption behavior and sources of information about inflation/prices.

2.1.3 Stage 3

In the third stage, immediately after reporting their information priors, 75% of respondents were randomly provided with true measures—defined as those published in publicly available data series—for which their information prior was elicited in Stage 2. These true measures are the treatment information. For the Food treatment, we used the series of average food and beverage prices for urban US consumers that are produced by the Bureau of Labor Statistics. Respondents saw the following information:

"*According to the most recent data available from the Bureau of Labor Statistics, the average prices of food and beverages in the US* ***INCREASED by 1.39%*** *over the last twelve months.*"

Thus, $\omega^{*}_{Food} \equiv 1.39$. For the SPF Forecast treatment, we used the median forecast of next-year Consumer Price Index (CPI) inflation from the Federal Reserve Bank of Philadelphia's quarterly Survey of Professional Forecasters (SPF). Respondents in this treatment saw the following information:

"The Survey of Professional Forecasters (SPF) is a quarterly survey of professional economists. According to the latest data, these professional economists expect, on average, ***inflation to be 1.96%*** *over the next twelve months.*

Not all of these professional economists agree about future inflation though. ***However, most (90%) of them expect inflation over the next twelve months to be between 1.19% and 3.03%."***

Thus, $\omega_{SPF}^{*} \equiv 1.96$. Our information treatments focus on either a backward-looking measure of inflation in a basket of specific goods (the food and beverages component of the CPI), or a forward-looking measure of general inflation (the overall CPI), which also includes the underlying distribution (precision) of the information. Our Food information treatment is motivated by empirical studies that show subjects' expectations of future prices/inflation are determined by past prices (Garner, 1982; Hey, 1994; Andersen, 2008; Lanne et al., 2009). Our SPF Forecast treatment is motivated by evidence and theory that consumers derive their expectations about future prices/inflation from news reports of the forecasts of experts (Carroll, 2003).

Also in the third stage, 25% of respondents were given no treatment information, and these respondents make up our control group.[11] In the analysis, we refer to the 75% of respondents who receive the objective information as being in the "treatment" group, and refer to the remaining 25% who do not receive the true information as being in the "control" group.[12]

2.1.4 Stage 4

Finally, all respondents were asked again for their *final* inflation expectations, using the same question-text that they received in the first stage. Again, these are the quantities we represent as $\pi'_{i,PP}$ or $\pi'_{i,RI}$ for consumer i. The reason for keeping a control group within each of the four treatment cells is that the simple act of taking a survey about inflation expectations (including receiving our questions in stage 2) may make respondents think more carefully about their responses and may lead them to revise their expectations even if they are not provided with any new information (see Zwane et al., 2011, for a discussion of how surveying people may change their subsequent behavior). Since we are interested in revisions in expectations that are directly attributable to the information, we identify that off of *differences* between the treatment groups' and control groups' changes in expectations.

2.2 Survey Respondents

[11] They are primarily used for identifying the causal effect of the information provided to all other respondents.

[12] Thus after the stage three random assignment, we have four treatment cells – RI × Food (respondents who report RI inflation expectations and receive the Food treatment), PP × Food, RI × SPF, and PP × SPF, each comprising 50% × 50% × 75% = 18.75% of the total sample – and four corresponding control groups, each comprising 50% × 50% × 25% = 6.25% of the total sample.

Among our 735 respondents, 705 finished the survey, of whom 667 gave answers for the minimum set of questions needed for our analysis: that is, answers for information priors, as well as both baseline and final inflation expectations for at least one of our three inflation questions (year-ahead point forecast, year-ahead density forecast, and three-year-ahead point forecast). We additionally exclude from our analysis 11 respondents with unusually high (greater than 50 percentage points) information priors (about Food or about the SPF Forecast) or baseline inflation expectations. Thus, we are left with a total sample of 656 respondents. Table 1 shows resulting sample sizes for each of the four treatment cells and corresponding control cells.

For these 656 respondents, average age is 52.7 years (standard deviation=14.0), with 43.1% being male, 87.7% non-Hispanic white, and 5.6% non-Hispanic black. The median annual family income is reported as “$60,000 to $74,999”, and 83.2% of respondents have an annual family income of $30,000 or more. Respondents hail from 48 different U.S. states, and 52.3% have a 4-year college degree. Hence our sample has higher income and higher educational attainment, and also has more white respondents, than the US population overall.

For the analysis, we define a respondent to be high income if the annual household income is at least $75,000; 42.8% of the sample falls in this group. We define a respondent to be “older” if the respondent is at least 55 years of age; 47.6% of the sample falls in this group.

We paid respondents a fixed compensation for completing the survey, and did not elicit respondents' inflation expectations or information priors using a financially incentivized instrument such as a scoring rule. This is because proper scoring rules may generate biases when respondents are not risk neutral (Winkler and Murphy, 1970). Moreover, incentivized belief elicitation techniques are not incentive-compatible when the respondent has a stake in the event that they are predicting (Karni and Safra, 1995), which is the case for inflation expectations. In addition, Armantier and Treich (2011) show that elicited beliefs are less biased (but noisier) in the absence of incentives.

3. Empirical Analysis

We begin our empirical analysis by reviewing data patterns in baseline inflation expectations, average revisions in expectations after the information treatments, and average absolute revisions; summary statistics for these are presented in Table 1.[13] Our goal here is to characterize the sample's updating behavior in aggregate. Later, we turn to documenting the sample's ex-ante informedness about our information treatments, and investigate respondents' responsiveness to new information, conditional on their ex-ante informedness and demographics.

Table 1 shows that median "prices you pay" (PP) expectations at baseline are substantially higher than "rate of inflation" (RI) expectations at baseline: 5 percentage points for PP versus 3 percentage points for RI. Bruine de Bruin et al. (2012) also find that expectations tend to be higher for PP than for RI. This highlights the importance of treating "own-basket" and "overall" inflation expectations separately in our analysis.

Median revisions are zero in all treatment and control groups, whereas the median *absolute* revision is nonzero in some cases – mostly, for respondents that were asked about the SPF Forecast – reflecting a combination of both upward and downward revisions in our sample. Meanwhile, in all four treatment groups and in all but one control group, mean revisions are negative, indicating average downward revisions. These downward revisions are larger in the treatment groups than in the control groups. For example, the mean downward revision in the RI × SPF treatment group is 1.96 percentage points, compared to a mean downward revision of 0.93 percentage points in the control group. If we had very large sample sizes, and if the distribution of perception gaps were asymmetric around zero, then we should expect these differences to be statistically significant. However, given our small sample in the control groups, we find that only one of the differences in revisions between the treatment and control group, PP × Food, is statistically significant (at the 10% level). In the regression analysis below, we control for the size of respondents' perception gaps which uses richer data than simply testing for average differences between groups.

As shown in the fourth row of each sub-panel, there is a sizable proportion of respondents who do revise their expectations, and a sizable proportion of respondents who do not, in both the treatment and control groups. And even though the proportion of respondents who do not revise their expectations is smaller in the treatment groups (as one would expect), nevertheless many treated respondents (between 40.15% and 52.83%) do not revise their expectations. We expect the treatment information to cause individuals to update their inflation expectations if i) individuals'

[13] Throughout this section, we first present results for respondents' point forecasts at the one-year horizon, and later extend our discussion to one-year density forecasts and three-year point forecasts. The one-year point forecast's format – that is, a single number representing inflation at a one-year horizon – most closely parallels the format of the information treatments, and is our primary outcome measure to test for meaningful updating behavior.

inflation expectations are based in part on their beliefs about the measures we use in our information treatments, i.e., food and beverage prices or professional forecasters' forecasts, ii) respondents find the provided information to be credible, and iii) respondents are not fully informed about the true values of these quantities. So, understanding respondents' heterogeneity in responsiveness – and the heterogeneity of inflation expectations in general – hinges on understanding whether some consumers are relatively under-informed about the inflation-relevant information in our treatments, or whether some consumers are relatively more responsive to this information. We turn to these questions next.

3.1 Perception gap

To measure respondents' ex-ante informedness, we calculate the gap between the information prior (which we defined above as the respondent's belief about the treatment information, and which is elicited prior to the revelation of the information) and the treatment information. We refer to this difference as a *perception gap* $\Delta\omega_i$ for respondent i. We sign perception gaps such that negative perception gaps indicate overestimation (treatment information minus information prior): that is, $\Delta\omega_{i,SPF} \equiv \omega^*_{SPF} - \omega_{i,SPF}$, and $\Delta\omega_{i,Food} \equiv \omega^*_{Food} - \omega_{i,Food}$. The median perception gap in the Food treatment was -2.61, i.e., past food/beverage price changes were over-estimated by 2.61 percentage points, whereas the median perception gap in the SPF Forecast treatment was -1.04.[14]

Our analysis reveals that the perception gaps are, on average, larger for respondents in our treatment groups who revise their inflation expectations versus those who do not. The average Food treatment perception gap for respondents who revise their expectations is -7.3 percentage points versus -5.8 percentage points for those who do not (difference statistically significant at 10%), while the SPF forecast treatment perception gaps are -3.61 for revisers and -1.87 for non-revisers (significant at 5%). This suggests that ex-ante informedness about the objective information in our treatments is important for explaining heterogeneity in respondents' revisions, and will be important for explaining heterogeneity in consumer inflation expectations in general. We return to this result in more depth later.

[14] Note that in the SPF Forecast treatment, respondents were also informed about the interval containing the forecasts of 90% of the professional economists: 1.19% – 3.03%. If we instead measure the SPF Forecast perception gap as the minimum (signed) distance between a respondent's prior and this interval, then the median perception gap in the SPF Forecast treatment is 0. However, all regression results below are qualitatively similar when we use this alternative perception gap.

In order to understand how the perception gap varies by information treatment and demographics, Table 2 reports a series of OLS regressions in which we regress the perception gap onto individual characteristics and treatment dummies. For ease in interpreting coefficients, we set perception gaps – in this regression only – to be log(treatment information / information prior), which preserves sign (relative to our linear distance measure) while giving the regression coefficients an elasticity interpretation: a coefficient of, e.g., 0.1 for a given demographic group indicates a 10% larger underestimate for that group.[15]

Column (1) of Table 2 shows results from regressing the perception gap onto treatment group dummies. The constant term in this regression shows the mean perception gap for the SPF Forecast treatment group. The coefficient of -0.561 indicates that respondents in the SPF Forecast treatment group, on average, over-estimate the forecast of professional forecasters by 56%. The average perception gap is -0.561 -0.825 = -1.39 in the Food treatment, indicating average over-estimates of 139%. Both estimates are statistically different from zero (at the 1% level), and the average Food treatment perception gap is statistically different from the average perception gap in the SPF Forecast treatment.

So, we find that our survey respondents are *on average* substantially misinformed about past changes in food/beverage prices as well as forecasts of professional forecasters, even as the median respondent's information prior is within 3 percentage points of the true information in the Food treatment, and within about 1 percentage point of the true information in the SPF treatment. For both mean and median respondents, we observe larger perception gaps in the Food treatment. Also, we find no significant differences in perception gap between control and treatment groups; this indicates our randomization successfully produced information and control groups with comparable information priors.[16]

To highlight the heterogeneity in perception gap by demographics, in the remaining five columns of Table 2 we regress the perception gap on dummies for gender, older age (55 years and older), high financial literacy,[17] education (college B.A. or more), high income ($75,000 and

[15] For respondents who report a zero or negative information prior, we recode their information prior as 0.1 for the sake of calculating log(info/belief). There are only fourteen such instances.

[16] To obtain the average perception gap for the control group in the SPF Forecast treatment, one has to add in the coefficient on control to the constant (–0.0431 – 0.561 = –0.641). Similarly the average perception gap for the control group in the Food treatment is –1.358 = –0.0431 + 0.0710 – 0.561 – 0.825.

[17] Our survey included a battery of 7 numeracy and financial literacy questions. The numeracy questions were drawn from Lipkus, Samsa, and Rimer (2001), while the questions about financial literacy were slightly adapted from Lusardi (2007). We coded a perfect score on these questions as "high financial literacy," which included 31.3% of the sample. See Appendix for the questions.

above), and also treatment and control group dummies. Two results are statistically significant when we control for all demographics at once (column 6). First, we find that college-educated respondents have 20% and 34% smaller perception gaps (less over-estimation) on average than their less-educated peers in the SPF Forecast and Food treatments respectively (though the difference *between* the two treatments is not statistically significant). Second, older respondents over-estimate the past change in food and beverage prices by 18% more relative to their younger counterparts.[18] While the remaining coefficients are not statistically significant, we find that female, low financial literacy, and low income respondents are more likely to over-estimate and have larger perception gaps. In the last column of Table 2, we reject the null hypothesis that the coefficients for all five of these demographic groups are jointly equal to zero (with a p-value of 0.097 for an F-test).

3.2 Inflation Expectations Revisions and Perception Gaps

If a respondent finds the treatment information relevant and credible and if she uses the treatment information sensibly to update her inflation expectations, we expect to see an under- (over-) estimation of treatment information leading to an upward (downward) revision in inflation expectations. For our purposes, under-estimations are signed as *positive* perception gaps. Therefore, inflation expectations' revisions should be positively related to the perception gap.

Figure 2 plots perception gaps and revisions, separately for each treatment and control group. More precisely, the figure shows the mean revision by perception gap decile[19], as well as a local linear regression of mean revision and perception gap decile. Data consistent with sensible updating behavior in response to the treatment information should have the following characteristics: (1) the data points should be in either quadrants 1 or 3 (the two shaded quadrants in the figure), i.e., mean revisions should be positive for positive perception gaps, and negative for negative perception gaps; and (2) there should be a positive relationship between mean revisions and perception gaps, i.e., the spline should be upward sloping in quadrants 1 and 3. Comparing the four graphs in the first column of Figure 2 (treatment groups) with the four corresponding figures in the second column (control groups), we see two patterns. First, all of the data points are either in quadrants 1 or 3 for the treatment groups (except for one duple in the PP x SPF treatment), while a

[18] To obtain the average perception gap for older respondents in the Food Treatment, one has to add the coefficients on older and older x Food treatment (0.120 – 0.304).

[19] In cases where deciles overlap, fewer than 10 points appear on the plot. Instances of overlapping deciles are indicated by circles that are larger in size.

substantial number of data points appear in the other quadrants for the control groups.[20] Second, the spline is upward sloping and in the predicted quadrants in two of the four treatment groups (RI x SPF; PP x Food), while we observe either a flat relationship or one that is not confined to the *predicted* quadrants in the control groups. These results indicate, nonparametrically, that our treatment groups had a greater prevalence (relative to the control groups) of sensible updating, and suggest that our information experiment caused respondents to revise their inflation expectations.

3.2.1 Baseline Updating Model

We next examine updating behavior in a regression framework. We estimate the slope of a fitted line for the individual-level data underlying each of the eight panels in Figure 2, regressing the revision in inflation expectations between stages one and four, $\Delta\pi_i \equiv \pi_i' - \pi_{i,}$, on the linear-perception gap $\Delta\omega_i$. We also include a set of interacted indicator variables as regressors, allowing us to estimate an updating slope separately for each treatment cell. Specifically, we estimate (separately for the Food and SPF Forecast treatments) the following regression:

$$\Delta\pi_i = \alpha_1 + \alpha_2 * I_{PP,i} + \beta_{RI}\left(T_{info,i} * I_{RI,i}\right) + \beta_{PP}\left(T_{info,i} * I_{PP,i}\right) + \gamma_{RI}\left(T_{info,i} * I_{RI,i} * \Delta\omega_i\right) + \gamma_{PP}\left(T_{info,i} * I_{PP,i} * \Delta\omega_i\right) + \epsilon_i \quad (1)$$

where $I_{PP,i}$ is an indicator for respondent *i* answering the "Prices you Pay" question, while $I_{RI,i}$ indicates the "Rate of Inflation" question. $T_{info,i}$ is an indicator that equals one if respondent *i* was in the treatment group, i.e., the treatment information was revealed to the respondent, and zero otherwise. Note that $\Delta\pi_i$ and $\Delta\omega_i$ are, respectively, the same variables that were plotted on the y-axes and x-axes in Figure 2. In this specification, α_1 is a constant capturing average updating in the RI question for respondents in the control group, and α_2 is a constant that similarly captures average (differences in) control-group updating for PP respondents (relative to RI respondents). Then, the sum $\alpha_1 + \beta_{RI}$ shows average updating for RI question respondents in the treatment groups with a zero perception gap. Similarly, the sum $\alpha_1 + \alpha_2 + \beta_{PP}$ is the average updating for PP question respondents in the treatment group with a zero perception gap. Inclusion of $\alpha_1$and α_2 in equation

[20] In the data, we also find a greater percentage of non-zero updating respondents in the shaded quadrants for the treatment groups than for the control groups: For example, 74.3% of non-zero updating in the PP × Food treatment group happens in the shaded quadrants, as compared with 53.3% for the PP × Food control, and 70.4% for the RI × SPF treatment group, as compared with 53.7% for the RI × SPF control.

(1) allows us to control for the revisions that are attributable to the other questions asked in the survey (as well as the mere act of taking the survey).

The coefficients γ_{RI} and γ_{PP} are our main coefficients of interest. They show updating behavior with respect to perception gap size, in RI and PP responses respectively, i.e., they provide an estimate of the causal effect of our information treatments on inflation expectations' revisions. For revisions to be consistent with meaningful expectation updating, as described above, we expect estimates of gammas to be non-negative.

Table 3 presents results from this baseline regression. We use weighted least squares to estimate equation (1), to ensure that our estimates are robust to the inclusion of outliers.[21] Focusing first on updating in one-year point forecasts for the Food group (column 1), we find a significant γ coefficient for PP responses (γ_{PP}): A perception gap of 10 percentage points in the Food treatment causes a revision of 0.35 percentage points for the PP question (significant at 5%). This estimate implies that a standard deviation increase in the perception gap results in a revision of 2.95 percent of a standard deviation (of the baseline expectations). Notably, we do not find a significant effect of the Food perception gap on RI responses: γ_{RI} is positive but not statistically different from zero.

Also, estimates of α are not significantly different from zero, i.e., there are no significant revisions in inflation expectations in the control group. Moreover, β_{PP} is negative but not statistically different from zero, which indicates that there is no significant effect of the Food information treatment on PP responses (relative to control group responses) *other than* what is explained by the size of respondents' perception gap size.

We find an analogous result for one-year point forecast updating in the SPF Forecast group (column 2). First, the estimate of γ_{RI} is significant: A 10 percentage point perception gap in the SPF Forecast treatment causes a 1.25 percentage point revision in forecasts for the RI question (significant at 1%). This estimate implies that an increase in the SPF Forecast treatment perception gap of one standard deviation leads, on average, to a revision that is 11.85% of a standard deviation of the baseline RI expectations. Second, γ_{PP} is positive but not significantly different from zero, i.e., the perception gap has a positive but insignificant effect on PP response. Third, both β coefficients in column 2 are not significantly different from zero; in particular, the insignificant β_{RI} coefficient indicates there is no effect of the SPF Forecast information treatment on RI responses

[21] We use weighted least squares (robust regressions) to estimate all expectation-updating regressions in this paper. Each of these regressions uses the minimum tuning constant possible for all regressions within a table. In other words, tuning constants (which determine the extent to which outliers are downweighted) are the same within tables, but not necessarily between tables. Therefore estimated coefficients may not directly comparable between tables.

(relative to control group responses) *other than* what is explained by the size of respondents' perception gap size.

It is notable that the Food treatment significantly affects PP responses (and not RI responses), whereas the SPF Forecast treatment affects RI (and not PP). That is to say, information about prices for "food and beverages" *only* significantly affects expectations about "the prices of things you usually spend money on", and information about "future inflation" *only* significantly affects expectations about the "rate of inflation." We discuss this divide in more detail in the next section.

We next estimate the baseline specification, but use the revisions in the fitted mean of the one-year density forecasts (columns 3 and 4 of Table 3), and three-year point forecasts (columns 5 and 6) as our dependent variable. We find our γ coefficients have a pattern similar to what we had seen for the one-year point forecasts: For both one-year density forecasts and three-year point forecasts, the Food treatment affects significantly only the PP question, and the SPF Forecast treatment affects significantly only the RI question.

In the last two columns of Table 3, using the fitted mean of the one-year density forecasts, we explore the relationship between revision of inflation expectations and the precision of baseline inflation expectations. In a Bayesian framework, *ceteris paribus*, respondents who are more uncertain about future inflation should be more responsive to the treatment information.[22] Using the variance obtained from fitting a beta distribution to each respondent's one-year baseline density forecast, we define a dummy variable, *Uncertain*, that equals 1 if the respondent's baseline variance is above the sample median. We add two new terms to equation (1): the RI and PP Gap terms are interacted with the dummy, *Uncertain*. Thus γ_{RI}, the coefficient on $\left(T_{info} * I_{RI} * Gap\right)$, shows updating behavior with respect to perception gap size in RI for low-uncertainty respondents, while $\gamma_{RI} + \gamma_{RI-U}$ shows updating behavior with respect to the perception gap for high-uncertainty respondents. A positive γ_{PP-U} and γ_{RI-U} would be consistent with Bayesian updating. That is indeed what we find: estimates of γ_{PP} and γ_{PP-U} in column (7) suggest that the Food treatment affects PP responses, but only for high-uncertainty respondents. Similarly column (8) shows that the

[22] In a Bayesian updating model, for beliefs that are characterized by the beta distribution, the posterior (updated belief) is

$$Posterior = \frac{\frac{1}{Variance(Prior)}}{\frac{1}{Variance(Prior)}+\frac{1}{Variance(Info)}} Prior + \frac{\frac{1}{Variance(Info)}}{\frac{1}{Variance(Prior)}+\frac{1}{Variance(Info)}} Info.$$

Then, the relative weight placed on the information is $\frac{Variance(Prior)}{Variance(Info)}$, i.e., responsiveness to information should be directly proportional to the uncertainty in baseline inflation expectations.

SPF Forecast treatment affects significantly the RI question, but only for high uncertainty respondents. These results indicate that the updating patterns shown in columns (3) and (4) are primarily driven by the revisions of high-uncertainty respondents.

3.2.2 Non-Linear Updating Model

Next, we test for non-linearity in the updating slopes. To do so, we estimate a separate γ coefficient for each tertile of perception gap size. This regression has the same constant terms as our baseline regression, but γ_{RI} and γ_{PP} are allowed to vary by the tertile of the absolute perception gap, i.e., we estimate separately γ_{RI-T1}, γ_{RI-T2}, and γ_{RI-T3} (with three similar γ_{PP} coefficients) where T1 denotes the first (lowest) tertile, T2 denotes the second (middle) tertile, and T3 denotes the third (highest) tertile.[23]

Results from this tertile-wise regression are presented in Table 4. The main result to highlight is the concentration of significant updating in the *middle* tertile of respondents, with a few exceptions where we also observe significant updating for the *highest* tertile. For example, at the one-year horizon for the SPF Forecast treatment (in column 2), no γ coefficients from the *first* tertile of absolute perception gap are significant. Instead, significant (and positive) updating is concentrated in higher tertiles for both the RI and PP question. Likewise, for one-year point forecast updating in the Food treatment (column 1) we find a statistically significant coefficient only for the middle tertile in the PP question.

We find that the one-year *density* forecast results (columns 3 and 4) are generally consistent with the one-year point forecasts results. Density forecast updating is positive and strongly significant in either the middle or the upper tertiles in the SPF treatment for both the PP and RI questions, and for the middle tertile of the Food treatment for the PP question. Also in Table 4, we see that three-year point forecast results (columns 5 and 6) show significant updating to be concentrated in middle tertiles. We only find significant three-year point forecast updating for the PP question, for both the Food treatment and the SPF Forecast treatment.

3.2.3 Heterogeneity in Updating

We next examine heterogeneity in updating behavior. Before studying this heterogeneity in a regression context, we document the differences in baseline inflation expectations, baseline

[23] Tertiles are calculated separately for each information treatment. By defining tertiles using absolute revisions, we treat positive and negative perceptions gaps of the same size symmetrically in the regressions.

uncertainty (in inflation expectations), perception gaps, and revisions – by gender, education, income, financial literacy, and age. We present these summary results in Table 5. We see first that female, low-income, low-education, low-financial literacy, and older respondents report higher baseline inflation expectations for both RI and PP. For example, females report a mean RI (PP) point forecast of 6.8 (7.1) percent, versus a mean forecast of 4.4 (5.9) for males.[24] That is, we see similar demographic patterns in expectations for future inflation, as we previously saw in Table 2 for informedness about current, inflation-related facts. For ease of comparison, the bottom panel of Table 5 also shows this heterogeneity in perceptions gaps,[25] and we observe that female, low-education, low-financial literacy, and low-income respondents have larger perception gaps in magnitude, though the differences are not statistically significant.[26] Meanwhile, as shown in the distribution of (pooled) baseline inflation expectations in Appendix Table A1, the demographic differences in average expectations are indeed driven by a greater percentage of female, low-income, low-education, and low-literacy respondents occupying the (high-expectation) right tail of the expectations distribution: for example, 54% of female respondents have expectations in the interval [5+), compared to 42% of male respondents in that range (proportions different at the 1% level, using a Chi-squared test).[27]

Table 5 also indicates that, especially for the RI question, the same demographic groups (i.e., female, low-education, low-literacy, and low-income respondents) not only report higher baseline inflation expectations, but they also have more uncertain baseline inflation expectations. For example, female respondents for the RI question have a mean (median) individual density forecast variance of 25.1 (4.0), while male respondents have an average forecast variance of 9.7 (2.4) (differences significant at the 10% (15%) level, using pairwise t- (median) tests).

The descriptive patterns of updating in Table 5 also indicate different *revisions* in inflation expectations (though differences are not all statistically significant). For example, females revise down their RI (PP) inflation expectations, on average, by 2.6 (2.3) percentage points, compared to a downward revision of 0.7 (1.4) percentage points for males. As a result of these different revisions,

[24] The lack of statistical significance in these gender differences is possibly a result of the small sample size in each of these cells.
[25] For simplification, we combine the perception gap for both treatments here. We obtain qualitatively similar patterns by demographics in perception gaps for both information treatments.
[26] This finding is similar to the results about forecast accuracy by demographics presented in Bryan and Venkatu (2001a), Souleles, 2004; Anderson (2008), Pfajfar and Santoro (2010), and Madeira and Zafar (2011).
[27] Systematically higher inflation expectations among these demographic groups have also been found in the literature: see Jonung (1981), Bryan and Venkatu (2001a, 2001b), Lombardelli and Saleheen (2003), and Blanchflower and Coille (2009).

the revised average inflation expectations converge within gender, income, education, and financial literacy groups after the information treatment. This pattern could be a consequence of demographic differences in (1) perception gaps, (2) precision of baseline inflation expectations (uncertainty), and/or (3) information-processing rules; to help disentangle these factors we return to regression analysis next.

We begin our regression analysis of heterogeneity by testing for gender differences in updating. Previous research on inflation expectation updating has found mixed results: Burke and Manz (2010) do not find significant differences by gender in information processing, but Madeira and Zafar (2011) find significant differences in expectation updating by gender and other demographic characteristics.[28]

We add two new intercept terms to equation (1), $\alpha_3 * Female_i$ and $\alpha_4 * Female_i * I_{pp,i}$, and estimate each γ coefficient separately for male and female respondents.[29] Table 6 presents the results for this baseline regression by gender. Nearly all significant updating that appears in Table 6 is driven by female respondents. In particular, we find that the effect of the SPF Forecast treatment on RI responses is driven by female respondents for one-year point and density forecasts, and the effect of the Food treatment on PP responses is driven by female respondents for one-year density forecasts and three-year point forecasts.[30] Moreover, the effect is quite large: a standard deviation increase in the SPF Forecast treatment perception gap leads on average to a revision of nearly 41 percent of a standard deviation of the baseline PP expectations. There is also a weak effect among females for the SPF Forecast treatment on PP responses at the 1-year point forecast horizon (column 2). These results indicate that, even when we condition on perception gap size, female respondents exhibit greater responsiveness to our treatment information than male respondents. In our discussion below, we examine whether this is a result of higher baseline uncertainty for females, different information-processing rules, or a combination of both.

[28] In other contexts of belief-updating, Mobius, Niederle, Niehaus, and Rosenblat (2011) find significant gender differences in information processing, while Wiswall and Zafar (2011) do not find gender differences.

[29] The intercepts from the baseline specification, α_1 and α_2, now capture average updating for male respondents in the control group, while α_3 and α_4 capture average updating differences between males and females in the control group. These four intercepts are additive: For example, average updating for females in the PP control group is given by $\alpha_1 + \alpha_2 + \alpha_3 + \alpha_4$, while average updating for *males* in the RI control group is given by α_1 only.

[30] Column 1 estimates of γ_{PP-F} (female) and γ_{PP-M} (male) are not significant, whereas the estimate of γ_{PP} in column 1 of Table 3 is marginally significant. This is probably a result of lower statistical power in Table 6 due to smaller cell sizes, and different downweighting of outliers in the two regressions (we use different tuning constants in Tables 3 and 6; see footnote 21). Meanwhile, OLS results are qualitatively similar between Tables 3 and 6, and OLS results similarly show that all significant updating in Table 6 is driven by female respondents.

We next test for heterogeneity in updating by other demographics. To simplify, we focus on updating at the one-year point forecast horizon. We estimate updating effects by some of the individual characteristics from Table 2 – income, age, education, and financial literacy. We also use an additional variable, based on the following question asked immediately after eliciting final inflation expectations (in stage 4): "*To what extent is your answer [to the PP or RI question] over the next twelve months the same or different because of the information provided to you [in the Food or SPF information treatment]?*" Responses are given on a 7-point scale (with a higher number indicating a larger effect of provided information); these responses are coded such that roughly 40% of respondents are flagged as "info-affected" (the cutoff for "info-affected" is 5 or more points out of 7).

For each of the five characteristics discussed above, we estimate the following regression (for simplicity, we drop the respondent subscript):

$$\begin{aligned}\Delta\pi = \alpha_1 + \alpha_2 * I_{PP} + \alpha_3 * C + \alpha_4 * C * I_{PP} + \beta_{RI}\left(T_{info} * I_{RI}\right) + \beta_{PP}\left(T_{info} * I_{PP}\right) \\ + \gamma_{RI}\left(T_{info} * I_{RI} * (1 - C) * Gap\right) + \gamma_{PP}\left(T_{info} * I_{PP} * (1 - C) * Gap\right) \\ + \eta_{RI}\left(T_{info} * I_{RI} * C * Gap\right) + \eta_{PP}\left(T_{info} * I_{PP} * C * Gap\right) + \epsilon,\end{aligned} \tag{2}$$

where C is a binary variable that represents one of the five individual-level characteristics – income, age, education, financial literacy, and "info-affected". The γ coefficients capture updating behavior for individuals without a given characteristic C, while the η coefficients capture updating behavior for individuals with "C". Each of these coefficients (γ and η) takes on a RI or PP subscript, as previously, to show distinct updating behavior for the two different question-texts. All characteristics' regression results are presented together in Table 7, with two columns for each characteristic such that we can again identify effects separately for the Food treatment and SPF Forecast treatment.

We analyze updating one characteristic at a time. While it would be ideal to test for updating differences by all of these characteristics simultaneously – for example, gender differences in updating could partly be a result of gender differences in income, education, or financial literacy – our sample size prevents us from exploring these channels. Fortunately, the correlation between each of these demographic variables is small.[31]

[31] The highest correlation that we observe is 0.29 between high income and college. Female and financial literacy has a correlation of -0.19; high financial literacy and college education has a correlation of 0.13; female and income has a correlation of -0.12. All other correlations are smaller than 0.1 in magnitude: -0.06 between female and college; -0.02 between female and older; 0.09 between high financial literacy and high income; -0.04 between older and college; -.03 between older and high income; and -0.08 between older and high financial literacy.

The first two columns of Table 7 present estimates of equation (2) by income (that is, C in the equation is High Income, defined as respondent's annual income being over $75,000). We find that Food × PP updating is significant only among lower income individuals; this is consistent with food and beverages receiving more weight in a basket of "things you usually spend money on" – that is, the focus of the PP question – as income decreases (McGranahan and Paulson, 2006). Meanwhile, we find that both high and low income individuals update significantly for SPF Forecast × RI. For high income individuals this updating coefficient is especially large, roughly four times larger than that for low income respondents (the estimates imply that an increase of one standard deviation in the SPF Forecast treatment perception gap leads to revisions of about 61% and 12% of one standard deviation of the baseline RI expectations, for high and low income individuals respectively). In columns (3)-(6) of Table 7, we observe an interesting split by our "ability" measures – education and financial literacy – in the SPF treatment. Whereas college-educated respondents update their RI expectations more than their less educated peers in the SPF Forecast treatment, high-financial literacy respondents update *less* than their low-financial literacy peers. Particularly interesting is the updating we observe among college-educated respondents for the PP question in the SPF Forecast treatment. High-education respondents update their PP expectations significantly in the SPF treatment – the estimate corresponds to a revision that is 40 percent of one standard deviation of baseline PP expectations in response to a standard deviation increase in the SPF Forecast treatment perception gap – whereas less-educated individuals do not.

Columns (7)-(8) of the table show updating by age. We find that older respondents are the only demographic subgroup among the four we examine that exhibits significant updating in the Food treatment. Moreover, older respondents are the only group for which we observe statistically significant Food × RI updating. Meanwhile, we find that both older and younger respondents update significantly for SPF Forecast × RI, with the coefficient for young being nearly twice as large as that for older; these coefficients imply that, on average, revisions of 12% and 6.4% of one standard deviation of baseline RI expectations result from a one standard deviation increase in perception gaps, for younger and older respondents respectively. We also obtain a statistically significant but negative estimate for SPF Forecast × PP updating for the young; this suggests counter-intuitive updating on their part for the PP question in response to the SPF treatment.

Estimates of the specification with the "info-affected" characteristic are shown in columns (9) and (10) of the table. If individuals in the information treatments are indeed changing their PP or RI expectations in response to the provided information, we expect to see stronger updating among respondents who report that the provided information "affected" their final PP or RI expectations. That is, we expect to see the magnitudes of η_{PP} (η_{RI}) to be greater than those of γ_{PP} (γ_{RI}), and expect them to be statistically positive. This is indeed the case: Individuals who report *not* being "info-affected" show no significant updating, whereas "info-affected" individuals exhibit the same updating behavior seen in our earlier baseline regressions, including a strong effect for the Food × PP, SPF Forecast × PP, and SPF × RI treatment cells. Moreover, point estimates of responsiveness to information among "info-affected" respondents (η) are six to ten times larger than those in the baseline regression results (Table 3, columns 1 and 2).

4. Discussion

In this section we discuss three main results derived from the analysis in Section 3. We also discuss the implication of our results for the modeling of consumer inflation expectations.

RESULT 1: Respondents' information priors have a strongly right-skewed distribution: while many respondents are well-informed, many other respondents substantially overestimate objective, inflation-relevant facts. Baseline inflation expectations are similarly distributed. There is also significant heterogeneity in information priors by type of information treatment.

We find that average perception gaps (the treatment information *minus* information priors) are negative and substantial – indicating overestimation of objective measures. On the one hand, a sizable proportion of respondents have information priors that are in line with the treatment information: 41% of respondents have perception gaps of 2 percentage points or less, and 22% of respondents have perception gaps of 1 percentage point or less. On the other hand, many respondents are substantially less informed: in fact, 38% of respondents expect professional forecasts of next-year inflation to be 5% or more, while our SPF benchmark was only 1.96% and had not been as high as 5% since 1984. Respondents' overestimates are even larger when we ask about food and beverage price inflation: 40% of respondents believe past-year food and beverage price inflation was 7% or more, while the published measure was only 1.39%, and has not risen as high as 7% since 1981. Thus the distribution of information priors is strongly right-skewed in both information treatments.

Similarly, we find that the distribution of respondents' expectations of future inflation is right-skewed. For example, the mean of baseline RI one-year point forecast responses is 5.6%, while the median is 3.0%, and the corresponding PP mean and median are 7.0% and 5.0%. Indeed, average consumers' subjective expectations of inflation have been consistently higher than actual inflation in recent periods (Geroganas, Healy, and Li, 2011),[32] even though *median* consumer inflation expectation survey responses generally track realized inflation (Thomas, 1999). Our results indicate that these right-skewed expectations may in part be due to a skewed distribution of perceptions about objective measures of realized inflation (and, hence, differences in information sets)– an explanation that we discuss more fully in Result 3 below.

We also find that perception gaps differ by type of inflation: When we ask respondents about food and beverage price inflation, the median perception gap is -2.61 (a 188% overestimation), whereas for SPF Forecast inflation the median perception gap is substantially smaller, -1.04 (a 53% overestimation). There are several possible explanations for larger perception gaps in the Food treatment. First, when respondents are asked about past changes in food and beverage prices, their responses are likely to suffer from recall bias, as respondents are likely to recall items for which perceived price changes were most extreme (Bruine de Bruin et al., 2011b). Second, frequency bias may lead respondents to report food inflation perceptions based on the frequency of purchase rather than the total dollar expenditures. Given that prices of frequently-purchased items inflate faster, this would bias their perceptions upwards (Georganas, Healy, and Li, 2011).[33]

RESULT 2: Respondents, on average, update their expectations in response to information we provide to them, and do so sensibly. The expectations distribution thus converges toward its center. Furthermore, updating patterns depend on the type of inflation – own-basket or overall – about which respondents are asked, and on the type of treatment information.

If respondents' inflation expectations are based in part on their perceptions of past food/beverage price changes and on experts' inflation forecasts, then our information treatments should lead respondents to revise their expectations if they are not fully informed about these benchmarks ex ante. That is, in fact, what we find holds on average in Section 3 (Figure 2; Tables 3 and 4). We find that updating in the treatment groups does significantly differ from updating in the control groups, and that this updating is explained by the size of respondents' perception gaps. This

[32] Furthermore, note the average of SPF cross-sectional medians of one-year-ahead inflation forecasts since 2000 has been 1.95%; the average of Michigan Survey of Consumers cross sectional medians of one-year-ahead inflation forecasts over the same period has been 3.01%.

[33] Whereas the perception gap in food prices may be partly explained by factors such as personal experiences, these patterns suggest that respondents are not fully informed of the objective inflation measures used in our information treatments. In particular, it is hard to explain the perception gap in the SPF Forecast treatment except as a lack of knowledge on the part of respondents.

importantly suggests that models explaining heterogeneous consumer inflation expectations need to incorporate consumers' different information sets, since we find that the distribution of expectations converges toward its center when all respondents are given information that, at baseline, only some respondents were well-informed of.

Respondents' updating behavior is also sensible, in the sense that it reflects the *sign* of respondents' perception gaps and also the *size* of these gaps. Moreover, we observe greater responsiveness to information for more uncertain respondents, consistent with a Bayesian updating model.[34] Furthermore, by finding consistently significant results in both information treatments, we find direct evidence that consumers can take into account forecasts of experts (as modeled in Carroll, 2003) and past price changes (Garner, 1982; Hey, 1994, Lanne et al., 2009) in their own forecasts, at least when they receive such information.[35] Note that while responding to information about past food and beverage prices can be rationalized in a model of adaptive learning, responding to forecasts of experts is not consistent with such a model.

We, however, find that information about food and beverage prices causes consumers to update expectations primarily for "prices you pay", whereas information about inflation forecasts causes consumers to update expectations primarily for the "rate of inflation". This result may be intuitive but is important: it indicates that consumers believe the price changes in their own consumption basket to be different from the overall rate of inflation, and suggests that consumers use different types of information to update their expectations about both. In particular, it is notable that food and beverage price information, at least as we have presented it here, has less relevance for consumers' expectations of overall inflation. This may be because some consumers have limited understanding of overall inflation, but also may be indicative that consumers are distinguishing (either correctly or incorrectly) between different types of information for what they perceive to be different types of forecasts.

The fact that the Food and SPF Forecast treatments' effects are primarily seen in, respectively, PP and RI responses, also suggests that respondents are processing the treatment information thoughtfully, rather than unconsciously anchoring to the new information (Tversky and

[34] Nevertheless, we are unable to investigate how respondents' revisions compare to some benchmark, such as Bayesian updating. This is because doing so requires richer, hard-to-acquire data on the underlying distributions of the information as well as detailed information on respondents' consumption bundles, such as the proportion spent on food and beverages.

[35] However, expectation updating may be asymmetric depending on whether the information is higher or lower than respondents' priors (Eil and Rao, 2011). We might have observed different expectation updating if food and beverage price inflation had been higher than respondents' priors, rather than lower. But in our setting, food and beverage price inflation was lower than prior beliefs for 94% of respondents.

Kahneman, 1974). While the direction of respondents' revisions is consistent with a naïve anchoring explanation, the differential effects by information treatment and expectation type are harder to reconcile with anchoring.

Furthermore, we find that the distribution of expectations converges toward its center, and the standard deviation of the expectations distribution falls from 7.0 percentage points to 4.5 for RI, and from 7.0 to 5.4 for PP. In fact, respondents' inflation expectations converge toward being near, or within 1 percentage point of, the actual realized CPI inflation between January 2011 and January 2012 (2.93%). Caution is warranted in using an ex-post realized outcome as a benchmark for accuracy of ex-ante expectations, since (1) inflation outcomes are uncertain, such that a single year's inflation realization may not coincide with an *objective* ex ante expectation,[36] and (2) respondents' point forecasts may refer to various statistics (i.e. mean, median, mode, or others) of their *subjective* probability distributions (Engelberg, Manski and Williams, 2009), and (3) respondents' expectations for PP (own-basket inflation) may use a different basket of goods than the CPI. Nevertheless, we find at the baseline that 39.6% of RI responses and 35.0% of PP responses are within 1 percentage point of ultimately realized CPI inflation, whereas post-treatment these percentages improve to 55.6% and 52.8%, respectively.

Our results for updating of one-year inflation expectations point forecasts also extend to the updating of the one-year density forecast and the three-year point forecast. We find that updating of the mean of the year-ahead density forecast is generally consistent with one-year *point* forecast updating, but that the updating coefficients are sometimes smaller. This is consistent with density forecasts allowing individuals to respond to additional information by changing one part of the distribution without translating the entire distribution along the axis.

In contrast, we find that updating at the three-year point forecast horizon is sometimes stronger than updating for one-year point forecasts. This result is difficult to interpret. One possible explanation for the strong three-year point forecast updating would be if short-term inflation expectations are affected by recent experiences, but long-term expectations are not. That is, if respondents have recently had salient experiences with price changes, such that they hold strong priors about next-year inflation, they might update less at the one-year horizon because of the strength of their priors (baseline inflation expectations), but might hold weaker priors at the three-year horizon, and hence update more.

[36] That is, a single year's realization is an inconsistent, albeit unbiased, estimator for the mean of the objective probability distribution.

We also, however, find that a substantial proportion of the respondents in the treatment groups do not revise their expectations. While these respondents have, on average, smaller perception gaps (that is, the information comes as less of a surprise to them), their perception gaps are still quite large: the mean (median) perception gap in the Food treatment is -5.80 (-3.61) percentage points, while the mean (median) gap in the SPF Forecast treatment is 1.86 (-1.04). For the most part, these "non-revisers" also do not revise their density forecast. If we regress an indicator for non-revision on treatment dummies and our demographics of interest, along with a level term for size of perception gap and perception gap squared, we also find that non-revision is generally difficult to predict:[37] while female respondents are about 12% less likely to be non-revisers (significant at the 5% level), we otherwise have no strong predictors for non-revision.[38] We conclude that some respondents likely found the treatment information to be either not credible, or not relevant to their process of forming inflation expectations.

RESULT 3. Females are on average more responsive to our information treatments relative to males. There is other heterogeneity in updating behavior by respondents' education, age, and income. This highlights the importance of consumers using heterogeneous information-processing rules.

The regression results (Table 6) for updating behavior by gender are striking: almost without exception, all significant updating behavior occurs for female respondents. Since we find only weakly different perception gaps by gender (Table 2 and lower panel of Table 5), and since we find that female respondents are more responsive to information than males are, even after we control for the size of perception gaps (Table 6), these results suggest that males and females use different information-processing rules, and/or that females are more uncertain than men about future inflation expectations at the baseline.[39] In Table 5, we find evidence of higher uncertainty (in baseline expectations) for females for RI (rate of inflation), but no significant gender difference in baseline

[37] Here, we restrict our sample to treatment-group respondents who had perception gaps of greater than 1 percentage point, since this offers a sample of respondents we could expect to revise their expectations. Regression results are qualitatively similar if we use a cutoff of 0 percentage points, or 2 percentage points.

[38] It should be pointed out that we find no evidence of differences in preferences for sources of information for the revisers and non-revisers. Our survey instrument included the question: "*When trying to come up with your answers to the questions about the prices of the things you usually spend money on [or "rate of inflation/deflation"], how much did you think about the information you received from the following sources?*" on a 1-7 scale. Options included TV/Radio; Newspapers; Internet; Financial advisors; Co-workers; Family, Friends; Shopping experience. Differences in how much relevance is attached to each one of these sources are not statistically different for the two groups of respondents.

[39] Our findings may also be consistent with the economics and psychology literature that finds that men are more (over)confident than women (Barber and Odean, 2001; Niederle and Vesterlund, 2007). These studies imply that, controlling for the information content of the signal, men respond less to information.

uncertainty in PP (prices you pay). This mixed result suggests that differences in ex ante uncertainty may be partly responsible for gender differences in expectation updating, but that differences in information-processing rules likely play an important rule as well.

Estimating gender-specific information-processing rules is beyond the scope of this paper, but our data provide some suggestive evidence on this question. For example, we find that female respondents are significantly (at the 1% level) more likely than males to answer that they "*thought a lot about…the price of groceries*" when initially reporting their inflation expectations, which suggests that they weight types of information differently than men. We also find that female respondents are significantly (at the 1% level) more likely to "*think about the information [they] received from…family and friends*" when reporting inflation expectations than males are. On the other hand, we find no significant gender difference in preference for some other information sources, such as "*shopping experience*," "*TV and radio*," or "*newspapers and magazines*". Therefore we find that females gather inflation-relevant information, and weight this information, differently than males.

Also notable is the finding that older respondents are less responsive to the SPF information treatment in updating their "rate of inflation" expectations than the young. This is consistent with models based on learning-from-experience (Malmendier and Nagel, 2009; Madeira and Zafar, 2011). These models posit that individuals are influenced by data realized during their lifetimes, and hence they imply that the young should rely to a greater extent on extrapolation of recent reports of inflation data – as we find. However, we also find a result that is harder to reconcile with learning-from-experience: younger respondents are *less* responsive to the Food treatment in their updating of PP expectations.

Next, our results by education provide insight into how different types of information affect updating for different types of inflation expectations. In the SPF information treatment, we find that both college-educated and non-college-educated respondents update their RI (rate of inflation) expectations,[40] but only college-educated respondents update their PP (prices you pay) expectations. This is somewhat surprising because, as shown in Table 2, college-educated respondents have smaller average perception gaps (Table 2). That is, the information content of our treatments is, on average, larger for non-college-educated respondents. This differential updating by education suggests that translating between information about inflation on the one hand, and price changes in

[40] The effects are, however, quite different: a standard deviation increase in the SPF Forecast treatment group leads to a revision of 32.4% (11.9%) of a standard deviation of baseline RI expectations for the college-educated (non-college-educated) respondents.

one's own consumption basket on the other hand, requires some sophistication. This may be because of respondents' different familiarity with the concept of inflation, or because of differences in how respondents process new information. In either case, these differences by education have important implications for our understanding of inflation expectations in general: low-education consumers' expectations seem to be shaped by factors *other* than information about published overall inflation indexes.[41]

Furthermore, despite no significant differences in average perception gaps by income, we find the notable result that lower-income respondents update significantly their "prices you pay" expectations in the Food treatment, while higher-income respondents do not. This is consistent with food and beverage purchases consistently making up a larger percentage of lower-income consumers' overall spending (McGranahan and Paulson, 2006).[42] Equally notable is the finding that both high and low income individuals update significantly their "rate of inflation" expectations in the SPF Forecast treatment. Together, these two results suggest that both high-income and low-income respondents successfully translate between the provided information and expected changes in different consumption baskets.

Looking across demographic groups, we also find that female, low-income, low-education, low-financial literacy, and older respondents have higher baseline inflation expectations than their counterparts, and disproportionately occupy the high-expectations right tail of the inflation expectations distribution (Table 5, Appendix Table A1). Importantly, we find (Table 2, Table 5) that these same groups also on average have higher perception gaps (that is, less ex ante informedness about the objective information in our treatments). Thus, we can offer a novel explanation for the right-skewed distribution of inflation expectations and the persistently high expectations of these demographic groups in particular: high inflation expectations may simply be a result of under-informedness about objective, inflation-relevant information.

The case of female respondents may be particularly encouraging for policy makers seeking to manage the high-expectation tail of the expectations distribution, since we find that females are also more responsive to the new information in our treatments. And, while not all of these high-expectation, high-perception-gap groups exhibit greater responsiveness to new information in a

[41] Meanwhile, we do not find any conclusive differences in updating by financial literacy, which is somewhat at odds with Burke and Manz (2010) and Bruine de Bruin et al. (2010a), who find that financial literacy is related with the expectations formation process.

[42] This result can also be explained if low-income respondents believe, either correctly or not, that the co-movement between food prices and the prices of other items in their basket is higher than what other respondents believe it is.

regression analysis where we control for their larger perception gaps (Table 7), we nevertheless find that these groups have larger average revisions in expectations, towards the center of the expectations distribution, than their counterparts (Table 5). That is to say, while these groups may need larger information shocks (than their counterparts need) in order to generate a unit revision in their expectations, these groups' perception gaps are large enough that providing true, inflation-relevant information to them leads to larger (sensible) expectations revisions than their counterparts exhibit.

Our updating results by "info-affected" help confirm our basic understanding of respondents' expectation updating process. These results are a consistency check in support of our baseline model, in which updating is a function of the provided information. Furthermore, these results indicate that respondents are aware of the provided information's effects on their revisions, which indicates that the updating process we observe here is a self-aware process, rather than subconscious belief updating (Hawkins, 1970). In particular, a naïve anchoring explanation for our observed updating (Tversky and Kahneman, 1974), which is usually explained as a subconscious process, is hard to reconcile with our "info-affected" results.

Overall, our analysis suggests that the demographic (and, in particular, gender) differences in inflation expectations that we observe are a result of both demographic differences in information sets and information-processing rules. Our findings partially validate recent modeling work that explains heterogeneous consumer inflation expectations through heterogeneous information sets and different updating rules (Nagel and Malmendier, 2010; Madeira and Zafar, 2011). These results also underscore the need for empirical work to study this heterogeneity in detail, at the individual level.

More generally, our findings also inform us about how consumer inflation expectations should be modeled. The literature on consumer inflation expectations broadly consists of three different models of expectations formation: rational expectations, where consumers use all available information efficiently when making a forecast; Bayesian updating, where consumers use all available signals and past realizations to make a forecast, and; adaptive expectations, where consumers use past realizations to make a forecast (the learning-from-experience model, as in Malmendier and Nagel, 2010, is a special case of adaptive expectations). In a rational framework (as in Muth, 1961), since all information known at the time of the forecast should already be incorporated in the forecasts, our information experiment (which provides publicly available information) should have had no systematic effect on individuals' forecasts. Therefore, we can reject perfect rationality for the average respondent. Imperfect knowledge about objective measures of inflation is, however, consistent with models of bounded rationality, such as sticky expectations or rational inattention (as in Barsky and Kilian, 2002; Mankiw and Reis, 2002; Carroll, 2003; Ball, Mankiw and Reis 2005). Our finding that consumers respond to information about past food prices is consistent with adaptive learning as well as Bayesian learning. That older respondents are less responsive to the SPF Forecast information (in other words, there is an age coefficient for how much people update) is consistent with a Bayesian learning model or a learning-from-experiences model; a standard adaptive learning model does not predict people's age to affect their updating. However, the fact that we find that additional information not present in past inflation rates matters – that is, forecasts of professional experts – is more consistent with Bayesian learning than with a learning-from-experiences model. Therefore, we conclude that are results are broadly consistent with Bayesian learning.[43]

5. Conclusion

A crucial aspect of monetary policy is managing inflation expectations. However, there is limited understanding of how individuals form these expectations – a primary question for economists and policy-makers. This paper, using a survey with an embedded information experiment, attempts to shed light on this question by exploring the causal determinants of inflation expectations. We find that respondents, on average, are not fully informed about past as well as future macroeconomic measures, and when provided with new inflation-relevant information, they update their inflation expectations. Moreover, the updating is meaningful in the sense that, on average, it is: (1) in the direction of the signal, (2) proportional to the strength of the signal (i.e., the revealed perception gap), and (3) greater when the baseline expectations are less precise.

We also find substantial heterogeneity by respondent characteristics in how fully informed respondents are about objective inflation measures, and in their updating behavior. Female, low income, less educated, lower financial literacy, and older respondents have larger perception gaps and tend to have higher expectations of future inflation. Therefore, our findings suggest a new explanation for these systematically high expectations previously found in the literature, whereby high expectations may be due in part to missing or inaccurate information about objective measures of actual inflation. Furthermore, we find that some of these same demographic groups – females in

[43] Since our design consists of only round of revisions, we are unable to say much about whether consumers have time-dependent rules (Carroll, 2003; Branch, 2007). A design with multiple rounds of revisions may also shed light on the speed of consumers' updating.

particular – tend to update their inflation expectations more than other respondents, even after conditioning on their larger perception gaps. The heterogeneous updating by demographics and information type are generally consistent with a Bayesian learning model of expectations formation (opposed to an adaptive learning model, or a rational expectations framework).

One policy implication of our results is almost immediate: Consumers respond to information about past prices as well as forecasts of experts by updating their inflation expectations, and so information campaigns might be effectively deployed to affect consumer inflation expectations. Since (1) keeping consumers' inflation expectations anchored is generally important for controlling inflation (Bernanke, 2004), and (2) consumers' inflation expectations may affect their economic decisions (Armantier et al., 2011), the large perception gaps in our sample suggest a role for public information campaigns about past and current inflation as part of prudent monetary policy. In particular, our results suggest that the (high-expectation) right tail of the distribution of public inflation expectations, consisting disproportionately of expectations from female, lower-education, lower-income, and older consumers, could be influenced and managed with public information campaigns, assuming we can find an effective way to deliver the information.

Our findings also underscore the results of Bruine de Bruin et al. (2012), who conclude that our PP question, like the similar Michigan Survey's question about "prices in general", causes respondents to focus more on price changes in their own consumption basket and hence to report expectations that are higher, more dispersed, and more correlated with gas and food price changes. The authors make the case for the RI question being a more reliable survey question, in that it is less sensitive to these transitory price changes. By finding direct evidence that information about food prices affects respondents' PP responses more than their RI responses, we find some support for their conclusion. However, as to which measure of inflation (PP or RI) is relevant for decisions taken by consumers, remains an open question.

While we have shown that respondents revise their inflation expectations sensibly in response to the provided information, we are unable to analyze whether the magnitude of their revisions is either an under- or over-reaction to the provided information. For example, without knowing the share of food and beverage expenditure in each respondent's consumption bundle, we cannot evaluate whether respondents should update their inflation expectations more or less than we observe in response to the Food information treatment. Nevertheless it is unsurprising that the SPF Forecast treatment results in higher-magnitude updating than the Food treatment does. Whereas the Food treatment provides information about past price changes in only a part of a typical consumption basket, the SPF Forecast treatment provides information in the *same time frame* for which we elicit respondents' expectations (next year), and provides information about price changes in a *whole* consumption basket. Furthermore, how respondents' revisions compare to some benchmark, say, if they were Bayesian updaters, is an important question – both for policy-makers and for understanding the heterogeneity in expectations – but one that requires richer data on respondents' consumption bundles as well as on the underlying distributions of the information provided to the respondents.

It should be pointed out that, in our study, respondents do not choose the type of information but are exogenously provided with a treatment. Observed heterogeneity in inflation expectations may partly arise because of demographic differences in information acquisition (Burke and Manz, 2010; Mobius et al., 2011). Moreover, belief-updating when presented with new information in a survey/experiment may be very different from instances where individuals acquire the information themselves (Hertwig et al., 2004). Finally, the long-term effects of new information on respondents' expectations are also unclear. Each of these areas requires further research.[44]

Also, providing information to respondents does not necessarily guarantee more accurate expectations. Whereas we do find in our experimental setting that revised expectations converge toward the range of recent years' inflation realizations (and indeed the actual realized CPI inflation between January 2011 and January 2012), information can have different effects in other contexts: sometimes, individuals presented with new information that is inconsistent with a prior belief may be *less* likely to revise their beliefs, and may even develop more polarized beliefs (Lord, Ross, and Lepper, 1979; Gentzkow and Shapiro, 2006).[45] Therefore, any public information campaigns to help anchor consumer inflation expectations need to be carefully designed indeed.

References

Anderson, R. (2008): "US Consumer Inflation Expectations: Evidence Regarding Learning, Accuracy and Demographics," Centre for Growth and Business Cycle Research Discussion Paper 099.

[44] An extension of the novel methodology presented here would be to re-survey respondents over regular intervals separated by, say, a few weeks. Changes in macroeconomic conditions may allow us to observe how inflation expectations change, especially if the surveys were randomly conducted before versus after substantial inflation-related current events, such as FOMC statements by the Federal Reserve or OPEC meetings. These surveys could occasionally incorporate experimental information treatments, generating an experimental panel of beliefs. This design would be helpful in distinguishing between short-term and long-term effects of information treatments such as the ones we use.

[45] Also, evidence from Germany suggests that while intensity of the coverage of inflation in the media increases inflation expectations' accuracy, they may become more biased if the media content is not neutral (Lamla and Lein, 2010).

Ang, A., G. Bekaert, and M. Wei. (2007); “Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?” *Journal of Monetary Economics*, 54 (4): 1163–1212.

Armantier O., Bruine de Bruin W., van der Klaauw W., Topa G. and B. Zafar (2011): “Inflation Expectations and Behavior: Do Survey Respondents Act on their Beliefs?” mimeo, Federal Reserve Bank of New York.

Armantier, O., and N. Treich (2011): "Eliciting Beliefs: Proper Scoring Rules, Incentives, Stakes and Hedging," Working Paper.

Barber, B., and T. Odean (2001): “Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment,” *The Quarterly Journal of Economics*, 116 (1): 261.292.

Ball, Laurence, Greg Mankiw, and Ricardo Reis (2005): "Monetary Policy for Inattentive Economies," *Journal of Monetary Economics*, 52, 703-725.

Barsky, Robert, and Lutz Kilian (2002): "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," *NBER Macroeconomics Annual 2001*, 137-183.

Bernanke B. (2004): “The Economic Outlook and Monetary Policy,*”* speech at the Bond Market Association Annual Meeting, New York, New York.

Bernanke B. (2007): “Inflation Expectations and Inflation Forecasting,” speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts.

Blanchflower, D., and C. Coille (2009): "The Formation of Inflation Expectations: An Empirical Analysis for the U.K," Paper presented at the Banco do Brasil X1 Annual Inflation Targeting Seminar, May 14--15, Rio de Janeiro.

Branch, W. (2007): “Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations,” *Journal of Economic Dynamics and Control*, 31(1), 245-276.

Brown, F., L. Hoffman, and J. Baxter (1975): “New Way to Measure Price Elasticity,” *Electrical World*, 184, 52–54.

Bruine de Bruin W., W. van der Klaauw, J. Downs, B. Fischhoff, G. Topa, and O. Armantier (2010a): “Expectations of Inflation: The Role of Financial Literacy and Demographic Variables,” *Journal of Consumer Affairs*, 44: 381-402.

Bruine de Bruin, W., S. Potter, R. Rich, G. Topa, and W. van der Klaauw (2010b): “Improving Survey Measures of Household Inflation Expectations,” *Current Issues in Economics and Finance*, 16(7), Aug/Sept.

Bruine de Bruin W., C. Manski, G. Topa, and W. van der Klaauw (2011a): “Measuring Consumer Uncertainty about Future Inflation,” *Journal of Applied Econometrics*, 26(3): 454-478.

Bruine de Bruin W., W. van der Klaauw, and G. Topa (2011b): “Expectations of Inflation: The Biasing Effect of Thoughts about Specific Prices,” Federal Reserve Bank of New York Working Paper, no. 489.

Bruine de Bruin W., W. van der Klaauw, J. Downs, B. Fischhoff, G. Topa, and O. Armantier (2012): “The Effect of Question Wording on Reported Expectations and Perceptions of Inflation,” Federal Reserve Bank of New York Working Paper, no. 443.

Bryan, M., and G. Venkatu (2001a): "The Demographics of Inflation Opinion Surveys," *Federal Reserve Bank of Cleveland Economic Commentary*, October: 1-4.

Bryan, M., and G. Venkatu (2001b): "The Curiously Different Inflation Perspectives of Men and Women," *Federal Reserve Bank of Cleveland Economic Commentary*, November: 1-4.

Burke, M., and M. Manz (2010): "Economic literacy and inflation expectations: evidence from an economic experiment," Working Paper.

Carroll, C. (2003): "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 118 (1): 269-298.

Carter, D. and J. Milon, (2005): "Price Knowledge in Household Demand for Utility Services," *Land Economics*, 81 (2), 265–283.

Chetty, R., and E. Saez (2009): "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipient," Working paper, Harvard University.

Eil, D., and J. Rao (2011): "The Good News-Bad News E¤ect: Asymmetric Processing of Objective Information About Yourself, " *American Economic Journal: Microeconomics*, 3(2): 114-138.

Engelberg J., C. Manski and J. Williams (2009): "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters," *Journal of Business and Economic Statistics*, 27: 30-41.

Evans, G., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

Gali, J. (2008): *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

Garner, A. (1982): Experimental evidence on the rationality of intuitive forecasts. In: Smith, V. L. (Ed.), *Research in Experimental Economics*. Vol. 2. J.A.I. Press, Greenwich, 113–128.

Gentzkow, M., and J. Shapiro (2006): "Media Bias and Reputation." *Journal of Political Economy*, 114(2): 280-316.

Georganas, S., P. Healy, and N. Li (2011): "Frequency Bias in Consumers' Perceptions of Inflation," Working Paper.

Hawkins, D. (1970): "The Effects of Subliminal Stimulation on Drive Level and Brand Preference," *Journal of Marketing Research*, 7(3): 322-326.

Hertwig, R., G. Barron, E. Weber, and I. Erev (2004): "Decisions from Experience and the Weighting of Rare Events," *Psychological Science*, 15 (8): 534-539.

Hey, J. (1994): "Expectations formation: Rational or adaptive or...?" *Journal of Economic Behavior and Organization*, 25(3): 329–350.

Hobijn, B., G. Topa, K. Mayer, and C. Stennis (2009): "Whose Inflation Is It? Household Level vs. Aggregate Measures of Inflation," Manuscript, Federal Reserve Bank of New York.

Jensen, R. (2010): "The (Perceived) Returns to Education and the Demand for Schooling," *The Quarterly Journal of Economics*, 125(2): 515-548.

Jonung, L. (1981): "Perceived and Expected Rates of Inflation in Sweden," *The American Economic Review*, 71 (5): 961--968.

Karni, E., and Z. Safra (1995): "The Impossibility of Experimental Elicitation of Subjective Probabilities," *Theory and Decision*, 38: 313-320.

Keane, M. and D. Runkle (1990): "Testing the Rationality of Price Forecasts: New Evidence from Panel Data," *The American Economic Review*, 80, 714-735.

Kokoski, M. (2000): "Alternative CPI Aggregations: Two Approaches," *Monthly Labor Review*, 123 (July): 31.39.

Lamla, M., and S. Lein. (2010): “The Role of Media for Consumers’ Inflation Expectation Formation,” Working Paper.

Lanne, M., A. Luoma, and J. Luoto (2009): “A Naïve Sticky Information Model of Households’ Inflation Expectations,” *Journal of Economic Dynamics and Control*, 33(6): 1332-1344.

Lipkus I., Samsa G. and B. Rimer (2001): “General Performance on a Numeracy Scale Among Highly Educated Samples,” *Medical Decision Making*, 21, 37–4.

Lombardelli, C., and J. Saleheen (2003): "Public Expectations of UK Inflation", *Bank of England Quarterly Bulletin*, 2003 (Autumn): 281-290.

Lord, C., L. Ross, and M. Lepper (1979): “Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence,” *Journal of Personality and Social Psychology*, 37 (11): 2098-2109.

Lusardi A. (2007): “Financial Literacy: An Essential Tool for Informed Consumer Choice?” presented at the conference “Understanding Consumer Credit: A National Symposium on Expanding Access, Informing Choices, and Protecting Consumers,” organized by the Harvard Joint Center for Housing Studies, Harvard University.

Madeira, C., and B. Zafar (2011): “Heterogeneous Inflation Expectations, Learning, and Market Outcomes,” Working Paper.

Malmendier, U., and S. Nagel (2009): “Learning from Inflation Experiences,” Working paper, Stanford University.

Mankiw, G., and R. Reis (2002): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *The Quarterly Journal of Economics*, 117(4), 1295-1328.

Mankiw, G., R. Reis, and J. Wolfers (2003): “Disagreement About Inflation Expectations,” In *NBER Macroeconomics Annual* 2003, ed. by M. Gertler, and K. Rogoff.

McGranahan, L., and A. Paulson (2006): “Constructing the Chicago Fed Income Based Economic Index – Consumer Price Index: Inflation Experiences by Demographic Group: 1983-2005,” Federal Reserve Bank of Chicago Working Paper.

Muth, J. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29(3): 315-335.

Nakov, A, and A. Pescatori (2010): "Monetary Policy Trade-Offs with a Dominant Oil Producer," *Journal of Monetary Economics*, 42(1), 1-32.

Niederle, M., and L. Vesterlund (2007) "Do Women Shy away from Competition? Do Men Compete too Much?" *The Quarterly Journal of Economics*, 122(3): 1067-1101.

Orphanides, A., and J. Williams (2006): "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations," Journal of Economic Dynamics and Control, 29, 1927-1950.

Pfajfar, D., and E. Santoro (2010): "Asymmetries in Inflation Expectations across Sociodemographic Groups," Working Paper.

Sargent, T. (1993): *Bounded Rationality in Macroeconomics*. Clarendon Press, Oxford, UK.

Sims C. (2009): "Inflation Expectations, Uncertainty and Monetary Policy," BIS Working Paper No. 275.

Souleles, N. (2004): "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 36 (1), 39-72.

Thomas, L. (1999): "Survey Measures of Expected U.S. Inflation," *Journal of Economic Perspectives*, 13 (4): 125–144.

Tversky, A. and Kahneman, D. (1974): "Judgment under uncertainty: Heuristics and biases," *Science*, 185: 1124-1131.

Winkler, R., and A. Murphy (1970): "Nonlinear Utility and the Probability Score," *Journal of Applied Meteorology*, 9: 143-148.

Zwane, A., J. Zinman, E.Van Dusen, W. Pariente, C. Null, E. Miguel, M. Kremer, D. Karlan, R. Hornbeck, X. Giné, E. Duflo, F. Devoto, B. Crepon and A. Banerjee (2011): "Being surveyed can change later behavior and related parameter estimates," *Proceedings of the National Academy of Sciences*, 10(1073): 1-6.

Wiswall, M., and B. Zafar (2011): "Determinants of College Major Choice: Identification Using an Information Experiment," Staff Report No. 500. Federal Reserve Bank of New York.

Appendix A: Financial Literacy Questions

On the following screens, you will receive questions that ask about financial topics. For each question, you must first decide if the statement is true or false and then choose a number to show how confident you are of your answer.

1) If the money on your savings account grows at an annual rate of 5%, then, regardless of inflation, you will be able to buy more with the money in this account in the future than you are able to buy today.

True
False

2) If your income doubles in the next ten years and prices of all goods and services also double, then you will be able to buy fewer goods in ten years than you can buy today.

True
False

Next we would like to ask you some questions which assess how people use numbers in everyday life. Please answer the following questions by filling in the blank. Please do not use a calculator for any of these questions.

If the chance of getting a disease is 10%, how many people would be expected to get the disease:

3) Out of 100 people _____________

4) Out of 1000 people _____________

5) Imagine that we roll a fair, six-sided die 1,000 times. Out of 1,000 rolls, how many times do you think the die would come up as an even number?

If you have $100 in a savings account, the interest rate is 10% per year and you never withdraw money or interest payments, how much will you have in the account after:

6) 1 year ____________

7) 2 years ____________

Figure 1: Survey Design

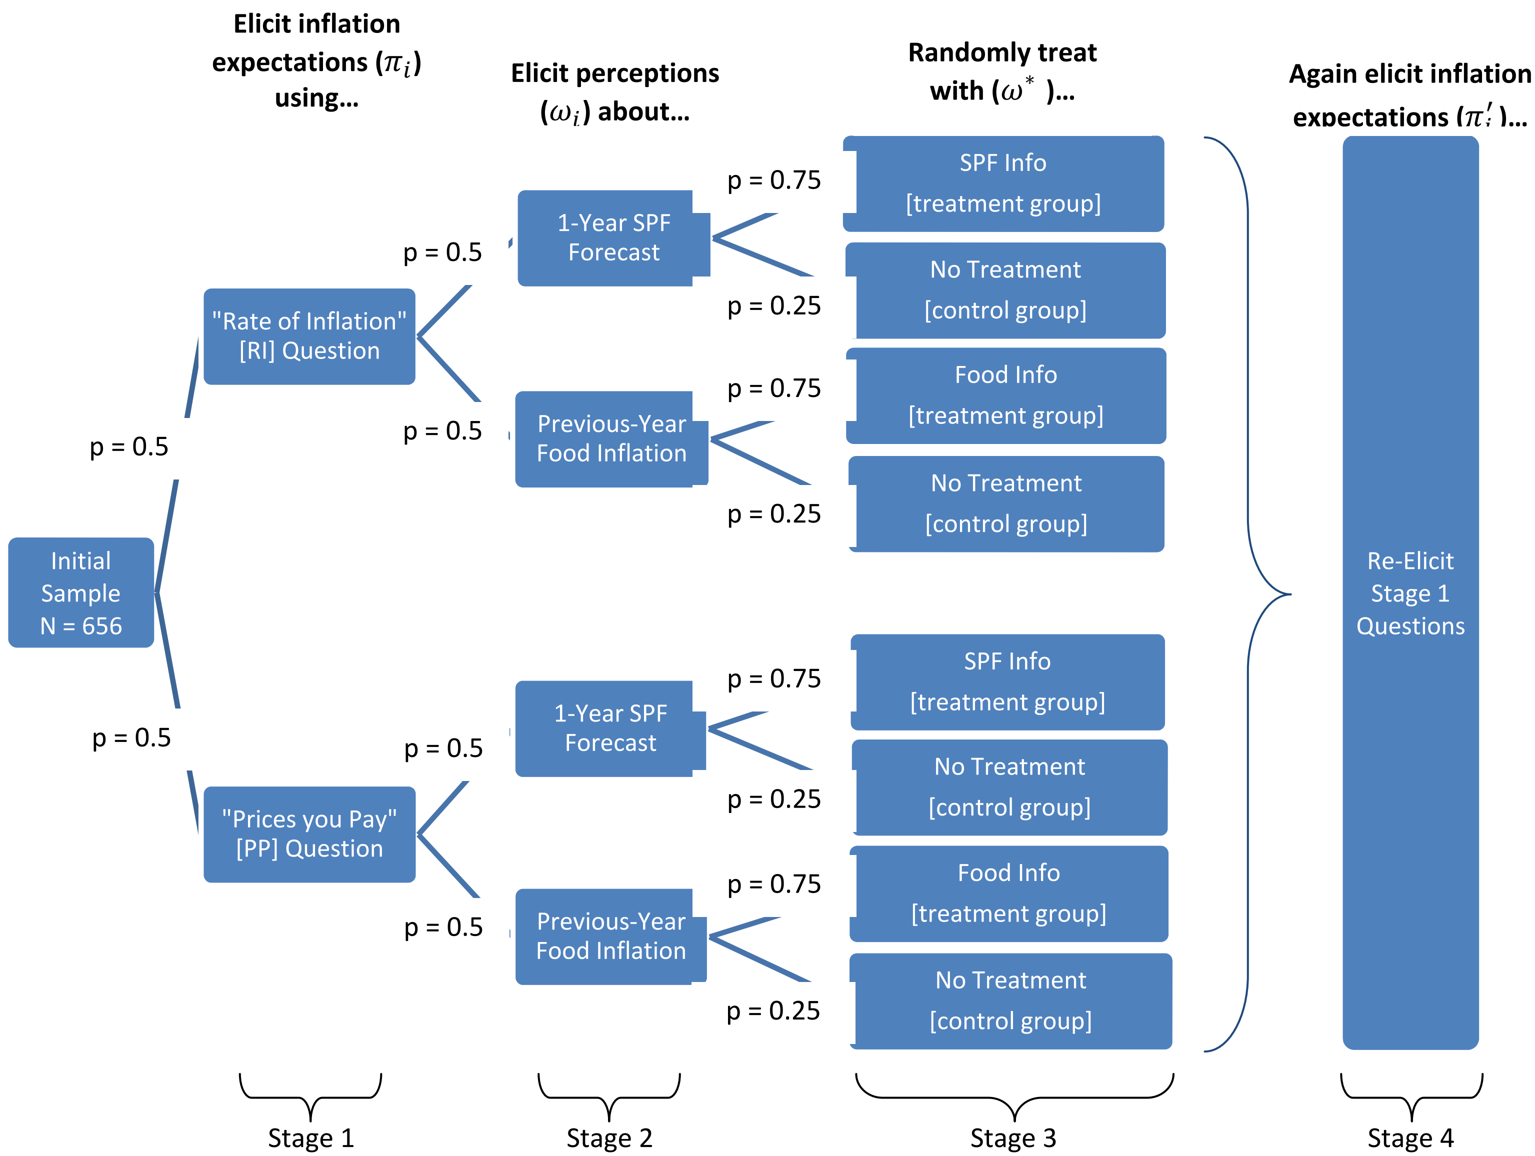

Figure 2: Inflation Expectations Revisions and Perception Gaps, for Information and Control Groups