

FREQUENCY BIAS IN CONSUMERS' PERCEPTIONS OF INFLATION: AN EXPERIMENTAL STUDY[†]

SOTIRIS GEORGANAS*, PAUL J. HEALY**, AND NAN LI***

ABSTRACT. We investigate whether the perception of economy-wide inflation is affected by the frequency with which various goods' prices are observed. We provide novel experimental evidence that consumers' perceptions of aggregate inflation are systematically biased toward the perceived inflation rates of the frequently-purchased items. This 'frequency bias' may affect consumers' consumption and investment decisions, and thus have important macroeconomic consequences. It may also explain why consumers typically over-estimate inflation in surveys during periods where frequently-purchased non-durable goods are inflating faster than durables.

Key words: Inflation perceptions; biases; laboratory experiment

JEL classification codes: C91; E31; C82.

[†]The authors thank John Duffy, Richard Lowery, James Peck, Julia Thomas and Huanxing Yang for their helpful comments.

*Dept. of Economics, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, UK; sotiris.georganas@rhul.ac.uk.

**Corresponding author. Dept. of Economics, The Ohio State University, 1945 North High street, Columbus, OH 43210, U.S.A.; healy.52@osu.edu. Fax: (614) 292-3906.

*** International Monetary Fund, 700 19th Street NW, Washington DC, 20431; nli@imf.org.

I INTRODUCTION

Perceptions of inflation play a central role in macroeconomic analyses because they directly affect the decisions of individuals, organizations and states. The efficacy of monetary policy is thought to depend more heavily on the perceived inflation rate than the actual, calculated inflation rate (e.g. Bernanke, 2007 and Blanchard et al., 2010). There is evidence, however, that the perceived rate is not necessarily consistent with the actual rate, and sometimes can deviate by a large margin.¹ Thus, it is important to understand how consumers perceive inflation, and uncover any systematic biases in those perceptions.

In this paper, we conduct an experimental and behavioral analysis of misperceptions and biases in perceived inflation. We use laboratory experiments to identify a specific, yet important bias in the perception of inflation: Consumers tend to overweigh price changes of frequently-purchased goods when forming economy-wide inflation perceptions. We refer to this phenomenon as the *frequency bias* in inflation perceptions. We also uncover a *correlation bias* in perceived individual-good rates, wherein these rates are perceived to be more similar than their true values.

The frequency bias generates testable predictions for field data. Items such as ‘food at home’ and ‘gasoline’ are frequently purchased, and are the most visible and best publicized components of the Consumer Price Index (CPI). But they constitute only 8.9% and 3.7% of total aggregate expenditures, respectively, and therefore only a small fraction of the aggregate inflation rate.² Frequency-biased consumers will perceive overall inflation rates that are biased toward the inflation rates of gasoline and food. Survey evidence on inflation perceptions broadly supports this conclusion. For example, Jonung (1981) finds that inflation perceptions of women are more biased than those of men, and attributes this to their more frequent retail shopping experiences and high food inflation rates at the time.³ More recently, Del Giovane and Sabbatini (2008) document that perceived inflation tends to be higher during periods in which prices of frequently purchased goods (e.g. basic food products like milk and vegetables) experienced larger increases. Jonung and Laidler (1988) show that perceptions are in general not rational,

¹For example, it is well documented that perceived inflation dramatically increased during the Euro changeover in 2002, even though the actual rate was stable (Fluch and Stix, 2005).

²These expenditure shares are based on the average weight over 1980-2010 from the U.S. Bureau of Labor Statistics.

³In a recent survey of Ohio consumers, Bryan and Venkatu (2001a) also find that women perceive higher rates of recent-past inflation than men, even controlling for age, socioeconomic and educational differences.

average perceptions of inflation are too high, and individual perceptions are strongly correlated with demographic characteristics and socioeconomic status.⁴ In fact, macroeconomic policy makers often stress the importance of food prices “in determining the wage demands of labor and the inflationary expectations of all consumers” (p.34, 1976 *Economic Report of the President*), suggesting some awareness of this bias. And it is the overall inflation rate that individuals must use when making financial investment and savings decisions, so a bias in perceptions of this rate should have real economic consequences through distorted allocations of financial assets.

Without the control of all relevant variables that the laboratory affords, however, it is difficult to provide anything more than suggestive evidence about such biases. For example, Ranyard et al. (2008) point out that an apparent bias toward essential goods like food and gasoline may arise because official CPI calculations use aggregate expenditure shares, which are more heavily influenced by the larger purchases of high-income individuals. The median consumer’s basket of goods contains more of these essential items than the CPI basket, so the median perception will appear to have a frequency bias. Our laboratory experiment rules out this explanation by forcing all subjects to consume the same basket. Field observations are further plagued by unobservable information sets, changing quality of goods, and lack of incentives in surveys (Mishkin, 1981).⁵ In contrast, all of these can be controlled and manipulated in the laboratory, allowing for cleaner measurement of perceptions.

Our experiment centers around a simulated shopping experience, where subjects have to choose to buy a variety of goods over time, each with its own rate of inflation. After shopping, subjects are asked their perceptions of the ‘economy-wide’ inflation rate during the experiment, as well as the rates of each individual good. In calculating the economy-wide rate, a well calibrated individual should take a weighted average of individual-good rates, weighing each by the fraction of their total expenditures spent on that good. Actual economy-wide inflation reports are clearly biased toward their reported rates of the more frequently purchased goods. Their individual-good reports are also more similar than the true rates. In a follow-up experiment, subjects who are told

⁴For example, perceptions tend to be higher for young people, women, unmarried individuals, minorities, and lower-income individuals. Similar patterns have also been documented in the U.S. (Bryan and Venkatu, 2001b), England (Blanchflower and MacCoille, 2009), Ireland (Duffy and Lunn, 2009) and New Zealand (Leung, 2009).

⁵Recent attempts have been made to correct for the upward bias introduced by substitution effect and quality improvement in the official CPI. Comparing the perceived inflation with chain-weighted hedonic CPI inflation would result in larger gap, as the improved CPI inflation is typically lower than the unadjusted CPI inflation.

the true inflation rates of the individual goods still exhibit the frequency bias, indicating that the frequency bias does not stem from misperceptions of individual-good rates, but occurs only when aggregating those rates.

In our experiment, information sets are both observable and manipulable by the researcher. Subjects' beliefs are elicited in an incentive-compatible way. The basket of purchased goods is chosen in advance and perfectly observed. The purchased goods are fictitious and therefore not subject to quality variations. This level of control allows us to eliminate confounds and isolate biases. The limitation of the laboratory is in the generalizability of results, given the level of abstraction and the self-selected subject pool. But our findings are consistent with previous empirical work, and with our own analysis of survey data on inflation perceptions (Section V). This suggests that the bias observed in the laboratory may well be present in real-world perceptions in the field.

We are not the first to propose that inflation perceptions may exhibit the frequency bias. For example, excessively high perceptions of inflation after the Euro cash changeover have been attributed to the fact that frequently-purchased goods happened to be inflating faster during that time period (Del Giovane and Sabbatini, 2008, e.g.). Bates and Gabor (1986) suggest that perceptions of inflation are altered by the availability of price change information in a person's memory.⁶ Del Missier et al. (2008) show that inflation perceptions can be increased by priming subjects to think first about highly-inflating goods. Ranyard et al. (2008) argue that availability should be greater for more frequently-purchased goods. Thus, availability provides an established psychological foundation for the frequency bias.

Inspired by these arguments, and by the loss aversion aspect of prospect theory (Kahneman and Tversky, 1979), Brachinger (2008) constructs an Index of Perceived Inflation (IPI) that uses frequency-based weights on individual goods (instead of expenditure-based weights), and an asymmetric perception of price increases versus decreases. Jungermann et al. (2007) provide empirical support for these assumptions using a novel survey on inflation perceptions and frequencies of purchase.⁷

A number of other factors have been identified that influence perceptions of inflation. These include socio-economic status, education level, personal income levels and income

⁶Tversky and Kahneman (1973) first formalized availability as influencing perceptions of lotteries, but the extension to price changes is transparent.

⁷Jungermann et al. (2007) find that absolute levels of perceived inflation correlate with purchase frequencies (which are self-reported on a three-point scale), but they do not correlate *errors* in perceived rates with the rates of frequently-purchased goods. Their findings would not represent a bias if frequently-purchased goods in fact had higher inflation during the survey period.

growth rates, personal experience with prices, social amplification by media, past expectations of current-period inflation, nationality, and the time horizon in question. A separate strand of research focuses on the formation of reference prices, where new prices are evaluated against past reference points. See Ranyard et al. (2008) and Mazumdar et al. (2005) for surveys of these respective literatures.

To our knowledge, the correlation bias in individual-good perceptions has not been documented, because almost all studies of inflation focus only on the economy-wide rate. But in other domains it is well documented that subjects employ a similarity heuristic, assessing independent objects as more similar than warranted. It may also be a reasonable heuristic in this setting because, in practice, inflation rates of individual goods often do exhibit correlation.

The closest work to ours is that of Huber (2011), who shows subjects cards with prices of various goods and asks them to guess the overall inflation rate with respect to previously-shown reference prices. As in our experiment, subjects report higher inflation when frequently-seen goods are inflating faster, documenting a frequency bias in that task. Huber also observed surprisingly accurate inflation reports for individual goods and, therefore, no correlation bias. This is likely because Huber's subjects were trained in performing inflation calculations at the start of the experiment and only had to compare observed prices against a known reference price. Our subjects, in contrast, faced a long stream of adjusting prices and were not trained in inflation calculations.⁸ Thus, we view Huber's result as similar to our finding that the frequency bias persists when individual-good inflation rates are known, but is less informative about settings where biases in individual-good rates are also present.

A handful of experimental studies investigate the formation of expectations of future prices (e.g. Schmalensee, 1976; Garner, 1982; Camerer, 1992; Hey, 1994). In these contexts, the rational expectations hypothesis is generally rejected in favor of adaptive expectations models. Adam (2007) studies forecasts of future inflation rates in a simulated macroeconomic model and finds that subjects adopt a 'Restricted Perceptions Equilibrium' in which agents use simple forecast functions only and outcomes and beliefs reinforce each other. Pfajfar and Zakelj (2009) find that, within a New Keynesian sticky price framework, subjects use various models of expectations formation of inflation, including sticky information, adaptive learning and rational rules. Burke and Manz (2011) study how economic literacy affects inflation expectations formation

⁸Huber also did not study price recall, and did not elicit reports in an incentive-compatible way.

through two specific channels: the choice of information and the use of given information. Malmendier and Nagel (2012) propose a personal experience-based model of expectations formation. Using cross-sectional survey data, they find differences in expectations across age groups can be explained by the variations in their life-time inflation experiences.

The rest of the paper is organized as follows. We formally define the frequency bias in Section II. In Section III we describe the experimental design and present the various results from the laboratory. In Section V, we relate our findings in the lab to survey observations. Section VI concludes.

II THE FREQUENCY BIAS DEFINED

Economy-wide inflation rates are generally calculated as the rate of change in the total price of a representative basket of goods. Formally, if each good i 's price at some point in time t is given by p_{it} , and if a basket is comprised of q_i units of each good i , then the period- t price of the basket is given by

$$(1) \quad P_t = \sum_i q_i p_{it}.$$

The inflation rate from period $t-1$ to t for each good is given by $\pi_{it} = (p_{it} - p_{i,t-1})/p_{i,t-1}$, and the economy-wide inflation rate is $\Pi_t = (P_t - P_{t-1})/P_{t-1}$. A bit of algebra shows that the aggregate inflation rate must be a convex combination of individual-good inflation rates, with the weight on each good equal to its share of the total expenditure in period $t-1$. Formally, we have

$$(2) \quad \Pi_t = \sum_i \theta_{i,t-1} \pi_{it},$$

where

$$(3) \quad \theta_{i,t-1} = \frac{q_i p_{i,t-1}}{\sum_j q_j p_{j,t-1}}.$$

We refer to θ_{it} as the *expenditure weight* for good i at time t . For our experiments θ_{it} does not vary with time, so we typically ignore the t subscript in the notation.

In reality, consumers may have perceptions of inflation for each good i , denoted π_i^P , that differ from the true good- i inflation rate π_i . Their perception of the economy-wide inflation rate (Π^P) may also differ from the true economy-wide rate (Π). We can relate

these perceived rates by

$$(4) \quad \Pi^P = \sum_i \omega_i \pi_i^P,$$

where ω_i is the weight the consumer actually places on π_i^P . Regardless of the accuracy of each π_i^P , if a consumer understands that economy-wide inflation rates are calculated by constructing a basket of q_i units of each good i , then it must be that $\omega_i = \theta_i$ for each i .⁹

A frequency bias occurs when the consumer's actual weights ω_i deviate from θ_i , with more weight put on goods that are more frequently purchased, and less weight put on goods that are less frequently purchased. To separate frequency of purchase from quantity purchased, we let $q_i = n_i \mu_i$, where n_i is the number of times good i was purchased in the given time period (measured as the number of distinct transactions), and μ_i is the average quantity per purchase. The *frequency weight* of good i is given by

$$(5) \quad \phi_i = \frac{n_i}{\sum_j n_j}$$

From these weights we can formally define the frequency bias:

Definition. A consumer's perceptions of inflation exhibit the **frequency bias** if there is some $\alpha > 0$ such that, for each good i ,

$$(6) \quad \omega_i = \alpha \phi_i + (1 - \alpha) \theta_i,$$

where $\phi_i = n_i / \sum_j n_j$ is the relative frequency with which good i is purchased and $\theta_i = P_i / P$ is the fraction of total expenditures spent on good i .

The degree to which consumers use frequency weights versus expenditure weights is captured by the parameter α . An unbiased consumer has $\alpha = 0$. Given α , the perception of the overall inflation rate is calculated as

$$\begin{aligned} \Pi^P &= \sum_i \omega_i \pi_i^P \\ &= \sum_i [\alpha \phi_i + (1 - \alpha) \theta_i] \pi_i^P. \end{aligned}$$

⁹If the consumer knows that a basket is used to calculate inflation rates, but they do not know the quantities q_i , then it still must be true that $\omega_i \in [0, 1]$ for each i and $\sum_i \omega_i = 1$. In our experiment the quantities q_i are clearly shown.

Letting $\Pi_{\text{EXP}}^p = \sum_i \theta_i \pi_i^p$ be the correct expenditure-weighted inflation rate and $\Pi_{\text{FREQ}}^p = \sum_i \phi_i \pi_i^p$ be the frequency-based inflation rate, we have that

$$(7) \quad \Pi^p = \alpha \Pi_{\text{FREQ}}^p + (1 - \alpha) \Pi_{\text{EXP}}^p.$$

Thus, the frequency bias can equivalently be expressed as a bias in Π^p toward Π_{FREQ}^p .¹⁰

Again, the parameter α provides a simple way to measure the magnitude of the bias. It is this parameter that we measure in our controlled laboratory experiments.

The frequency bias represents an error in how individuals aggregate the inflation rates within their own consumption basket. The bias will be observed at an aggregate level (aggregating across individuals) as long as most individuals have similar relative frequencies of purchasing various goods.

III THE LABORATORY EXPERIMENT

Experimental Design

The experiment is designed to measure perceived inflation rates in a simulated economy. The frequency bias can be estimated by comparing reported perceptions of economy-wide inflation against the reported perceptions of inflation for each individual good. If the frequency bias is strong, then economy-wide inflation reports will be biased toward the inflation rates of goods that are purchased more frequently; therefore, we compare a baseline treatment with fairly flat inflation (Treatment EQ) to two treatments where the most-frequently purchased goods inflation rates are either large and positive (Treatment POS) or large and negative (Treatment NEG).

Nine experimental sessions were conducted at Ohio State University in November and December of 2009.¹¹ All subjects were Ohio State undergraduate students recruited via e-mail.¹² All sessions took place in the Ohio State Experimental Economics Laboratory. In total, 186 subjects participated in the experiment in sessions of roughly 21 subjects per session. Each subject was only allowed to participate in one session of this experiment.

¹⁰The Index of Perceived Inflation (IPI) developed by Brachinger (2008) is similar to Π_{FREQ}^p , except all individual-good rates are multiplied by $c = 2$ if they are positive to account for loss aversion.

¹¹One small pilot session and one session with a technical flaw are excluded. Results from these sessions appear qualitatively similar to the nine reported sessions. The data are available upon request.

¹²Existing evidence shows that undergraduate students behave similarly to other populations in most economic decisions; there are some settings with systematic subject-pool differences, though there seems to be little guidance about which settings will generate differences and in which directions these differences will operate. Thus, from an *ex-ante* perspective, subject pool effects can be treated as unbiased noise. See Fr chet te (2009) for details.

For each session, all subjects arrived at the laboratory simultaneously, were seated at computer terminals, and told to log into the experiment website.¹³ The website then provided specific instructions regarding the procedures for the experiment, which subjects read at their own pace. They then proceeded to make a series of decisions through the experiment website. Once every subject had completed the experiment, each was paid in cash privately based on their earnings and left the laboratory. Earnings during the experiment were recorded in ‘points’, with each point being worth one penny of actual payout. Final earnings ranged from \$8.40 to \$22.59, with a mean of \$18.15. Sessions took roughly one hour to complete.

The experiment consists of two phases. The first phase is broken into 96 periods, referred to as ‘days’. Sixteen days constitute a ‘month’, for a total of six months in the first phase. In each day subjects are shown a 4×3 table of prices. An example table is shown in Figure I. Each row corresponds to a different type of good, labeled abstractly as goods A, B, C, and D, and each column corresponds to a different brand, labeled as 1, 2, and 3. Each day subjects are told which type of good they are to purchase (A, B, C, or D) and are asked to select the cheapest price for that good. They could then click on any of the twelve prices in the table. If they click on the lowest price of the correct good then they earn five points. The middle price of the correct good earns them three points, and the highest price earns them one point. Clicking on any price of an incorrect good earns them zero points. After clicking a price, the experiment proceeds to the next day, where a new table of twelve prices is shown and subjects are again told which good to buy. If a subject does not click any price within 30 seconds then they earn zero points for that ‘day’ and the experiment automatically proceeds to the next day (no time limit was imposed on the first day).

Over the 96 days, subjects shop for the different goods with different frequencies. Specifically, in each 16-day month they are asked to buy good A seven times, good B six times, good C two times, and good D one time; see Table I. We refer to each month’s bundle of purchases as a ‘basket’. The ordering of the purchases in the basket was randomized within each month.

The simulated shopping experience is designed to mimic key aspects of actual consumer purchases. When shopping for an item, consumers focus only on a single type of good, though other goods’ prices are available for perusal. Some items—such as gasoline and food—are purchased more frequently than others. Multiple prices for the desired

¹³The website is available at <http://healy.econ.ohio-state.edu/exp/shopping>. To experience this experiment and view the instructions, log in using session password ‘test’.

Shopping

Shopping Day 1 of Month 1.
(There are 16 shopping days per month and 6 months total.)

Today you need to buy Good B.
Please select the cheapest price for Good B.

IT IS NOW MONTH 1

Time remaining for this day:
(There is no time limit for your first shopping day.)

GOOD	BRAND 1	BRAND 2	BRAND 3
A	\$1.06	\$1.09	\$0.91
B	\$6.52	\$7.57	\$6.69
C	\$131.70	\$115.46	\$127.33
D	\$487.89	\$460.40	\$487.34

Your Total Points Earned: 0 Points.

FIGURE I. Phase I of the experiment: The shopping decision.

Good	A	B	C	D	Basket
Purchases per month	7	6	2	1	16
Month 1 mean price	\$1	\$7	\$122	\$470	\$763
Monthly inflation rate:					
Treatment EQ	4%	4%	4%	4%	4%
Treatment POS	10%	9%	7%	1%	3.63%
Treatment NEG	-10%	-2%	1.5%	5.5%	3.80%

TABLE I. Frequencies of purchase, starting prices, and inflation rates for the four goods used in the experiment.

good may be offered, adding noise to inflation perceptions, and consumers benefit by choosing the lowest-priced option. No notion of quality is introduced so that prices need not be adjusted for varying quality levels. We reward purchases using a fixed point system rather than giving shoppers a total budget because recalling basket inflation rates with a fixed budget would amount to observing the total change in the budget. This would oversimplify the problem of recalling inflation rates since, in reality, liquid asset balances are affected by much more than expenditures.

Each of the four goods $i \in \{A, B, C, D\}$ is given an initial mean price \bar{p}_{i1} for the first month; the values of \bar{p}_{i1} used in the experiment are given in the third row of Table I. At the beginning of each subsequent month, the mean price for each good jumps by the monthly inflation rate π_i^* , which does not vary during the experiment. It then stays constant through the month. In Treatment EQ (Sessions 1–3) all four goods have an equal inflation rate $\pi_i^* = 0.04$. In Treatment POS (Sessions 4–6), the inflation rates are positively correlated with the frequency of purchase, so that the more frequently-purchased goods have higher inflation rates. In Treatment NEG (Sessions 7–9) the inflation rates are negatively correlated with frequency of purchase, with goods A and B actually experiencing deflation on average.

Although inflation occurs from month to month, the mean price does not change within the month. Thus, for any day t in month m the mean price of good i is \bar{p}_{im} , and in every day of month $m + 1$ the mean price of good i is $\bar{p}_{i,m+1} = \bar{p}_{im} (1 + \pi_i^*)$.

The three daily prices for each good offered to the subject each day are uniform random draws centered at the current month's mean price. Specifically, in each day t of month m the realized price of brand $b \in \{1, 2, 3\}$ is a value p_{ibmt} drawn from a uniform distribution over the interval $[0.9\bar{p}_{im}, 1.1\bar{p}_{im}]$, and then rounded to the nearest penny. Each brand's daily price is drawn independently of all other prices, conditional on that good's mean price for the month. All twelve prices (three brands of four goods) for each day are shown in a single table so subjects can easily see all prices for all goods each day. See Figure I for an example of the actual table presented to subjects in the experiment.

If $\iota(m, t) \in \{1, 2, 3, 4\}$ identifies the good a subject is asked to buy on day t of month m , and if $p_{imt} = \min\{p_{i1mt}, p_{i2mt}, p_{i3mt}\}$ denotes the minimum price for good i on day t of month m , then the total expenditure on good i in month m is given by

$$P_{im} = \sum_{\{t:\iota(m,t)=i\}} p_{imt}.$$

The realized total basket price for month m is then the total expenditure for the month, $P_m = \sum_i P_{im}$.

The realized inflation rate for the entire basket of goods over the six months is given by $\Pi = (P_6 - P_1)/P_1$. The realized inflation rate for each good i over the six months is the change in total expenditures on good i between the first and last month, or $\pi_i = (P_{i6} - P_{i1})/P_{i1}$. Here, the realized inflation rates π_i may differ slightly from the fixed, underlying inflation rates π_i^* given in Table I because of randomness in the actual price draws observed by a subject. Note that all inflation rates are calculated based on subjects' actual purchases, and not on prices they did not pay.

As described in Section II, the basket inflation rate must be a convex combination of individual inflation rates, using the expenditure shares as weights. Thus, if $\theta_i = P_{i1}/P_1$ is the expenditure weight of each good i , then

$$(8) \quad \Pi = \sum_i \theta_i \pi_i.$$

Phase one of the experiment ends after all six months of shopping were complete, which typically takes about twenty minutes. At no point during the first phase are subjects told that they are buying an identical basket of goods each month—though an astute subject could deduce this fact—and subjects are never told in phase one that they will be asked inflation-related questions in phase two.

Phase two consists of two decisions made sequentially: A guess of the basket inflation rate and a guess of each good's inflation rate.

Before the first decision, subjects are told that they had just purchased an identical quantity of each good in each month, thus forming a 'basket' of goods that they had purchased in each month. They are then asked: "What was the TOTAL percentage change of the price of a basket of goods from month 1 to month 6?" Subjects then enter a guess of the six-month basket inflation rate, which we denote here by Π^p .¹⁴ At the end of the experiment they are told the realized inflation rate Π and receive $425 - 500 |\Pi^p - \Pi|$ points for their guess. Thus, a perfect guess earns \$4.25, while a guess that is off by ten percentage points (where $|\Pi^p - \Pi| = 0.10$) earns \$3.75. Earnings are truncated below zero, so no subject can earn negative payoffs for this decision. Subjects do not learn the true inflation rate or their earnings for this guess until the experiment is complete.

After submitting their estimate of the basket inflation rate, subjects are asked to guess the six-month inflation rate for each of the four goods. At the end of the experiment, the subject is paid $125 - 500 |\pi_i^p - \pi_i|$ points for each of their four guesses π_i^p . Thus, four perfect guesses earns \$5.00, and subjects lose five cents for every percentage

¹⁴Before continuing, subjects are asked to verify all decisions that require keyboard input in order to minimize the occurrence of typographic errors.

	Theoretical Weights	Estimated Weights	95% Confidence Interval
Frequency-Based Inflation Rate	0.000	0.440	[0.241,0.639]
Expenditure-Based Inflation Rate	1.000	0.560	[0.361,0.759]

TABLE II. Estimate of the size of the frequency bias in reported inflation rates.

point difference between a guess and that good's true inflation rate. Again, earnings were truncated below zero, so no subject could earn negative payoffs for this decision. Subjects do not learn the true inflation rate or their earnings for this guess until the experiment was completed.

At the end of the experiment subjects are shown their earnings in points from each decision in the experiment, along with the true inflation rates for each good and for the entire basket of goods. The point earnings are then converted to dollars (at a rate of one cent per point) and rounded up to the next whole dollar amount. Subjects are paid their earnings in cash privately, sign a receipt, and leave the laboratory individually.

In our analysis, eight subjects (out of 186) are removed from the data as outliers for having at least one guess whose error was greater than 100 percentage points. Analyzing medians without removing outliers yields qualitatively similar conclusions, but is less amenable to regression analysis.

Measuring the Frequency Bias

We begin by measuring the degree of frequency bias across all treatments. This is done by estimating the parameter α in the relationship

$$\Pi^p = \alpha \Pi_{\text{FREQ}}^p + (1 - \alpha) \Pi_{\text{EXP}}^p$$

that was derived in equation (7) above. The values Π_{EXP}^p and Π_{FREQ}^p are calculated from the individual-good inflation reports for each subject. The announced basket rates (Π^p) are then regressed on these two values. We constrain the two regression coefficients to sum to one, though we do not require that α be between zero and one. The results are shown in Table II.

On average, subjects put 44% weight on the frequency with which goods are purchased and only 56% weight on the (theoretically-correct) expenditure weights. These estimates have a standard error of 0.101 and are therefore significantly different from

the theoretical predictions at the five percent level. Thus, the frequency bias is both statically significant and economically meaningful in size.

This result is robust to the specification of the linear regression. Removing the constraint that the coefficients sum to one gives an estimated relationship of

$$\Pi^p = 0.428\Pi_{\text{FREQ}}^p + 0.528\Pi_{\text{EXP}}^p$$

with a standard errors of 0.102 and 0.109, respectively. Also allowing for a constant gives an estimated relationship of

$$\Pi^p = 8.187 + 0.419\Pi_{\text{FREQ}}^p + 0.257\Pi_{\text{EXP}}^p$$

with standard errors of 1.905, 0.975, and 0.121, respectively. The positive constant indicates a general tendency to report basket rates that are high relative to the reported individual-good rates, and both slope estimates remain significantly different from both zero and one.¹⁵

Breaking the result down by treatment yields somewhat noisier results because the sample sizes are smaller. In Treatment POS the estimated α (the weight on Π_{FREQ}^p) is 0.387, with a p -value of 0.002. In Treatment NEG the estimated α is 0.291 with an insignificant p -value of 0.124. In Treatment EQ α is not identified because $\Pi^{\text{FREQ}} = \Pi^{\text{EXP}}$.

Treatment Differences

Figure II presents the difference between actual and reported inflation for each treatment. It shows that people overestimate basket inflation rates when the frequently-purchased goods have the highest inflation rates (Treatment POS), and that people underestimate overall inflation rates when the frequently-purchased goods have the lowest inflation rates (Treatment NEG). When all goods have the same inflation rate, subjects are reasonably well calibrated (Treatment EQ).

Table III shows the average reported and actual inflation rates in each treatment for each individual good and for the total basket, as well as the average error for each. The treatment effects from Figure II are apparent in the last column of the table; reported rates for the basket are too high in Treatment POS, roughly accurate in the Treatment

¹⁵Because Π_{EXP}^p and Π_{FREQ}^p have an estimated correlation coefficient of 0.723, one might worry that these regressions are impacted by multicollinearity problems. Diagnostic tests show that multicollinearity is not a serious problem here: The variance inflation factor of the last regression is 2.09 and the tolerance of α is 0.478. Both are below most thresholds for concern. Finally, regressing Π^p on Π_{FREQ}^p alone gives an estimated relationship of $\Pi^p = 10.286 + 0.568\Pi_{\text{FREQ}}^p$ with standard errors of 1.644 and 0.068, respectively.

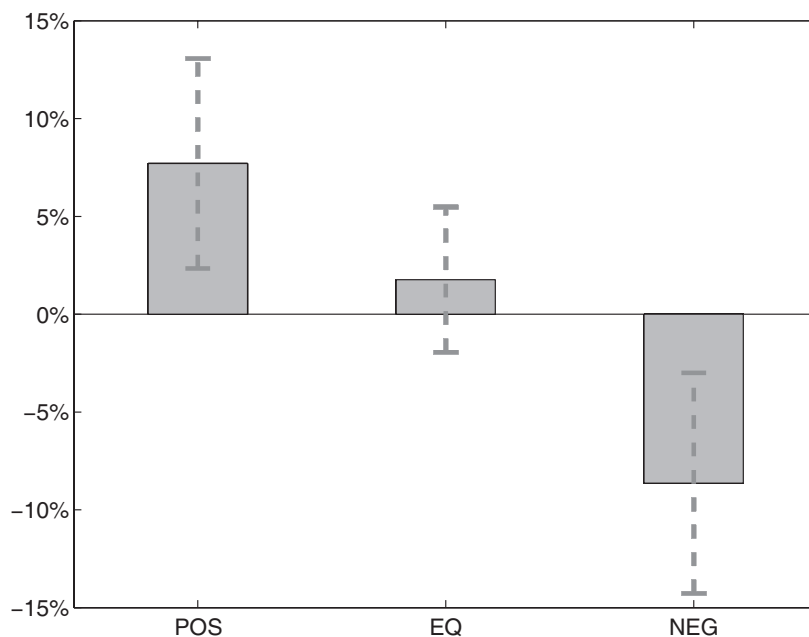


FIGURE II. Average errors (in percentage points) in economy-wide inflation perceptions by treatment. The dashed lines represent 95% confidence intervals.

Trt		Good				Basket
		A	B	C	D	
	Purchases/Period	7	6	2	1	
POS	Reported Inflation	29.30	23.70	22.90	22.49	25.47
	Actual Inflation	61.02	53.24	40.42	5.16	19.57
	Average Error	-31.72*	-29.54*	-17.52*	17.33*	5.90*
EQ	Reported Inflation	19.87	19.19	19.88	19.17	23.01
	Actual Inflation	21.96	21.46	21.08	21.27	21.16
	Average Error	-2.09	-2.27	-1.21	-2.10	1.86
NEG	Reported Inflation	-2.08	7.57	8.62	12.62	10.47
	Actual Inflation	-40.99	-9.75	8.28	30.33	20.37
	Average Error	38.91*	17.32*	0.34	-17.71*	-9.90*

Note: *Average error is significantly different from zero at the 1% level.

TABLE III. Mean six-period inflation rates for each good and for the entire basket.

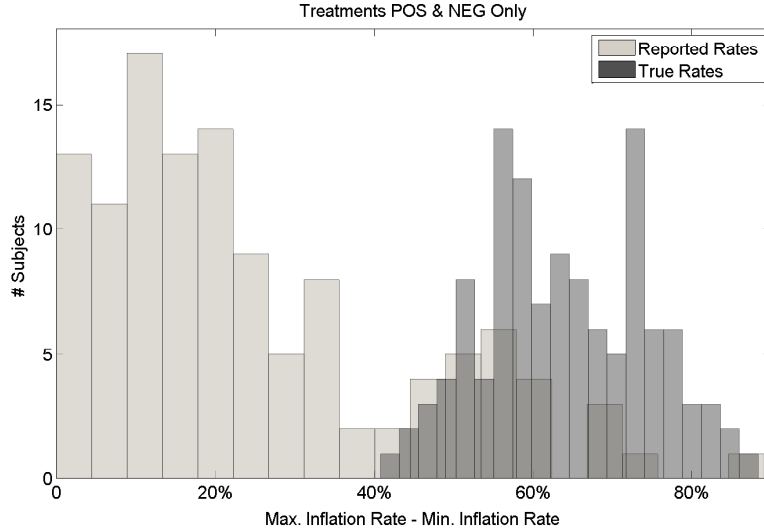


FIGURE III. Histograms of the difference between the maximum and minimum of individual good inflation rates, reported and actual.

EQ, and too low in Treatment NEG. The average errors (reported rates minus true rates) are significantly different from zero (at the 5% level) in Treatments POS and NEG, but not in EQ. Errors in Treatment NEG are significantly lower than in the other two treatments (large-sample permutation test p -values of < 0.001 and 0.001 , respectively; see Siegel and Castellan, 1988, p.154), though errors in Treatments POS and EQ are not significantly different (p -value of 0.157).

Individual-Good Inflation Rates

The frequency bias is one of aggregation. Table III also reveals a systematic bias in the accuracy of individual-good inflation rates: Subjects report individual-good inflation rates that are biased toward the overall basket rate. For example, in Treatment POS, subjects grossly underestimate the rate of the highest-inflation goods and overestimate the rate of the lowest-inflation goods. In Treatment NEG the same phenomenon occurs. In Treatment EQ, however, each individual good’s inflation rate equals the basket rate, and so subjects’ individual-good reports are well calibrated. Figure III shows that, for Treatments POS and NEG, the range of reported individual-good rates is generally smaller than the true range. In fact, 94% of these subjects report a smaller range than the range they actually experienced. Eight of these subjects (6.8%) report the exact same rate for all four goods.

We refer to this bias in the accuracy of individual good reports as the *correlation bias*, since reported individual-good rates are more correlated than the true individual-good rates.¹⁶ In all three treatments, the actual prices are independently drawn for each good, and so no correlation exists between the true inflation rates of the four goods.¹⁷ Subjects' reports, however, are highly correlated. For every treatment and for every pair of goods, the subjects' reported inflation rates for those two goods have a positive correlation coefficient that is significant at the five-percent level. The estimated coefficients are all greater than thirty percent. Figure IV shows the relationship between actual inflation rates and reported rates for the four goods. A linear regression shows a relationship of 0.27 with a standard error 0.025, significantly less than the one-to-one relationship that would be exhibited by a well-calibrated individual.

It is possible that the correlation bias is driving the treatment differences observed in Figure II. Suppose subjects exhibit the correlation bias but not the frequency bias, so that their individual-good rates are correlated but, given those incorrect rates, they form their perceptions of the basket rate using the (correct) expenditure weights on each good. In Treatment POS, the infrequently-purchased Goods C and D have the lower inflation rates, so such a subject would overestimate those rates. But those goods also constitute 93 percent of total expenditures, so these two overestimates would result in an overestimation of the basket rate. Similarly, in Treatment NEG, Goods C and D have high inflation rates and would be underestimated, leading to an underestimated basket rate. In Treatment EQ all goods' rates would be correctly perceived, as would the basket rate. These predictions exactly match Figure II.

To disentangle the correlation bias from the frequency bias, we regress basket inflation reports on individual-good inflation reports to estimate the weights subjects place on each good. These estimates can then be compared to the (average) expenditure-based weights that subjects would use if they exhibit no frequency bias, as well as the frequency-based weights for each good. The results appear in the first panel of Table IV.

Indeed, subjects' actual weights exhibit a significant frequency bias for Goods A, C and D (one-sided p -values of 0.004, 0.035, and 0.032, respectively), with a marginally

¹⁶A simple explanation for the correlation bias is that subjects have a prior over inflation rates that is common across goods and correct when averaging across goods. This is particularly justifiable in a context-free laboratory experiment with fictitious goods. If subjects observe the true inflation rate with noise (perhaps due to inattention) then, assuming the distributions of the prior and noise are symmetric and quasiconcave (see Chambers and Healy, 2010), the average posterior expectation of each good's inflation rate will lie between its true rate and the common prior. Perceived inflation rates will be biased toward the overall mean, generating the correlation bias.

¹⁷Pairwise tests for correlation confirm this expected result in our data.

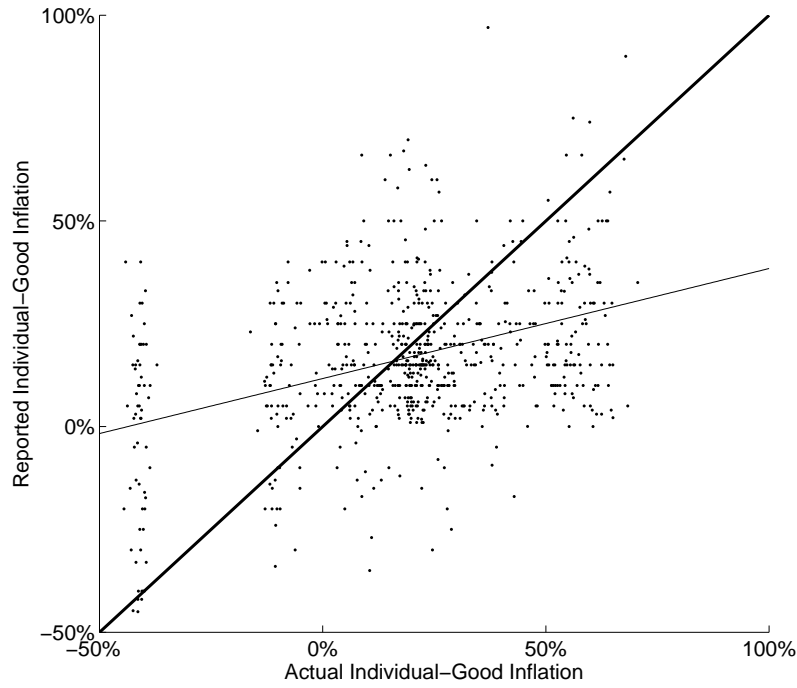


FIGURE IV. Reported versus actual inflation rates across the four individual goods.

	I. Reported Individual Inflation in Good:			
	A	B	C	D
Reported Basket Inflation (Standard Error)	0.214 (0.075)	0.234 (0.121)	0.102 (0.119)	0.450 (0.088)
	II. Loss Aversion-Adjusted Reported Rates:			
	A	B	C	D
Reported Basket Inflation (Standard Error)	0.119 (0.071)	0.329 (0.108)	0.210 (0.106)	0.342 (0.081)
Avg. Expenditure Weights	0.009	0.055	0.319	0.616
Frequency Weights	0.438	0.375	0.125	0.063

TABLE IV. Estimates for reported basket inflation regressed on reported individual-good inflation rates, compared to the expenditure-based and frequency-based weights.

significant frequency bias in the weight on Good B (p -value 0.071). Thus, we find that both the frequency bias and the correlation bias operate simultaneously, and both work together to generate the treatment effects seen in Figure II.

We can also estimate the weights in Brachinger’s (2008) Inflation Perceptions Index (IPI), in which agents use only frequency weights ($\alpha = 1$) and positive individual-good inflation rates are doubled because of loss aversion.¹⁸ The second panel of Table IV displays estimates from this regression specification. The frequency bias is significant in goods B and D (p -values 0.006 and < 0.001) but not in goods A and C (0.061 and 0.152). The root mean squared error of the second regression is noticeably larger than the first (28.389 versus 16.884), indicating that the inclusion of loss aversion does not provide a better estimate for subjects’ inflation weights. Furthermore, Brachinger’s assumption of $\alpha = 1$ appears too extreme for our data.

IV A FOLLOW-UP EXPERIMENT

The shopping phase of our main experiment is designed to be as simple as possible. Subjects are explicitly told which good to buy each day. The actual price levels do not affect earnings, and subjects are not told that they will be asked to guess inflation rates, so subjects have no incentive to pay close attention to price levels while shopping. This may exacerbate inflation perception errors, and perhaps introduce systematic biases. To examine the robustness of our main results, we ran a second experiment in which subjects are given a monthly shopping budget and can choose which good to buy each day based on their prices. They are paid for any unspent budget, and therefore have a strong incentive to watch prices and buy each good when a relatively low price is offered.¹⁹ Their only constraint is that they must buy a total of seven units of Good A, six units of Good B, two units of Good C, and one unit of Good D each month. See Figure V for a screenshot of this shopping interface.²⁰ In a debriefing survey, many subjects indicated that their strategy was to focus on the prices of the most expensive good, buying it only when a low price appeared.

We also increased awareness of price levels and inflation by initially informing subjects that, after shopping, they “will be presented with several questions about [their] shopping experience, prices, inflation, etc.” We did not explicitly tell them they would be asked to guess inflation rates, however, because subjects could make perfect guesses

¹⁸More specifically, Brachinger (2008) proposes the IPI as (using our notation) $IPI_t = \sum_{i:p_{i,t} > p_{i,0}} \left[c \left(\frac{p_{i,t} - p_{i,0}}{p_{i,0}} \right) + 1 \right] \phi_{i,0} + \sum_{i:p_{i,t} < p_{i,0}} \frac{p_{i,t}}{p_{i,0}} \phi_{i,0}$, where $c = 2$ and reference prices $p_{i,0}$ are set to be the prices of the base period. Here, we take the month-1 price as the reference price.

¹⁹Subjects’ monthly budgets were randomly drawn from a uniform distribution on $[0.95\bar{P}, 1.05\bar{P}]$, where $\bar{P} = (\sum_m \sum_i \bar{p}_{im} q_i) / 6$ is the average monthly expenditure if the subject always buys at the mean price.

²⁰The computer interface prevents subjects from buying more than the required quantity of any good by making the purchase buttons inactive for that good.

Shopping						
Shopping Day 8 of Month 1. (There are 16 shopping days per month and 6 months total.)						
Time remaining for this day: <input type="text" value="39"/>						
GOOD	Total # Needed	# Bought So Far	# Left to Buy	BRAND 1	BRAND 2	BRAND 3
A	7	1	6	\$0.96	\$1.07	\$0.96
B	6	4	2	\$6.46	\$6.35	\$6.40
C	2	2	0	\$113.86	\$123.32	\$133.86
D	1	0	1	\$493.01	\$431.50	\$489.09

FIGURE V. The shopping decision in the follow-up experiment.

simply by writing down prices in months one and six. For similar reasons, we did not show them their monthly shopping budgets or the amount remaining in those budgets until the experiment ended.

Finally, we added two new questions to the end of the experiment, after guessing basket inflation rates and individual inflation rates. First, subjects are asked to guess the average price they paid for each good in month one and month six.²¹ Then subjects are shown the actual inflation rates for each good and asked once again to guess the basket rate. This allows us to isolate the frequency bias by ‘shutting down’ the correlation bias. For each guess, subjects received 125 points minus 5 points for each percentage point of error, with a minimum payment of zero on every guess.

We ran two treatments, POS2 and NEG2, which are otherwise identical to treatments POS and NEG described in Table I. A total of 43 subjects participated in POS2 and 47 in NEG2. Five subjects (two in in POS2 and three in NEG2) are discarded for having at least one inflation guess with over 100 percentage points error. Price guesses with more than 100% error, or that are closer to another good’s true price, are removed from

²¹For example, with Good A, they are asked the average of the six Good-A prices they chose in that month.

Trt		Good				Basket	Basket (2nd)
		A	B	C	D		
	Purchases/Period	7	6	2	1		
POS2	Reported Inflation	37.74	35.05	22.77	18.27	27.01	30.89
	Actual Inflation	60.73	52.45	38.16	4.76	18.60	18.60
	Average Error	-22.99*	-17.40*	-15.39*	13.51*	8.41	12.29*
NEG2	Reported Inflation	-13.72	-1.68	8.99	22.15	19.23	13.99
	Actual Inflation	-40.97	-9.92	7.58	30.28	20.07	20.07
	Average Error	27.25*	8.24*	1.41	-8.12*	-0.84	-6.08

Note: *Average error is significantly different from zero at the 1% level.

TABLE V. Mean six-period inflation rates for each good and for the entire basket in the follow-up experiment.

our analysis.²² Sessions lasted roughly one hour. Earnings ranged from \$6.53 to \$25.40, with an average of \$16.72.

With this new data, our estimate of the magnitude of the frequency bias is given by the following regression:

$$\Pi^p = 0.206\Pi_{\text{FREQ}}^p + 0.794\Pi_{\text{EXP}}^p.$$

This estimated frequency bias of 20.6% is lower than the 44.0% seen in the main experiment, but still significantly positive with a p -value of 0.009.²³ When using subjects' second guesses of the basket rate (where actual individual-good rates are shown), the estimate increases to 25.3% and remains significant (p -value < 0.001).

Average inflation reports and errors for each treatment are shown in Table V. Results are similar to the original experiment (Table III), though initial basket errors are greater in Treatment POS2 (8.41 compared to 5.90) and much smaller in NEG2 (-0.84 compared to -9.90). Neither initial basket error rate is significantly different from zero in the new experiment (partly due to a smaller sample size), but they are significantly different from each other (large-sample permutation test p -value of 0.044). Individual-good reports are less extreme than actual inflation rates, demonstrating that the correlation bias is still present. The ranges of individual-good reports are again too small, with 7.3% of subjects reporting the same rate for all four goods. When individual-good rates are shown to subjects, their guesses of the basket rate actually become *worse* in both

²²This constitutes 8.7% of our data. Almost all of these are obvious transposition errors, where the subject's guess of Good A's price is entered as a guess for Good D, for example.

²³The estimated α is 22.4% if coefficients are not constrained to sum to one, and 21.6% if, in addition, a constant is added to the regression. Both remain significant at the one percent level.

	Reported Individual Inflation in Good:			
	A	B	C	D
1st Reported Basket Inflation (Standard Error)	0.080 (0.116)	0.127 (0.139)	0.316 (0.179)	0.477 (0.128)
	Actual Individual Inflation in Good:			
	A	B	C	D
2nd Reported Basket Inflation (Standard Error)	-0.181 (0.183)	0.576 (0.446)	0.276 (0.409)	0.330 (0.142)
Avg. Expenditure Weights	0.009	0.056	0.325	0.611
Frequency Weights	0.438	0.375	0.125	0.063

TABLE VI. Estimates for basket inflation reports regressed on (reported or actual) individual-good inflation rates, compared to the average expenditure-based and frequency-based weights.

treatments. The error in POS2 is significantly different from zero, and from the error in NEG2 (both p -values < 0.001). Thus, shutting down the correlation bias exacerbates subjects' errors in aggregating the inflation rates, in exactly the direction predicted by the frequency bias.

Regressing the (initial) reported basket rate on the individual-good rates (top third of Table VI) again reveals point estimates between the expenditure weights and the frequency weights. Estimated weights are negatively correlated with frequency of purchase, indicating that estimates are based more heavily on expenditure shares than frequency weights. With the smaller sample size, we cannot reject the hypothesis that the actual weights equal the expenditure weights, but we can reject the frequency weights at the 5% level for goods A and D. These facts are consistent with our estimate of α being relatively small. Once again, multiplying positive individual-good inflation rates by $c = 2$ to account for loss aversion (Brachinger, 2008) only increases the regression's root mean squared error (to 27.811 from 18.126), indicating a worse fit.

The middle third of Table VI show a regression of subjects' second guess of the basket inflation rate on the true individual-good rates. The pattern of estimated weights is less clear, and only the frequency weight for good A can be rejected at the 5% level.

Average errors in average price guesses, expressed as a percentage of the true average price, are shown in Table VII. Surprisingly, the correlation bias appears in these price guesses as well. If subjects recall only the overall average price for each good (averaging across all months), and believe inflation rates to be similar, then for goods with high inflation they will overestimate month-1 prices and underestimate month-6 prices. For goods with low inflation they will underestimate month-1 prices and overestimate

Month	Trt.	Good			
		A	B	C	D
1	POS2	3.61%	5.17%	12.90%	-3.95%
	NEG2	-9.84%	-3.92%	8.40%	5.38%
	<i>p</i> -value	(< 0.001)	(0.005)	(0.150)	(0.999)
6	POS2	-0.98%	-1.95%	6.08%	4.73%
	NEG2	18.14%	3.25%	2.25%	-1.44%
	<i>p</i> -value	(> 0.999)	(0.992)	(0.193)	(< 0.001)

TABLE VII. Average percentage error in price guesses by treatment and month. *p*-values from one-sided permutation tests for differences in means shown in parentheses.

month-6 prices. This is consistent with the data: One-sided large-sample permutation tests confirm significant differences in price errors between treatments for goods A, B, and D, in both month 1 and month 6.²⁴ Again, this correlation bias can partially explain the treatment effect in the first basket rate guess, because subjects are overestimating price growth of the expensive goods (C and D) in treatment POS2, and underestimating price growth of expensive goods in treatment NEG2.

Subjects who incorrectly recall prices should also have incorrect memories of expenditure shares. Thus, even if a subject is trying to use the correct expenditure-based weights, their basket inflation estimates will appear systematically biased. For example, in POS2, subjects overestimate prices of Good A and therefore should overestimate the month-1 expenditure share of Good A. Since Good A has the highest inflation, second basket guess errors would be positive, even if subjects have no frequency bias. And this is exactly what we observe. In NEG2, however, subjects would overestimate the expenditure share of Good D, which now has the highest inflation rate, and so second basket errors should again be positive. But we observe negative errors in NEG2, contradicting this theory. In other words, the apparent frequency bias cannot be explained away by assuming erroneous expenditure shares based on recalled prices.

We conclude the existence of a frequency bias appears robust to details of the shopping experience, and whether subjects are aware that questions about inflation will follow. The magnitude of the bias, however, is diminished when prices are more salient and subjects are more focused on inflation. The fact that the correlation bias appears in recalled prices raises an interesting question about its source: Are prices calculated from biased perceptions of inflation, or are inflation rates calculated from biased perceptions

²⁴Significance obtains if the *p*-value is below 0.05 for goods with high inflation, or above 0.95 for goods with low inflation.

of price? Regardless, the observed treatment effects in the second basket rate guess indicate that the recalled prices are not used to formulate expenditure shares, leaving room for the frequency bias to explain the observed pattern of errors.

V SUGGESTIVE EVIDENCE FROM SURVEY DATA

Do biases observed in a hypothetical laboratory setting with college undergraduates extend to real-world inflation perceptions? We briefly examine inflation perceptions survey data in search of evidence that the frequency bias is in fact present. First, it has been well documented by surveys of European consumers' perceptions of inflation that general public throughout the EMU member countries perceived significantly higher inflation in the years 2002–06 following the introduction of the euro (Figure VI).²⁵ In general, prices of goods were rounded up after the conversion from local currencies to the euro. This meant that the prices of low-value but frequently-purchased goods increased significantly (e.g. from 1.70 to 2 euro) while for more valuable but infrequently shopped goods the increase was insignificant (e.g. from 980.70 to 981 euro). In terms of total expenditure, this rounding effect was trivial; however, for individuals whose daily purchases became noticeably more expensive, the perceived effect was large.²⁶

In a monthly survey conducted in Ohio for a shorter time period (August 1998 to November 2001), consumers were asked how much they think prices rose over the past 12 months. The average reported rate was 6 percent, while the actual increase in CPI was only 2.7 percent (Bryan and Venkatu, 2001b). But inflation rates for non-durable goods have been systematically higher than inflation rates for durable goods over the past twenty years, with the difference becoming large after 2002. Since non-durable goods are purchased more frequently, the frequency bias correctly predicts that consumers' perceptions of inflation over this time period are greater than the CPI inflation rate calculated using expenditure-weight.

To expand on this argument, consider consumer experiences in the United States for the last two decades. There have been notable increases in the prices of low-price, everyday goods, e.g. food and beverage (3.1 percent), energy (4 percent), and transportation (2.6 percent), and a much smaller rise or even decline in the prices of the relatively expensive consumer goods, e.g. apparel (0.58 percent), audio-visual devices (0.4 percent),

²⁵In this survey, participants are asked "How do you think prices have developed in the last 12 months?" and are given possible answers "risen a lot", "risen moderately", "risen slightly", "stayed about the same", and "fallen". The reported indicator is a linear combination of the frequency of responses given to each answer.

²⁶See Del Giovane and Sabbatini (2008), for example.

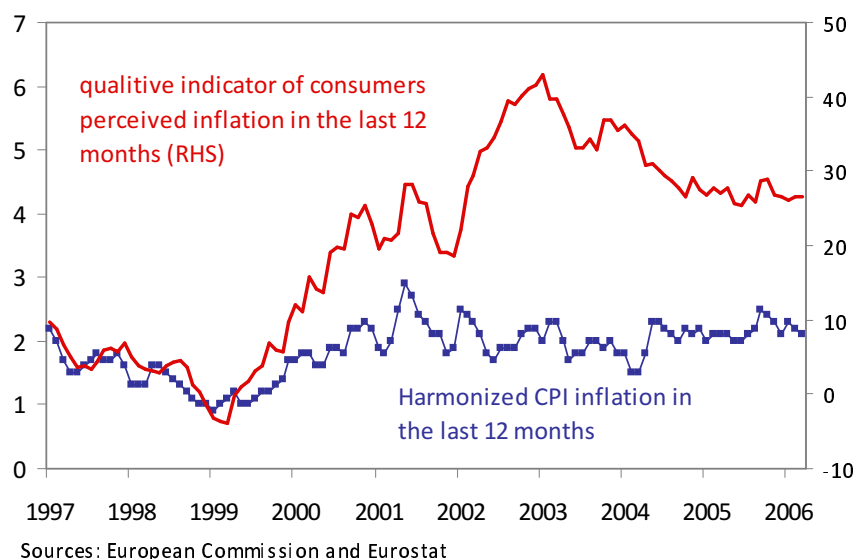


FIGURE VI. Actual inflation vs. a qualitative indicator of perceived Inflation in the Euro area

and information technology (-11.1 percent). The aggregate effect is a relatively low overall CPI inflation rate. Consumers more frequently experienced the goods with higher inflation rates, however, leading to the apparent upward biases in inflation perceptions and expectations.

Although this correlation is merely suggestive, it does support the claim that the experimental results may have macroeconomic consequences.²⁷ A thorough investigation of the ‘frequency weighted’ price statistics using micro-data on inflation perception and shopping frequency of individual goods would be an interesting topic for future research.

VI DISCUSSION AND FUTURE RESEARCH

We have shown that people misperceive inflation in a controlled lab experiment, biasing their perceptions of economy-wide inflation toward the inflation rates of the more-frequently purchased goods. One implication is that macroeconomic analysts should

²⁷It is interesting to note that although we observe frequency bias in average consumers, professional forecasters do *not* appear to exhibit any such bias. In fact, professional inflation forecasts are generally accurate (Keane and Runkle, 1990, e.g.). However, consumers only occasionally pay attention to news reports of inflation forecasts (Carroll, 2003). In most macro policy applications, what matters is consumers' inflation perceptions. It is in the failure to predict the consumers' responses that policymakers' models may suffer in predictive power. It is for this reason that policymakers should care about the existence of the frequency bias.

consider adjusting for the frequency bias in inflation perceptions. In this sense, our paper provides experimental evidence supporting a modified version of Brachinger's (2008) Index of Perceived Inflation, in which the frequency bias is weakened ($\alpha < 1$) and loss aversion is ignored ($c = 1$).

An open question is whether the frequency bias is attenuated with experience. Although our experiment cannot address this question, we conjecture that adjustments in perceptions would be very slow. Learning is fastest when feedback about mistakes is clear (see Weber, 2003, e.g.). Small mistakes in consumption-savings decisions, however, are unlikely to provide informative negative feedback. Thus, consumers will feel little to no pressure to adapt their method of aggregation. As evidence that the bias may not attenuate, recall that the over-estimation of inflation observed in the field has persisted over the past 25 years, suggesting little learning from experience. We hypothesize that the accuracy of expert forecasts (Keane and Runkle, 1990) does not come from recalling past shopping experiences, but rather from considering prices and inflation analytically.

The frequency bias that was documented in the present study may also be present in other situations where agents have to aggregate different pieces of relevant information to form a perception of current trends. Investors in financial markets observe the movement of individual prices in multiple occasions over a given time period. As far as the frequency of price information they get about specific shares is not equal to the weight they have in the general index, investors' perceptions of the general trends in the stock market could be biased. Another example can be found in the field of mass media: Receivers of news will get the same news item multiple times, which might bias their perception of reality. Such a phenomenon can be relevant in the field of political economy. Suppose a given candidate has one good and one bad characteristic. Even voters with neutral priors might underestimate the relative merits of this candidate if they receive reminders about the negative characteristic more frequently than they receive reminders about the positive one. This justifies the extensive use of advertising in political campaigns. An extensive study of such questions remains for future research.

REFERENCES

- Adam, K., April 2007. Experimental evidence on the persistence of output and inflation. *Economic Journal* 117 (520), 603–636.
- Bates, J. M., Gabor, A., 1986. Price perception in creeping inflation: Report on an enquiry. *Journal of Economic Psychology* 7, 291–314.
- Bernanke, B. S., 2007. Inflation expectations and inflation forecasting. Speech at NBER Summer Institute Monetary Economics Workshop.
- Blanchard, O., Dell’Ariccia, G., Mauro, P., 2010. Rethinking macroeconomic policy. *Journal of Money, Credit and Banking* 42, 199–215.
URL <http://dx.doi.org/10.1111/j.1538-4616.2010.00334.x>
- Blanchflower, D. G., MacCoille, C., 2009. The formation of inflation expectations: An empirical analysis for the u.k., nBER working paper No. 15388.
- Brachinger, H. W., 2008. A new index of perceived inflation: Assumptions, method and application to germany 29, 433–457.
- Bryan, M. F., Venkatu, G., November 2001a. The curiously different inflation perspectives of men and women. Tech. rep., The Federal Reserve Bank of Cleveland.
- Bryan, M. F., Venkatu, G., October 2001b. The demographics of inflation opinion surveys. Tech. rep., The Federal Reserve Bank of Cleveland.
- Burke, M. A., Manz, M., 2011. Economic literacy and inflation expectations: Evidence from a laboratory experiment, discussion Paper No. 11-8, Research Center for Behavioral Economics, Federal Reserve Bank of Boston.
- Camerer, C. F., March 1992. The rationality of prices and volume in experimental markets. *Organizational Behavior and Human Decision Processes* 51 (2), 237–272.
- Carroll, C. D., 2003. Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics* 118 (1), 269–298.
- Chambers, C. P., Healy, P. J., 2010. Updating toward the signal, forthcoming, *Economic Theory*.
- Del Giovane, P., Sabbatini, R., 2008. The Euro, Inflation and Consumers’ Perceptions: Lessons from Italy. Springer-Verlag, Berlin.
- Del Missier, F., Ranyard, R., Pietroni, A., 2008. Fruitful interactions between psychology and behavioral economics: The case of perceived inflation. *Sistemi Intelligenti* 20, 39–56.
- Duffy, D., Lunn, P. D., 2009. The misperception of inflation by irish consumers. *The Economic and Social Review* 40 (2), 139–163.

- Fluch, M., Stix, H., 2005. Perceived inflation in Austria—extent, explanations, effects. *Monetary Policy & The Economy* 3, 22–47.
- Fréchette, G. R., 2009. Laboratory experiments: Professionals versus students, New York University working paper.
- Garner, A. C., 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, pp. 113–128.
- Hey, J. D., December 1994. Expectations formation: Rational or adaptive or...? *Journal of Economic Behavior and Organization* 25 (3), 329–350.
- Huber, O. W., 2011. Frequency of price increases and perceived inflation. An experimental investigation. *Journal of Economic Psychology* 32, 651–661.
- Jonung, L., 1981. Perceived and expected rates of inflation in Sweden. *American Economic Review* 71, 961–968.
- Jonung, L., Laidler, D., 1988. Are perceptions of inflation rational? Some evidence for Sweden. *American Economic Review* 78 (5), 1080–1087.
- Jungermann, H., Brachinger, H. W., Belting, J., Grinberg, K., Zacharias, E., 2007. The euro changeover and the factors influencing perceived inflation. *Journal of Consumer Policy* 30, 405–419.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decisions under risk. *Econometrica* 47, 263–291.
- Keane, M. P., Runkle, D. E., September 1990. Testing the rationality of price forecasts: New evidence from panel data. *American Economic Review* 80 (4), 714–735.
- Leung, C., June 2009. The demographics of household inflation perceptions and expectations. Reserve Bank of New Zealand: Bulletin 72 (2), 34–42.
- Malmendier, U., Nagel, S., 2012. Learning from inflation experiences, Stanford University Working Paper.
- Mazumdar, T., Raj, S., Sinha, I., 2005. Reference price research: Review and propositions. *Journal of Marketing* 69, 84–102.
- Mishkin, F. S., June 1981. Are market forecasts rational? *American Economic Review* 71 (3), 295–306.
- Pfajfar, D., Zakelj, B., 2009. Experimental evidence on inflation expectation formation, discussion paper 2009-07, Tilburg University, Center for Economic Research.
- Ranyard, R., Del Missier, F., Bonini, N., Duxbury, D., Summers, B., 2008. Perceptions and expectations of price changes and inflation: A review and conceptual framework. *Journal of Economic Psychology* 29, 378–400.

- Schmalensee, R., January 1976. An experimental study of expectation formation. *Journal of Money, Credit and Banking* 44, 17–41.
- Siegel, S., Castellan, Jr., N. J., 1988. *Nonparametric Statistics for the Behavioral Sciences*, 2nd Edition. McGraw-Hill, New York, NY.
- Tversky, A., Kahneman, D., 1973. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology* 5, 207–232.
- Weber, R. A., 2003. ‘learning’ with no feedback in a competitive guessing game. *Games and Economic Behavior* 44, 134–144.