FREQUENCY BIAS IN CONSUMERS' PERCEPTIONS OF INFLATION†

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

ABSTRACT. We provide two independent sources of evidence that consumers' perceptions of inflation are systematically biased toward the inflation rates of the most-frequently purchased items. Surveys of inflation perceptions and expectations show that consumers over-perceive economy-wide inflation during periods where frequently-purchased non-durable goods are inflating faster than durables. Separately, in a controlled laboratory experiment, subjects' estimates of aggregate inflation is clearly biased toward the inflation rates of the most-frequently purchased goods. Taken together, these results suggest that the frequency bias is a fundamental feature of consumers' inflation perceptions that can distort optimal life-cycle decisions.

Key words: Inflation perceptions; biases; lab experiment
JEL classification codes: C91; E31; C82.

August 30, 2011

†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.
**Dept. of Economics, The Ohio State University, 1945 North High street, Columbus, OH 43210, U.S.A.; healy.52@osu.edu.
***Dept. of Economics, The Ohio State University, 1945 North High street, Columbus, OH 43210, U.S.A.; nanli@mail.econ.ohio-state.edu.

I Introduction

Expectations about inflation play a central role in a wide range of macroeconomic analyses because these expectations influence many economic decisions, such as consumption, savings, investment, wage bargaining, and capital flows. The efficacy of monetary policy depends more heavily—in most cases—on the anticipated inflation than on the actual realized inflation rate. Consequently, understanding the formation of these expectations has become increasingly important to understand and influence macroeconomic outcomes.

To form expectations about future inflation, consumers often rely on perceptions about recent price changes.[1] If those perceptions and resulting expectations are systematically biased, then so too will be consumers' intertemporal decisions. Responses to macroeconomic policy changes will be affected. Policymakers therefore ought to take these biases into account. In order to design optimal policies, it is important to understand how consumers perceive price changes, and discover any systematic mistakes they may exhibit in forming their beliefs about future inflation.

In this paper, we use both survey evidence and laboratory experiments to identify a specific and important bias in the perception of inflation: When forming economy-wide inflation perceptions, consumers tend to overweight price changes of frequently-purchased goods. We refer to this phenomenon as the *frequency bias* in inflation perceptions.

A rational, unbiased consumer who is calculating the overall inflation rate should weight each individual good's inflation rate by the percentage of total expenditures spent on that good. Our evidence suggests that consumers' actual weights are biased, placing higher weights on those goods that are purchased more frequently. Goods like food and gasoline are frequently purchased but constitute a relatively small share of total expenditures. Housing expenditures, on the other hand, are substantial but made infrequently. We suggest that consumers' perceptions of overall inflation are biased toward the inflation rates of goods like gasoline and food, and away from the inflation rates of low-frequency goods like housing.

The presence of a frequency bias distorts the choices of economic agents and can reduce welfare. When a frequently-purchased commodity experiences an idiosyncratic

[1]In fact, economists have used the past (weighted) inflation rates to approximate the expected rate, implicitly assuming that agents correctly perceive past inflation rates and then use those perceptions to form their expectations. Examples of autoregressive formation of inflationary expectations include adaptive expectations, extrapolative or regressive schemes, *et cetera*.

price change, consumers' perceptions of overall inflation—and, therefore, their expectations of future inflation—will be disproportionately affected. Distorted inflation expectations lead in turn to a suboptimal intertemporal allocation of resources; for example, consumers misperceiving inflation to be too high may save less than they should or invest in the wrong asset classes.[2]

The distorting effect of the frequency bias on individual choices has important implications for monetary and fiscal policy. The present study suggests clear ways for policy makers to adjust their models of inflation expectations to better inform their own policy decisions. Anticipating frequency effects by estimating consumer inflation expectations with frequency-biased weights will lead to more accurate predictions about expectations and, ultimately, more effective policy decisions and inflation controls.

Using existing survey data on consumers' perceptions and expectations of inflation, we find that over-perception of inflation is correlated with the difference between inflation rates of non-durable and durable goods. Since non-durables are purchased more frequently than durables, the observed correlation suggests that consumers may bias their perception of economy-wide inflation toward the inflation rate of more frequently-purchased goods. In other words, consumers who see fast-rising gasoline prices may erroneously conclude that the economy-wide inflation is high.

In a separate and novel laboratory experiment, we ask subjects to engage in a sequence of simulated 'shopping' decisions. Subjects' retrospective estimates of the inflation rate of a representative bundle of goods is systematically biased toward the inflation rates of the most-frequently purchased goods. The existence of this bias in a controlled laboratory setting suggests that it may be a fundamental attribute of how individuals aggregate various price changes. Such fundamental biases in calculation would still apply when the sterile context of the laboratory is removed.

Evidence for the existence of the frequency bias is bolstered by the fact that we observe it using these two separate methodologies. There are certainly alternative explanations for the correlation observed in the survey data, and there may be concerns that observed phenomena in the controlled laboratory setting do not generalize to the field. But alternative explanations for the survey data cannot explain the measured bias in the laboratory, and concerns about the generalizability of laboratory results must be reconciled with the observed phenomenon in the field.[3] In other words, the laboratory

[2]For example, some assets such as gold are widely believed by the lay public to be safeguards against high inflation. Investors perceiving inflation to be rising might over-invest in such assets.

[3]For example, if non-durables experience higher rates of quality improvement over time and households fail to quickly adjust to these quality changes, then households may incorrectly perceive higher inflation

environment provides the control necessary to isolate and measure the bias suggested by the survey data, and the survey data confirms that the bias found in the laboratory may generalize to the field setting.

The laboratory is especially useful for the current study because biases in expectations and perceptions are particularly difficult to identify in the field. Without the control of the laboratory, expectations can only be elicited through surveys or inferred from behavioral data. Information sets are difficult or impossible to measure and cannot be controlled. To illustrate, consider the long literature testing the rational expectations hypothesis using survey data. Most studies generally reject the claim that expectations are rational (see Figlewski and Wachtel, 1981, e.g.), but such inquiries are ultimately inconclusive because nearly any expectation can be rationalized by some unobservable beliefs about unrealized events. Perceptions of inflation may also appear biased because the current representative basket of goods used to calculate official inflation estimates may not accurately represent the true purchases of many consumers in the economy. Consumers may substitute between goods, shifting the true basket composition away from the basket used by the statistical agency, or the quality of goods may change over time, affecting perceptions of inflation. It is often difficult to disentangle these possible explanations. Furthermore, there may be a disconnect between survey responses and the actual expectations used to make economic decisions (Mishkin, 1981).

In the experiment, however, information sets are both observed and manipulated. Resulting beliefs can be elicited in an incentive-compatible way. The basket of purchased goods is chosen by the experimenter in advance and perfectly observed. Quality of goods can be abstracted away. Accurately measuring biases is entirely feasible with this level of control and measurement.

Although we observe the 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.). Since policymakers typically rely on professional forecasts, this component of their models is likely unbiased. 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.

It may seem that biases such as the frequency bias should be transitory because agents can easily observe official inflation estimates and should adjust their perceptions

rates on those goods. But any quality-based explanation such as this cannot explain the frequency bias in the laboratory because quality is held constant in the experiment.

accordingly. Even if agents do use this information, however, it may not fully mitigate the bias since consumers can recognize that the representative basket of goods used to calculate official estimates typically varies substantially from their own expenditure habits. The more a consumers' basket differs from the representative basket, the less informative the official estimates will be for his personal economic decisions, and the less his bias will be moderated by observing those estimates.

Finally, note that the frequency bias is a bias of *aggregation*. It neither claims nor requires that perceptions of individual-good inflation rates be accurate. It is not caused by individuals remembering only the prices of more-frequently purchased goods and forgetting the prices of large, infrequent purchases. It only describes cognitive errors in aggregating the perceived individual-good rates to form a perception of the economy-wide inflation rate. It is the economy-wide inflation rate, however, that individuals must use when making financial investment and savings decisions, and so the bias will manifest in distorted allocations of financial assets.

Our experimental data clearly show that subjects also make errors in estimating individual-good rates. Specifically, the perceived individual-good rates are too highly correlated, relative to true rates. Since our experimental treatments assign more extreme inflation rates to the more frequently-purchased goods, the estimates for those goods contain the most error. This result refutes the claim that the frequency bias is simply a bias in memory; the most frequently-encountered goods are the ones for which subjects' inflation estimates are the most erroneous. The observed *correlation bias* in individual-good inflation rates is discussed briefly in Section V.

We review related literature in the following section, and we formally define the frequency bias in Section III. The survey evidence is explored in Section IV. In Section V we describe the experimental design and present the various results from the laboratory. Section VI concludes.

II Related Literature

To our knowledge, this is the first experimental study that directly investigates behavioral biases in the perception of past inflation. A handful of experimental studies have looked at expectations of future inflation in specific contexts. Schmalensee (1976) showed twenty-three subjects price data from an eighteenth-century British wheat market. He found that predictions of subsequent wheat prices roughly followed an adaptive expectations process, and that forecasters' confidence was affected by past performance and 'turning points' in the price data. The study had no scope for the frequency bias since prices for only one good were observed. Garner (1982), Hey (1994), and others also explore how subjects' forecasts of future prices are determined by past prices, and Adam (2007) looks at forecasts of future inflation rates in a simulated macroeconomic model where forecasts affect future rates. Several market experiments have subjects making predictions about future prices (see Williams, 1987 or Camerer, 1992, for example); the rational expectations hypothesis is generally rejected in favor of adaptive expectations models. None of these studies feature the aggregation of multiple prices to generate predictions or perceptions about economy-wide inflation rates. Thus, none has the scope to study issues related to the frequency bias.

There are many non-experimental studies that use survey data to examine economic agents' inflation perceptions and expectations (e.g., Wachtel, 1977; Friedman, 1979; Figlewski and Wachtel, 1981; Gramlich, 1983; Batchelor, 1984; Jonung and Laidler, 1988). These studies show that perceptions and expectations of inflation are quite disperse, particularly across different demographic groups. For example, Bryan and Venkatu (2001) document that individuals with income below the 20th percentile predict inflation twice as high as the highest income quintile.

Using an early sub-sample of the Swedish survey data, Jonung (1981) finds that Swedish women in 1977 had higher perceptions of inflation than did Swedish men. Jonung even suggests a frequency-bias explanation for his results: Since Swedish women during that time period were five times more likely to shop for food than men, and since the 1977 inflation rate for food was significantly greater than that of the overall consumer price level, women were led to have upwardly biased expectations. When both genders are pooled, Van Duyne (1982) finds no evidence that Americans placed excess weight on recent food prices when forming inflation expectations during the 1970s, though he does cite the 1976 *Economic Report of the President* as saying:

> "Food prices are the most visible and best publicized of all the components of the CPI. For this reason they may be especially important in determining the wage demands of labor and the inflationary expectations of all consumers." (p. 34; reprinted in Van Duyne, 1982, p. 419)

Thus, the frequency bias phenomenon has been suggested in past work, but without the control of the laboratory it was difficult to provide anything more than suggestive evidence about its existence.

Our paper is also related to the 'sticky information' literature in macroeconomics. There, the apparent inertia in inflation expectations can be explained by assuming that the acquisition of inflation-related information is costly, so rational agents may sample this information infrequently, leading to more sluggish inflation perceptions (Feige and Pearce, 1976; Galbraith, 1988; Mankiw and Reis, 2002, 2003). The frequency bias would work within the framework of the sticky information literature: If agents' samples are comprised of disaggregated price data, then their aggregate inflation perceptions would be biased toward frequently-purchased goods. Thus, inflation would be both sluggish to update and frequency biased.

III 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

$$P_t = \sum_i q_i p_{it}. \tag{1}$$

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 $\Pi_t = (P_{t-1} - P_t)/P_{t-1}$. 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 period-t expenditure. Thus, we must have

$$\Pi_t = \sum_i \theta_{it}\pi_{it}, \tag{2}$$

where

$$\theta_{it} = \frac{q_i p_{it}}{\sum_j q_j p_{jt}}. \tag{3}$$

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

$$\Pi^p = \sum_i \omega_i \pi_i^p, \tag{4}$$

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.[4]

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, and μ_i is the average quantity per purchase. The *frequency weight* of good i is given by

$$\phi_i = \frac{n_i}{\sum_j n_j} \tag{5}$$

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,

$$\omega_i = \alpha \phi_i + (1-\alpha)\theta_i, \tag{6}$$

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 $\alpha = 0$, for example. Given this α, the perception of the overall inflation rate is calculated as

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

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

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

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

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.

[4] If 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.

IV Survey Evidence

In practice, inflation is measured through the changes in a consumer price index (CPI), most often the All Items CPI for All Urban Consumers (CPI-U), reported by the United States Bureau of Labor Statistics. The CPI-U is calculated by forming a basket of goods that represents the purchases of a typical American consumer living in an urban area. The price for each good in the basket is surveyed, the total price of the entire basket is calculated, and the resulting total expenditure is normalized to that of some base year.[5] The 1996 Boskin Commission Report concluded that reported CPI-U inflation is systematically overestimated mainly due to quality improvement in consumption goods. Since the Commission's report, many improvements were introduced into the CPI. Nevertheless, today's upward bias is still estimated to be at least 1.0 percent per year (Gordon, 2006). The Chain Weighted CPI (C-CPI) updates the base year frequently, but is only available for limited time period.

As an example of how the frequency bias operates in practice, consider the rapid increase in gasoline prices that occurred during the Spring and early Summer of 2008. In most cities in the United States, gas prices are displayed prominently outside each gas station, making gasoline inflation particularly salient to the average consumer. During that same period, consumers' expectations of inflation spiked well above the realized inflation rate. This is seen in Figure I, which compares the actual quarterly CPI inflation rates (both CPI-U and the C-CPI) to the average household's inflation expectation taken by the Survey Research Center at the University of Michigan, from 1985Q1 to 2010Q4.[6] In early 2008, consumers apparently placed more weight on gas price changes than did the CPI-U and C-CPI, perhaps because these gas price changes were among the most visible and salient.

Figure I reveals a general trend of expectations significantly exceeding actual inflation over the past twenty years.[7] There have been several attempts to rationalize the systematic error in consumers' expectations. For example, the "Peso problem" explanation (see Rogoff, 1980; Krasker, 1980; Lizondo, 1983; Campbell and Shiller, 1991;

[5]The current base year is the average of a 36-month period covering 1982, 1983, and 1984.

[6]These respondents are asked to provide forecasts of inflation rate over the next year of "things you buy".

[7]Lower income groups typically report higher perceived inflation than higher income groups. It appears that the higher inflation perception of the average consumer is an artifact of disproportionately overvaluing the inflation opinion of the lower-income respondents. However, Kokoski (2000) shows that using population demographics as an alternative weight to construct CPI does not significantly affect the calculation result, suggesting that one needs to look elsewhere for explanations.

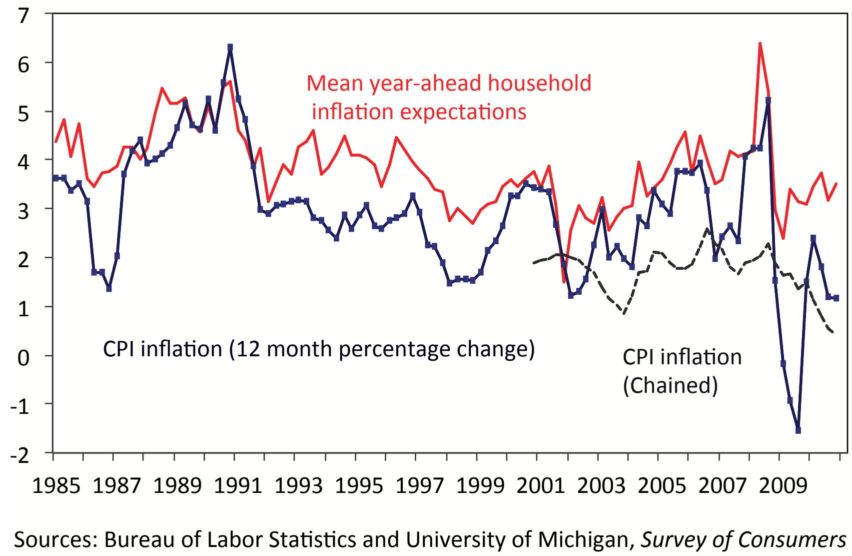

FIGURE I. Actual inflation and consumers' inflation expectations in the U.S.

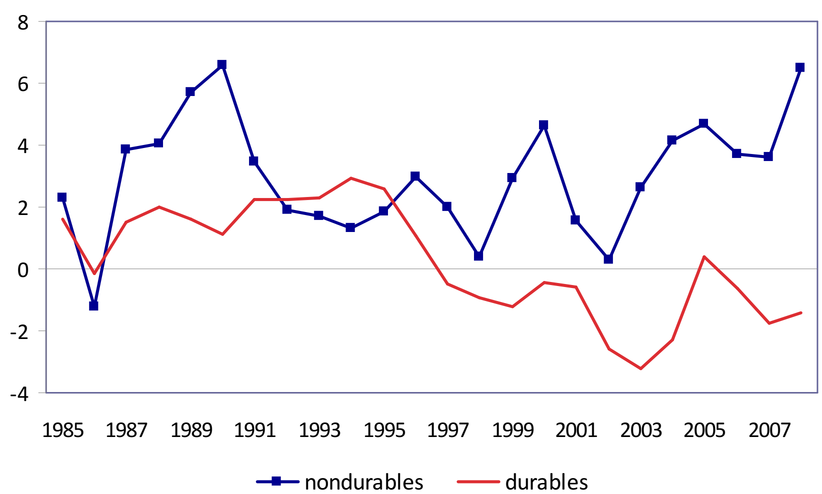

Source: Bureau of Labor Statistics, all urban consumers

FIGURE II. Annual inflation rates of durable goods and non-durable goods

Bekaert et al., 2001, e.g.) claims that agents rationally believe there is a small probability of a very large increase in inflation, leading to expectations that almost always look biased, *ex post*.

We suggest instead that this systematic error in inflation expectations can be explained by the frequency bias in inflation perceptions. Figure II shows that inflation

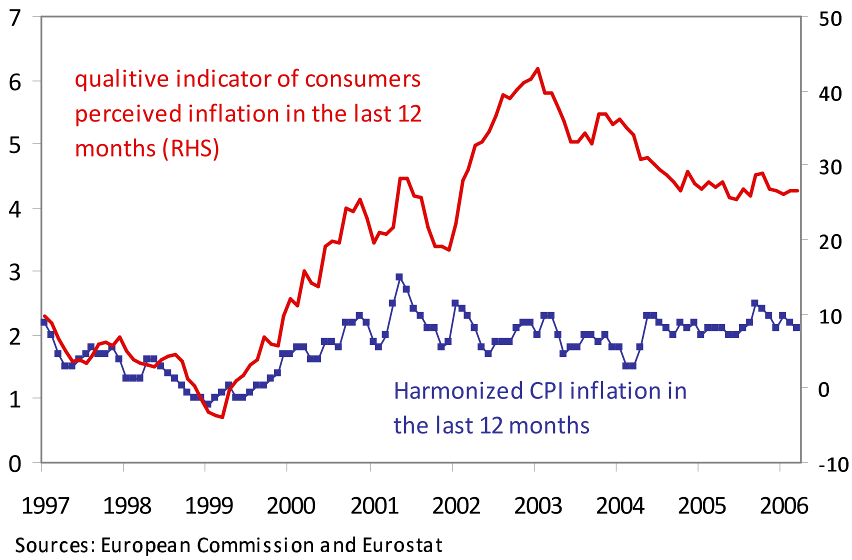

FIGURE III. Actual inflation vs. a qualitative indicator of perceived Inflation in the Euro area

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 predicts that consumers' perceptions of inflation over this time period are greater than the inflation rate calculated using an expenditure-weighted index like the CPI-U or C-CPI. A monthly survey in Ohio of consumers' perceptions of past price increases confirms this prediction. Respondents reported about 6 percent on average for the period where the actual increase in CPI was only 2.7 percent (Bryan and Venkatu, 2001). Similarly, a survey of European consumers' perceptions of recent inflation shows significant overestimation of recent inflation in the years 2002–06 following the introduction of the euro (Figure III).[8] 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).

[8]In 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.

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.[9]

To expand on this argument, consider consumer experiences in the United States from 1985–2009. 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), and information technology (-11.1 percent). The aggregate effect is a relatively low overall inflation rate, according to the CPI. Consumers more frequently experienced the goods with higher inflation rates, however, leading to the apparent upward biases in inflation perceptionsand expectations.

V LABORATORY EVIDENCE

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.[10] All subjects were Ohio State undergraduate students recruited via e-mail.[11] All sessions took place in the Ohio State Experimental Economics Laboratory. In total, 186 subjects participated in the experiment in sessions of roughly 21

[9]See Del Giovane and Sabbatini (2008), for example.

[10]One 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.

[11]Existing 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 Fréchette (2009) for details.

subjects per session. Each subject was only allowed to participate in one session of this experiment.

For each session, all subjects arrived at the laboratory simultaneously, were seated at computer terminals, and told to log into the experiment website.[12] 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 IV. 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.

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

[12]The website is available at `http://healy.econ.ohio-state.edu/exp/shopping`. To experience this experiment, log in using session name 'test'.

FIGURE IV. 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.

and food—are purchased more frequently than others. Multiple prices for the desired 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. In each subsequent month, the mean price for each good is inflated by a monthly inflation rate π_i^* that does not vary during the experiment. 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 $\bar{p}_{i,m+1} = \bar{p}_{im}\left(1+\pi_i^*\right)$.

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 IV 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.

As described in Section III, 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

$$\Pi = \sum_i \theta_i \pi_i. \tag{8}$$

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.[13] 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 $125 - 500\left|\pi_i^p - \pi_i\right|$ 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 point difference between a guess and that good's true inflation rate. Again, earnings

[13]Before 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.

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

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

that was derived in equation (7) above. The values Π^p_{EXP} and Π^p_{FREQ} 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 are significantly different from the theoretical predictions of 0% and 100%, respectively, with p-values less than 0.001. Thus, the frequency bias is both statically significant and economically meaningful in size; nearly half of agents' expectations are derived from their frequency of purchase.

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

$$\Pi^p = 0.528\Pi^p_{\mathrm{EXP}} + 0.428\Pi^p_{\mathrm{FREQ}}.$$

Both coefficients are significantly different from both zero and one at the five-percent level. Also allowing for a constant gives an estimated relationship of

$$\Pi^p = 8.187 + 0.257\Pi^p_{\mathrm{EXP}} + 0.419\Pi^p_{\mathrm{FREQ}}.$$

The positive constant is significant, indicating 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.[14]

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^p_{FREQ}) 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 the estimated α is 1.01 with a p-value less than 0.001.

As is common in experimental studies, the exact magnitude of the effect is difficult to pin down, but its presence is apparent; the real power of the experimental method comes in studying treatment effects, which we analyze in the following subsection.

Treatment Differences

Figure V 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).

[14]Because Π^p_{EXP} and Π^p_{FREQ} 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 is 0.478. Both are below most thresholds for concern. Finally, regressing Π^p on Π^p_{FREQ} alone gives an estimated relationship of $\Pi^p = 10.286 + 0.568\Pi^p_{\mathrm{FREQ}}$ with a p-value of less than 0.001 on the slope coefficient.

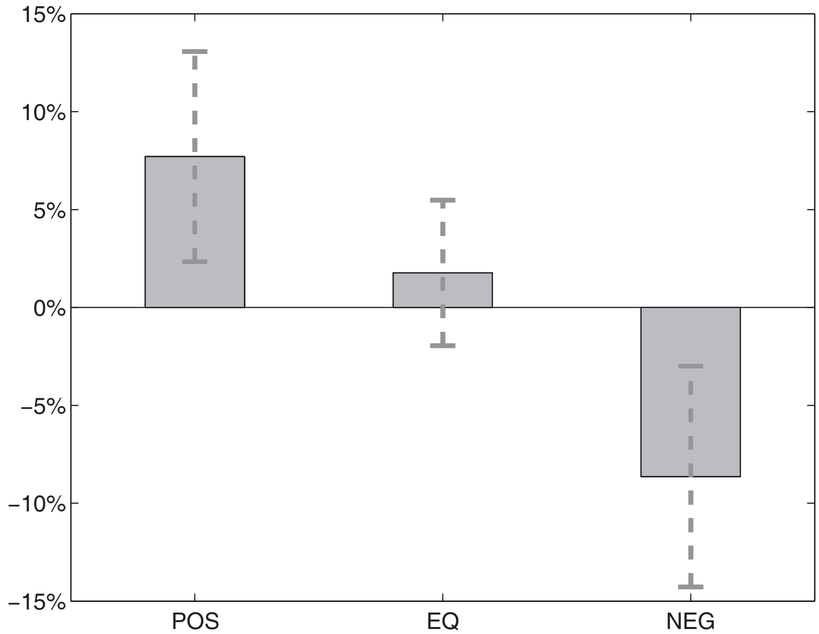

FIGURE V. Average errors (in percentage points) in economy-wide inflation perceptions by treatment. The dashed lines represent 95% confidence intervals.

Trt		Good A	Good B	Good C	Good D	Basket
	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.

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 V 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 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 (Wilcoxon rank-sum test p-values of < 0.001 and 0.004, respectively), though errors in Treatments POS and EQ are not significantly different (p-value of 0.202).

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-inflating good and overestimate the rate of the lowest-reporting good. 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.

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.[15] 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.[16] 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 VI shows the relationship between actual inflation rates and reported rates for the four goods. A linear regression shows a relationship of 0.27, significantly less than the one-to-one relationship that would be exhibited by a well-calibrated individual.

[15]A 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.

[16]Pairwise tests for correlation confirm this expected result in our data.

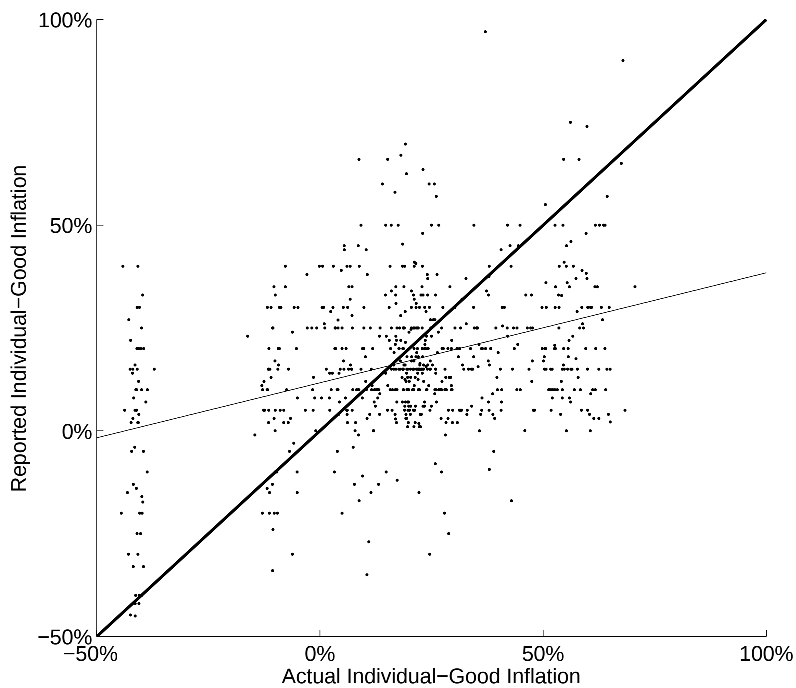

FIGURE VI. Reported versus actual inflation rates across the four individual goods.

It is possible that the correlation bias is driving the treatment differences observed in Figure V. 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 V.

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 Table IV.

	Reported Individual Inflation in Good:			
	A	**B**	**C**	**D**
Reported Basket Inflation	0.214	0.234	0.102	0.450
(Standard Error)	(0.075)	(0.121)	(0.119)	(0.088)
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.

Indeed, subjects' actual weights differ from the correct expenditure-based weights. In all four goods, the actual weight used is biased toward the frequency weight. 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 V.

Finally, we observe that subjects' precision of their perceptions is affected by the frequency of purchase. For each subject, we ask whether their relative ordering of reported inflation rates for goods A and B matches the true ordering of inflation rates for goods A and B. We then regress this binary variable on the absolute difference between the two goods' inflation rates using a probit regression. The estimated coefficient (0.0236) is highly significant (p-value of 0.002), indicating some sensitivity between the rates of these most frequently-purchased goods. When we perform the same regression on the infrequently-purchased goods C and D, however, the estimated coefficient (0.0012) is not significant (p-value of 0.858). Thus, subjects appear unaware of differences between infrequently-purchased goods' inflation rates. These results are consistent with the suggestion that subjects focus attention on frequently-purchased goods and virtually ignore infrequently-purchased goods' prices.

VI DISCUSSION

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. This result is congruent with suggestive survey evidence that upwardly-biased inflation expectations are contemporaneous with higher non-durable goods inflation rates, since non-durables are more frequently purchased than durables.

One implication of the frequency bias is that using current inflation as a proxy for inflation expectations (as in Cagan, 1956, e.g.) is potentially flawed unless the model adjusts for the frequency bias. Even economic research that is interested in inflation *expectations* rather than current perceptions should take these findings into consideration, since the frequency bias in perceptions likely extends to a bias in expectations.

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. Slight over- or under-investments in housing, for example, are unlikely to cause foreclosure or financial distress. 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 (Figure I) 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 experience, but rather from considering prices and inflation analytically, rather than by recalling past shopping experiences.

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 can 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 exactly one good characteristic and a bad one of the same magnitude. 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. An extensive study of these 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.

Batchelor, R. A., 1984. Consumer perceptions and expectations. In: The Measurement and Impact of Inflation and Growth Expectations. Report to the Commission for the EEC, Chapter 5. City University, London.

Bekaert, G., Hodrick, R. J., Marshall, D. A., 2001. Peso problem esxplanations for term structure anomalies. Journal of Monetary Economics 48, 241–270.

Bryan, M. F., Venkatu, G., October 2001. The demographics of inflation opinion surveys. Tech. rep., The Federal Reserve Bank of Cleveland.

Cagan, P. D., 1956. The monetary dynamics of hyperinflation. In: Friedman, M. (Ed.), Studies in the Quantity Theory of Money. University of Chicago Press, Chicago.

Camerer, C. F., March 1992. The rationality of prices and volume in experimental markets. Organizational Behavior and Human Decision Processes 51 (2), 237–272.

Campbell, J. Y., Shiller, R. J., 1991. Yield spreads and interest rate movements: A bird's eye view. Review of Economic Studies 58, 495–514.

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.

Feige, E. L., Pearce, D. K., 1976. Economically rational expectations: Are innovations in the rate of inflation independent of innovations in measures of monetary and fiscal policy? Journal of Political Economy 84, 499–522.

Figlewski, S., Wachtel, P., 1981. the formation of inflationary expectations. Review of Economics and Statistics 63, 1–10.

Fréchette, G. R., 2009. Laboratory experiments: Professionals versus students, new York University working paper.

Friedman, B. M., 1979. Optimal expectations and the extreme information assumptions of rational expectations macromodels. Journal of Monetary Economics 5, 23–42.

Galbraith, J. W., 1988. Modelling expectations formation with measurement errors. Economic Journal 98, 412–428.

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.

Gramlich, E. M., 1983. Models of inflation expectations formation: A comparison of household and economist forecasts. Journal of Money, Credit and Banking 15, 155–173.

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

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.

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.

Kokoski, M., November 2000. Alternative cpi aggregations: Two approaches. Montly Labor Review 123, 31–39.

Krasker, W. S., 1980. The 'peso problem' in testing the efficiency of forward exchange markets. Journal of Monetary Economics 6, 269–276.

Lizondo, J. S., 1983. Foreign exchange futures prices under fixed exchange rates. Journal of International Economics 14, 69–84.

Mankiw, N. G., Reis, R., 2002. Sticky information versus sticky prices: A proposal to replace the new keynsian phillips curve. Quarterly Journal of Economics 117, 1295–1328.

Mankiw, N. G., Reis, R., 2003. Sticky information: A model of monetary non-neutrality and structural slumps. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps. Princeton University Press, Princeton, NJ.

Mishkin, F. S., June 1981. Are market forecasts rational? American Economic Review 71 (3), 295–306.

Rogoff, K., 1980. Essays on expectations and exchange rate volatility. Ph.D. thesis, Massachusetts Institute of Technology.

Schmalensee, R., January 1976. An experimental study of expectation formation. Journal of Money, Credit and Banking 44, 17–41.

Van Duyne, C., 1982. Food prices, expectations, and inflation. American Journal of Agricultural Economics 64, 419–430.

Wachtel, P. A., 1977. Survey measures of expected inflation and their potential usefulness. In: Popkin, J. (Ed.), Analysis of Inflation: 1965–1974. National Bureau of Economic Research, Cambridge, Massachutses, pp. 361–95.

Weber, R. A., 2003. 'learning' with no feedback in a competitive guessing game. Games and Economic Behavior 44, 134–144.

Williams, A. W., February 1987. The formation of price forecasts in experimental markets. Journal of Money, Credit and Banking 19 (1), 1–18.