Indirect Consumer Inflation Expectations: Theory and Evidence*

Ina Hajdini† Edward S. Knotek II‡ John Leer§ Mathieu Pedemonte¶
Robert Rich|| Raphael Schoenle**

July 5, 2023

Abstract

Based on indirect utility theory, we introduce a novel methodology to meausre inflation expectations indirectly. This methodology asks consumers about the change in their income required to buy the same amounts of goods and services one year ahead. Analytically, our methodology possesses smaller ex-post aggregate inflation forecast errors relative to forecasts based on conventional survey questions. We ask this question in a large-scale, high-frequency survey of consumers in the US and 14 countries, and we show that indirect consumer inflation expectations perform well along several empirical dimensions. We then show that individual experiences matter for inflation expectations. For example, age and gender have different effects internationally, while individual inflation and local experiences are generally highly relevant. In an application to gasoline price changes, we find positive effects of gasoline prices on inflation expectations, but in line with the expenditure share of gasoline.

JEL codes: E31, D84, E37, E71

Keywords: Inflation, expectations, surveys, consumers, heterogeneous beliefs

---

*The views expressed here are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. We thank Kenza Benhima, Mark Bils, Olivier Coibion, Nicola Gennaiola, Ricardo Reis, Dmitriy Sergeyev and Michael Weber, as well as seminar participants at the JME-SNB-SCG Conference 2022 "Inflation: Expectations Dynamics" and the 13th ifo Conference on Macroeconomics and Survey Data for helpful comments and suggestions. We thank Grant Rosenberger for excellent research assistance.

†Federal Reserve Bank of Cleveland, ina.hajdini@clev.frb.org

‡Federal Reserve Bank of Cleveland, edward.knotek@clev.frb.org

§Morning Consult, jleer@morningconsult.com

¶Federal Reserve Bank of Cleveland, mathieu.pedemontelavis@clev.frb.org

||Federal Reserve Bank of Cleveland, robert.rich@clev.frb.org

**Brandeis University, schoenle@brandeis.edu

1 Introduction

Inflation expectations play a crucial role in macroeconomic models and policymaking. Conventionally, surveys elicit inflation expectations by directly asking respondents – consumers, professional forecasters, and businesses alike – about "prices in general" or "inflation." However, such questions contain concepts that are abstract and may not be well-defined for some respondents, especially on the household side.

Building on Hajdini et al. (2022b), the present paper proposes an alternative, novel methodology to elicit inflation expectations, a methodology that is also particularly suited to exploiting the breadth of individual data that online survey technologies can provide. Rather than asking directly about an individual expectation for overall inflation, our methodology exploits indirect utility theory and elicits inflation expectations indirectly by asking consumers about the change in income required to buy the same amounts of goods and services a year from the date of the survey. Respondents' answers about their expected required income growth reveal their personally relevant inflation expectations, which when combined with a large sample, can be aggregated up to measure aggregate inflation expectations.

Our methodology is based on indirect consumer utility: all else being equal, a consumer would be able to generate the same amount of utils if her nominal income changes proportionally with the price index of the consumption basket. Extending this line of reasoning to expectations, we show that the expected change in nominal income the consumer finds necessary to purchase the same basket of goods and services over the next 12 months, while preserving her welfare, is equivalent to the anticipated change in the price index of the consumption basket.

The immediate implication of our novel question is that consumers are being asked about their anticipated annual growth rate of the price index of their *individual* consumption basket. This is different from conventional survey questions on *aggregate* inflation expectations. We show that, once responses are aggregated, our novel method of eliciting inflation expectations yields a similar point estimate of aggregate inflation expectations as compared to the conventional question. Indeed, we show that over time, our measure of indirect consumer inflation expectations measure fluctuates closely with the Michigan Survey of Consumers and the Federal Reserve Bank of New York's Survey of Consumer Expectations. Importantly, we show analytically that the advantage of our methodology is that, over time, it should generate a more accurate measure of aggregate inflation expectations, where accuracy is measured by the variance of the ex-post forecast errors. The rationale is as follows: consumers have imperfect information about the evolution of aggregate inflation due to, for example, rational inattention, sticky information, cognitive biases, etc.[1] However, it is reasonable to believe that consumers have better knowledge of the evolution of the prices related to the goods and services they consume rather than other prices in the economy. Therefore, as we show in the paper, aggregating consumers' responses to our question would yield a smaller variance in the ex-post inflation forecast errors relative to aggregating respondents' answers to the standard question. Put differently, we argue that our methodology could offer a more precise estimate of aggregate inflation expectations relative to the direct question about aggregate inflation expectations.

Morning Consult, a data intelligence company, has implemented the indirect consumer inflation expectations question in its proprietary surveys, both in the US and internationally. Data collection started in February 2021 and is ongoing. The US survey poses the question weekly to a cross-sectional, nationally representative sample of approximately 20,000 US adults. In addition to indirect consumer inflation expectations, the survey also records demographic characteristics such as age, ethnicity, gender, educational attainment, income, and region. Several supplementary modules, such as questions about longer time horizons, complement the main survey question. Using the identical question as in the US survey, Morning Consult also surveys respondents in 14 other countries: Australia, Brazil, Canada, China, France, Germany, India, Italy, Japan, Mexico, Russia, South Korea, Spain, and the UK. The international survey also records key demographic characteristics but is run monthly, sampling approximately 1,000 individuals per country. Both the US and international samples are designed to be representative according to the respective population targets. Moreover, when we compute statistics from the US and international samples, our analysis uses survey weights to correct for any remaining sampling inaccuracy.

Our measured indirect consumer inflation expectations perform well along several empirical dimensions, as our empirical analysis shows. First, our series exhibits trends similar to those of established series of inflation expectations, such as the University of Michigan's Survey of Consumers (MSC) and the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE), as well as actual measured consumer price index inflation rates. Our aggregate ICIE measure is of a similar magnitude, and the aggregate time series also matches the trend of these measures of expectations. Second, a one-time survey supplement contains a concurrent question on

[1] See, for example, Mankiw and Reis (2002), Sims (2003, 2010), and Gabaix (2020), among many others, for microfoundations. For evidence on inflation expectations biases due to imperfect information, see Coibion and Gorodnichenko (2015a), Bordalo et al. (2020), Kohlhas and Walther (2021), Angeletos, Huo, and Sastry (2021).

"prices in general" similar to the Michigan question. At the individual level, the answers to that question are highly correlated with our ICIE measure.

As is commonly found in the literature on inflation expectations, there is a meaninful correlation between our ICIE measure and socio-demographic characteristics. We show that individual inflation expectations decrease in income, as in Bruine de Bruin et al. (2010b), and that younger respondents tend to have lower inflation expectations as in Malmendier and Nagel (2016), though the relationship is not linear in age. Education and being a woman are positively associated with inflation expectations as in D'Acunto, Malmendier, and Weber (2021). In contrast with the literature, however, we find a positive association between inflation expectations and numeracy.

A few influential papers have argued that individual experiences may help explain expectations, such as Malmendier and Nagel (2016) in the case of specific cohorts, or D'Acunto, Malmendier, and Weber (2021) in the case of gender mediated by shopping experiences. Recent work situates the role of experiences more broadly into the context of memory and recall of experiences (Bordalo et al. (2022a,b)), for which inflation constitutes one domain of application. Exploiting the rich national and international cross-sectional heterogeneity of experiences in our data, we find that a nuanced, positive picture emerges regarding the effects of age and gender. In terms of age, we find that not all countries exhibit a monotonic relationship between an age interaction and inflation expectations. In general, younger cohorts have lower inflation expectations internationally. However, older individuals do not necessarily hold higher inflation expectations. For example, in China, the oldest cohort tends to have inflation expectations similar to those of the youngest cohort, which are different from to those of the middle-age cohorts. In terms of the relationship with gender, we also observe international differences. The pattern in the US where female respondents hold higher inflation expectations is not common internationally. Russia, Australia, and Canada share this pattern, but in countries like China, Germany, India, Italy, Japan, Mexico, South Korea, and Spain, the pattern is reversed. In the case of Brazil, France, and the UK, there appear to be no statistically significant gender differences. These results suggest that individual experiences captured by age and gender differ internationally, calling for a potentially more nuanced explanation of the gender effect than the role of women as the main shoppers in households (D'Acunto, Malmendier, and Weber (2021)).

In several exercises, we analyze further this role of personal experiences in inflation expectations. We are able to identify this role through our large sample size and the high-frequency variation contained in our data. We find that experiences of inflation are highly correlated with inflation expectations. This result holds when we break down country fixed effects by quantifying the average inflation experienced by respondents, based on Malmendier and Nagel (2016), while taking into account common cohort characteristics such as optimism or numerical bias. This result also holds when we consider detailed local inflation experiences, down to the city level. A clear and statistically significant positive correlation between both inflation expectations and local inflation experience persists, even after removing time fixed effects.

To demonstrate the potential of our geographically dense, high-frequency survey data to identify effects of specific shocks on inflation expectations, we complement our analysis by focusing on gasoline prices, a series that has received a lot of attention especially in light of recent geopolitical events. To do so, we exploit the geographic variation in our data and commuting intensities to show that gasoline price changes are associated with sizable effects on inflation expectations. In addition, our analysis suggests that consumers' inflation expectations overeact to gasoline prices. The effect of gasoline price changes on inflation expectations also tends to be very persistent.

The remainder of this paper is organized as follows. We introduce our methodology of measuring indirect inflation expectations in the next section. Second, we present the survey and discuss its main features and robustness. Third, we discuss the role of experiences for inflation expectations in our data before presenting a specific application in the context of gasoline prices and inflation expectations. We conclude with a short discussion of ICIE expectations at longer horizons.

2 A Novel Methodology to Elicit Inflation Expectations

Surveys of inflation expectations, such as the Federal Reserve Bank of New York's Survey of Consumer Expectations or the University of Michigan's Survey of Consumers, elicit inflation expectations using a direct question about inflation or prices. Our new methodology proposes an indirect approach: rather than the direct question, we ask consumers about the change in income they deem necessary to be able to buy the same amounts of goods and services a year from the date of the survey. Their answer about the expected required income growth, as we show below, then reveals their personally relevant inflation expectations, which can be aggregated up to measure aggregate inflation expectations.

Specifically, our approach uses the following question to elicit indirect consumer inflation expectations:

Next we are asking you to think about changes in prices during the next 12 months in relation to your income. Given your expectations about developments in prices of goods and services during the next 12 *months, how would your income have to change to make you equally well-off relative to your current situation, such that you can buy the same amount of goods and services as today? (For example, if you consider prices will fall by 2% over the next 12 months, you may still be able to buy the same goods and services if your income also decreases by 2%.) To make me equally well off, my income would have to ...*

Respondents then select from three options, filling in the percentages if they select (1) or (3):

1. *Increase by* __%;
2. *Stay about the same*;
3. *Decrease by* __%.

In what follows, we motivate this new methodology by describing two desirable properties it contains. First, we show that our question delivers individually relevant inflation expectations, that is, about the growth rate of prices for respondents' individual consumer baskets. Second, we show that eliciting consumers' individual inflation expectations can yield a more precise measure of aggregate inflation expectations, relative to asking them directly about aggregate inflation expectations.

2.1 Indirect Utility Foundations

Indirect utility theory directly links expected changes in prices – inflation expectations – to changes in expected income required to make consumers equally well off. This link therefore provides a straightforward avenue to elicit inflation expectations from individual survey respondents by eliciting required expected income growth to keep them equally well-off.

To show this result, we assume that consumer i's utility at any period t is given by $u(c_{it})$, where c_{it} is the consumption basket in real terms, and function $u(.)$ is increasing and concave in consumption.[2] Each period, the consumer faces the budget constraint

$$P_{it}c_{it} + S_{it} = Y_{it} + R_{t-1}S_{i,t-1} \tag{1}$$

where P_{it} is the price index of our consumer's *consumption basket*, S_{it} are nominal savings in period t, R_t is the nominal interest rate associated with savings, and Y_{it} is nominal income. For the

[2] We highlight that the same foundations would apply if the utility is a function of labor hours in addition to real consumption. The reason why is that fixing the consumption basket as well as the welfare between any two periods implies that labor hours worked in the two periods are also equalized.

consumer to be equally well-off in period $(t+1)$ as in period t, it must be that

$$u(c_{it}) = u(c_{i,t+1}) \tag{2}$$

Given that $u(.)$ is increasing and concave in consumption, the condition in (2) is satisfied *if and only if* $c_{it} = c_{i,t+1}$. To see what this means for the consumer's expectations, consider the expected budget constraint in period $(t+1)$:

$$\mathbb{E}_{it}(P_{i,t+1}c_{i,t+1}) = \mathbb{E}_t\left(Y_{i,t+1} - (S_{i,t+1} - R_t S_{it})\right) \tag{3}$$

where $(S_{i,t+1} - R_t S_{it})$ denotes net nominal savings. We assume that, in every period, net nominal savings are a fixed share s_i of nominal income. As a result, $(Y_{it} - (S_{it} - R_{t-1}S_{i,t-1})) = (1-s_i)Y_{it}$ and $(Y_{i,t+1} - (S_{i,t+1} - R_t S_{it})) = (1-s_i)Y_{t+1}$. Then, dividing (3) by the constraint in (1), we have

$$\mathbb{E}_{it}\left(\frac{P_{i,t+1}c_{i,t+1}}{P_{it}c_{it}}\right) = \underbrace{\mathbb{E}_{it}\left(\frac{P_{i,t+1}}{P_{it}}\right)}_{\text{expected inflation}} = \mathbb{E}_{it}\left(\frac{(1-s_i)Y_{i,t+1}}{(1-s_i)Y_{it}}\right) = \underbrace{\mathbb{E}_{it}\left(\frac{Y_{i,t+1}}{Y_{it}}\right)}_{\text{expected nominal income growth, } g} \tag{4}$$

where the first equality follows from the assumption that the consumer enjoys the same amount of utils in period $(t+1)$ as in period t. Our analysis shows that there exists a unique income growth rate g that exactly reveals the expected inflation rate of our consumer.

2.2 Indirect Consumer Inflation Expectations versus Conventional Measures

A distinct advantage of our approach – relative to conventional questions that elicit aggregate inflation expectations – may lie in the following: asking consumers about their own, personal inflation expectations, rather than aggregate inflation, and then aggregating across consumers can deliver a more precise estimate of expectations about aggregate inflation, as measured by the variance of the ex-post forecast error.

The reason is that consumers might already face uncertainty around *current* aggregate inflation, which amplifies uncertainty around expectations about future aggregate inflation. By contrast, if a survey of individual "personal" inflation expectations is used to compute an expectation of aggregate inflation, the econometrician can circumvent the layer of uncertainty inherent in the process of aggregate inflation. This principle is quite general and can potentially be applied to other survey domains. We now demonstrate its advantage in the context of inflation expectations.

Suppose the data-generating process for inflation is given by an AR(1) process:

$$\pi_t = \rho \pi_{t-1} + \varepsilon_t,\, \varepsilon_t \sim \mathcal{N}(0,1) \tag{5}$$

Consider a uniform distribution of consumers $i \in [0,1]$ who receive a private signal about aggregate inflation,

$$s_{it} = \pi_t + u_{it},\, u_{it} \sim \mathcal{N}(0, \tau^{-1/2}) \tag{6}$$

where u_{it} is idiosyncratic noise, assumed to be uncorrelated with ε_t for any i and t; and τ denotes the precision of the signal. Similar to Sims (2003, 2010) and Woodford (2003), we think of noise u_{it} representing rational inattention. We then interpret the signal s_{it} as describing the i^{th} consumer's "experience" with inflation. Combining (5) and (6), the law of motion for the signal is given by

$$s_{it} = \rho s_{i,t-1} + (u_{it} - \rho u_{i,t-1}) + \varepsilon_t$$

In this setup, we will compare two measures of aggregate inflation expectations: first, the aggregated measure of inflation expectations if respondents are surveyed about expectations of their own "personal" inflation experience, $\pi^e_{t,personal} = \int_i \mathbb{E}_{it} s_{i,t+1} di$; and second, the aggregated measure of inflation expectations if respondents are surveyed about their expectations of aggregate inflation, $\pi^e_{t,aggregate} = \int_i \mathbb{E}_{it} \pi_{t+1} di$.

Starting with the first, the aggregated measure of inflation expectations if respondents are surveyed about expectations of their own personal inflation experience, we find that

$$\pi^e_{t,personal} = \int_i \mathbb{E}_{it} s_{i,t+1} di = \rho \int_i (s_{it} - u_{it}) di = \rho \pi_t \tag{7}$$

Regarding the second, the aggregated measure of inflation expectations based on respondents' expectations of aggregate inflation is given by

$$\pi^e_{t,aggregate} = \int_i \mathbb{E}_{it} \pi_{t+1} di = \frac{\rho\kappa}{1 - \rho(1-\kappa)\mathbb{L}} \pi_t \tag{8}$$

where $\kappa = 1 - \frac{2}{1+\tau+\rho^2+\sqrt{(1-\rho^2-\tau)^2+4\tau}} \leq 1$ is the Kalman gain, i.e., the weight that the consumer puts on the signal, and $\mathbb{L}$ is a lag operator.[3] Note that κ is increasing in the signal precision τ, such that $\lim_{\tau\to 0} \kappa = 0$ and $\lim_{\tau\to\infty} \kappa = 1$.

[3]See Appendix C for the derivation of $\pi^e_{t,aggregate}$ and Kalman gain κ.

How is the measure of inflation expectations in (8) different from the one in (7)? We first measure the difference between the two:

$$\pi^e_{t,personal} - \pi^e_{t,aggregate} = \frac{\rho(1-\kappa)}{1-\rho(1-\kappa)\mathbb{L}}\varepsilon_t$$

In inflationary (disinflationary) environments with $\varepsilon_t > 0$ ($\varepsilon_t < 0$), asking respondents about aggregate inflation will deliver a smaller (higher) point estimate relative to asking them about their anticipated personal experiences with inflation. On average, however, the two methods will deliver the same point estimate.

We now turn to the uncertainty around the two estimates, measured as the variance in the distance between ex-post inflation and the aggregated inflation expectations:

$$\mathbb{E}\left(\pi_{t+1} - \pi^e_{t,aggregate}\right)^2 = \frac{\rho^2(1-\kappa)^2}{1-\rho^2(1-\kappa)^2} + 1 \geq \mathbb{E}\left(\pi_{t+1} - \pi^e_{t,personal}\right)^2 = 1$$

with the term in red denoting the *excess* uncertainty in inflation expectations due to respondents being asked about aggregate inflation, rather than expectations about their own experiences. Note that the level of excess uncertainty is increasing in the underlying persistence of inflation, ρ, but decreasing in the Kalman gain, κ.

The intuition for why the aggregation of expectations about individual experiences yields a more accurate measure of inflation expectations is that it avoids uncertainty about the *current realizations* of individual experiences. In contrast, when surveyed about aggregate inflation, consumers face uncertainty about current inflation as well as future inflation.

3 Survey Implementation

The question to elicit indirect consumer inflation expectations was introduced by the data intelligence company Morning Consult in its US and international consumer surveys. We describe the related survey structure and survey results on indirect consumer inflation expectations below.

3.1 Survey Description

Morning Consult first introduced the inflation expectations question in the last week of February 2021 in its US and international surveys. Data collection is ongoing. The US survey poses the question to respondents weekly to a representative opt-in sample of approximately 20,000 US adults via Morning Consult's proprietary survey infrastructure. The interviews are conducted online through multiple nationally recognized vendors. The US survey uses a stratified sampling process based on age and gender to reach a broad, nationally representative audience. In addition to indirect consumer inflation expectations, the survey also records demographic characteristics such as age, ethnicity, gender, educational attainment, income, and region. Survey responses come from repeated cross-sections; while some repeat sampling likely occurs given the size of our samples, we do not have a panel aspect to the data per se.

Using the identical question as in the US survey, translated to the local language, Morning Consult also asked the survey question of respondents in 14 other countries. These countries include Australia, Brazil, Canada, China, France, Germany, India, Italy, Japan, Mexico, Russia, South Korea, Spain, and the UK. The international survey employs a similar stratified sampling process and also records key demographic characteristics. The international survey is run monthly, containing approximately 1,000 observations per country. Both the US and international samples are representative according to the respective population targets, as Table 1 shows for the US. To correct for any remaining sampling inaccuracy when computing statistics from the US and international samples, we use survey weights.

Table 1: US Survey Respondent Characteristics

	Survey	US Population		Survey	US Population
Age			**Education**		
18-34	27.17 %	28.99%	<College	62.20%	58.3%
35-44	16.78%	16.56%	Bachelor's degree	23.66%	23.50%
45-64	34.17%	32.21%	Post-grad	16.82%	14.4%
65+	21.88%	22.24%			
			Income		
Gender			**Under 50k**	48.44%	37.8 %
Male	46.17%	48.70%	50k-100k	32.94%	28.6%
Female	53.83%	51.29%	100k+	18.62%	33.6%

Notes: Entries report statistics for the survey respondents and the US population, as obtained from the US Census Bureau. Household income and education (25 years and older): CPS ASEC, 2021; gender: ACS, 2019, which does not report gender other than "male" and "female"; age, race, region: National Population Estimate, 2019.

The US survey has also included supplementary questions on several occasions. First, during the weeks ending on July 31, August 7, and August 14, 2021, the US survey included 3 additional questions to gauge respondent numeracy. Section B.1 in the Appendix lists these questions. The questions were not asked simultaneously and were randomized in a sub-sample of around 6,000 respondents. Higher-income, younger, male, and highly educated respondents tend to answer the numeracy questions correctly, as Table 10 in the Appendix shows. Second, in February 2022, a group of 2,209 respondents were randomly assigned an answer option of "stay the same" instead of "stay about the same" for our main question. Third, in April and June 2022, the US survey elicited long-run inflation expectations using the indirect approach. Horizons included 3 years, 5 years or 10 years into the future. The exact questions are also listed in Section B.1. The analysis exploits the additional information from these supplements below.

3.2 Survey Results

Focusing on the US data, we next provide an overview of our indirect consumer inflation expectations measure. Both the associated aggregate time series and the individual-level data suggest that we capture meaningful information, a point on which we elaborate in the subsequent section.

Figure 1: Indirect Consumer Inflation Expectations and Comparison over Time

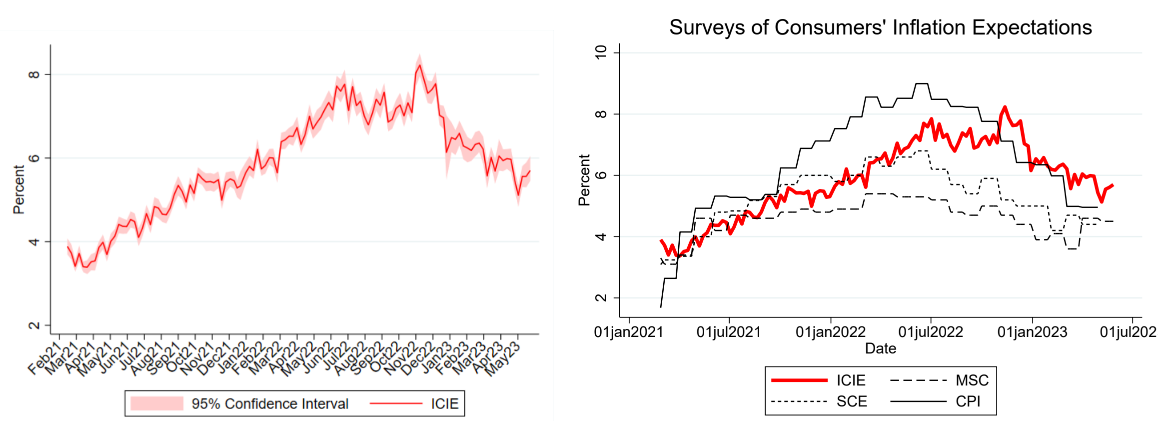

Notes: The left panel shows aggregate indirect consumer inflation expectations with a 95% confidence band. The right panel shows the aggregate indirect consumer inflation expectations (solid red line) along with median expectations from the Federal Reserve Bank of New York's Survey of Consumer Expectations (dashed line), the University of Michigan's Survey of Consumers (long-dashed lined) and the Bureau of Labor Statistics CPI inflation over the past 12 months (solid line).

Figure 1 exhibits the evolution of our measure of aggregate indirect consumer inflation expectations (ICIE) since February 2021. The left panel plots the series together with a 95 percent confidence interval. Aggregate ICIE are computed as the weighted average of all responses in a given week, after trimming the sorted top 10 percent and bottom 10 percent of responses.[4] The confidence interval is constructed as follows: every week, we generate 1000 bootstrapped samples from our incoming data; we compute the weighted average for each sample, sort them, and identify the ones at the bottom and top 2.5 percent. The right panel of the figure shows that our series compares favorably with established series of inflation expectations, such as the University of Michigan's Survey of Consumers (MSC) and the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE), as well as inflation over the prior year as measured by the growth in the consumer price index. Given the evolution of realized inflation and inflation expectations from the MSC and the SCE since March 2021, our ICIE measure appears to be of a similar, relevant quantitative magnitude. Furthermore, our series matches the *trend* of the other two sources of inflation expectations and realized CPI inflation quite well.

Further comparison with established approaches to measuring inflation expectations comes from comparisons of the individual-level data. In a supplement to the Morning Consult survey in related work by Hajdini et al. (2022a), we asked a question similar to the Michigan Survey of Consumers "In the next year, do you think prices in general will increase, decrease or stay the same?" – while also posing our ICIE question. We did in two waves in February 2022 and September 2022. Figure 2 shows the responses to both questions in a "bin-scatter" plot, based on answers of the control group in the experiment of that paper.

[4]The survey uses a stratified sampling process based on age and gender to reach a broad, nationally representative audience. Survey results for the week ending each Saturday are then weighted to match the US adult population according to age, ethnicity, gender, educational attainment, and region; all results shown use these survey weights.

Figure 2: ICIE and “Prices in General”

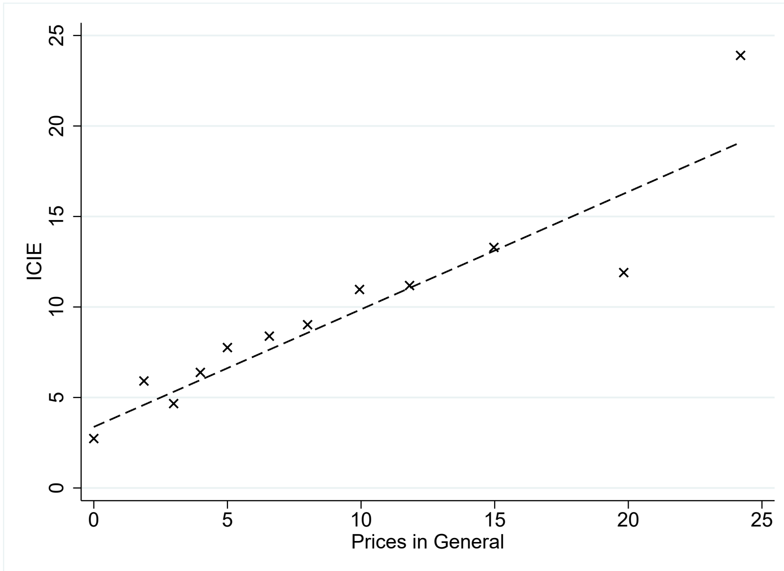

Notes: The figure shows a bin-scatter plot comparing the ICIE (y-axis) and the MSC-type question (x-axis). The dashed line represents a linear projection of the data, with a slope of 0.4.

A high correlation coefficient clearly characterizes the relationship between the ICIE and the measure of “prices in general.” The answers to the “prices in general” question are somewhat less dispersed: answers range from 0 to 25 percent, while they range from 0 to 40 percent for the ICIE, after trimming as described above.

Table 2: ICIE, Demographic Characteristics and Numeracy

		Numeracy 1		Numeracy 2		Numeracy 3	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income (50k-100k)	-0.248***		-0.406*		-0.788***		-0.502**
	(0.018)		(0.243)		(0.269)		(0.243)
Income (Over 100k)	-1.326***		-1.067***		-1.404***		-1.264***
	(0.021)		(0.299)		(0.320)		(0.294)
Age (35-44)	1.238***		1.365***		0.338		0.629*
	(0.024)		(0.333)		(0.388)		(0.336)
Age (45-64)	2.620***		2.030***		0.662**		2.151***
	(0.020)		(0.266)		(0.326)		(0.269)
Age (65 and more)	2.078***		1.572***		0.116		1.510***
	(0.020)		(0.289)		(0.356)		(0.289)
Education (Bachelor's)	0.557***		0.779***		-0.118		0.588**
	(0.019)		(0.265)		(0.277)		(0.263)
Education (Post-grad)	0.463***		0.283		0.130		0.514*
	(0.021)		(0.289)		(0.322)		(0.303)
Gender (Female)	1.233***		0.209		0.109		0.757***
	(0.015)		(0.210)		(0.231)		(0.209)
Numerate		0.655***	0.749***	3.063***	3.072***	1.107***	1.171***
		(0.240)	(0.245)	(0.238)	(0.258)	(0.206)	(0.211)
Constant	4.680***	5.107***	3.777***	3.529***	3.710***	4.762***	3.241***
	(0.015)	(0.119)	(0.263)	(0.192)	(0.335)	(0.148)	(0.270)
Sample	All	Num	Num	Num	Num	Num	Num
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,136,442	6,764	6,764	5,571	5,571	6,841	6,841
R-squared	0.040	0.003	0.014	0.028	0.033	0.005	0.019

Notes: The table shows results of different regressions. Column (1) shows a regression using the full sample of the ICIE on demographic characteristics. For income, the lower-income bracket (less than 50k) is excluded. For age, the group between age 18 and 34 is excluded. For education, the group with less than college is excluded and for gender, males are excluded. Column (2) shows the effect of a dummy equal to one for respondents who answered question (1) correctly. Column (3) adds demographic controls. Columns (4) and (5) are for answering numeracy question 2 right, and columns (6) and (7) are for numeracy question 3. All specifications use time FE and robust standard errors.

At the same time, 35 percent of respondents answer with a zero for the "prices in general" question while 48 percent do so for the ICIE question. Finally, the standard deviation of the "prices in general" question, is 6.48 percent, which is similar to the standard deviation of 7.09 percent for the ICIE. When we exploit the individual-level data and run a regression of the ICIE on the "prices in general" question, we find a statistically significant relationship at the individual level. The estimated regression coefficient is 0.40 with robust standard errors of 0.03 (p=0.000). In addition, the intercept is 3.98 (with a 0.28 standard error). Excluding outliers increases the strength of the relationship: For example, when we exclude answers of the ICIE above 15 percent, the coefficient of the regression is 0.52.

Another comparison for our measure of inflation expectations lies in the meaningful correlation between measured inflation expectations and socio-demographic characteristics as well as measures of numeracy. Table 2, Column (1) presents results from a regression of individual inflation expectations on categories of income, age, education, and gender. We find that inflation expectations decrease with the level of income in a statistically significant way. Individuals in the lowest income group (incomes below $50,000) on average have a difference as large as 1.4 percentage points. This finding is in line with results in Bruine de Bruin et al. (2010a) partaining to the SCE. Moreover, respondents age 35 and above tend to have higher inflation expectations as compared to the ones under age 35. Importantly, this outcome broadly aligns with findings in Malmendier and Nagel (2016). However, we find that the increase in our measure of inflation expectations in the respondents' age is not linear in the age category, and the difference is not always statistically significantly different from 0. The age effect can reach up to 2.47 percentage points. The table also shows that more highly educated respondents, relative to respondents with less than a college education, as well as women, tend to have higher inflation expectations. Finally, looking at Columns (2) through (7), each of our three measures of numeracy is highly statistically significantly and is *positively* associated with higher inflation expectations, by between 0.75 and 3 percentage points.[5]

Figure 7 in the Appendix illustrates the evolution of our measure of inflation expectations across various demographic groups. Generally, the relationships are stable in terms of ordering over time and consistent with the regression findings. Differently from the regression results, consumers in the lower income bracket tend to have unconditionally higher inflation expectations than the rest. In line with the regression results, less educated respondents tend to have lower

[5] Bruine de Bruin et al. (2010b) find a different result, where people with higher financial literacy tend to have lower inflation expectations. That finding was based on a question about aggregate inflation and in a context of lower inflation.

inflation expectations than more highly educated respondents. In addition, women have consistently higher expectations, by approximately 1 percentage point, starting in the middle of our sample, consistent with findings in D'Acunto, Malmendier, and Weber (2021).

Underlying these aggregate features and average relationships of indirect consumer inflation expectations is rich heterogeneity in the distribution of individual inflation expectations. Table 3 reports summary statistics of the distribution of indirect consumer inflation expectations for US respondents up to July 23, 2022. The table reports data for the whole and two trimmed samples: one that trims the top 5 percent and the bottom 5 percent of responses and a second (our favored specification) that drops the top 10 percent and the bottom 10 percent.[6] For the whole sample, we do not report averages because the whole sample contains a few extreme outliers.

Table 3: Distribution of ICIE, with Different Cuts

	All		Trimmed mean (95-5 pct)		Trimmed mean (90-10 pct)	
Average	-	-	8.57	8.15	6.25	5.82
1 percentile	-20	-25	0	0	0	0
5 percentile	0	0	0	0	0	0
10 percentile	0	0	0	0	0	0
25 percentile	0	0	0	0	0	0
50 percentile	0	0	0	0	0	0
75 percentile	12	10	10	10	10	10
90 percentile	49	50	25	25	20	20
95 percentile	75	78	50	50	25	25
99 percentile	200	160	60	60	40	40
Observations	2,441,383	2,441,383	2,441,383	1,952,989	1,952,989	1,291,006
Weights	No	Yes	No	Yes	No	Yes

Notes: The table shows the unweighted and weighted distributions of answers for the ICIE. Columns "All" show the results for the whole sample. Columns under "trimmed mean (x-y pct)" shows the distribution from a sample, where the answers below the y^{th} and above x^{th} percentile of answers were excluded. If more than one respondent is in that group, we randomly remove respondents with similar answers until getting an x or y percentage of the sample.

The distribution of answers has several characteristics that are important to highlight. First, the distribution contains a big concentration of zero answers. The median answer is zero across

[6]For these cuts, we select the percentiles by ordering the answers. As many respondents can provide the same number in that percentile – as is the case – we randomly select some of those answers clustered at the same number and drop them. This methodology leads us to drop exactly the percentile of answers needed.

all samples. This is consistent with other surveys, as described by Andrade, Gautier, and Mengus (2020). In addition, modal workers keep reporting zero nominal wage growth according to the Current Population Survey, as shown in Krohn, Rich, and Tracy (2022). Second, only a small share of respondents provide negative answers. This finding is consistent with other surveys, as described recently in Gorodnichenko and Sergeyev (2021). Third, the table shows that average inflation expectations are influenced by values in the upper tail.[7]

While these characteristics are relatively different from those in other surveys, they are not constant and reflect changes in the economic environment. We find that the distribution evolves systematically as the economy goes from a low-to a high-inflation environment during our sample period: respondents at the beginning of the survey period were exposed to an inflation rate below 2 percent, a number that was similar to pre-pandemic levels. But inflation started to increase rapidly post-pandemic, reaching values above 5 percent in June 2021. After that, the inflation rate continued to increase until reaching a peak of close to 9 percent in June 2022. Figure 3 plots the evolution of the share of responses in different answer brackets.

Figure 3: Distribution of Share of Responses over Time

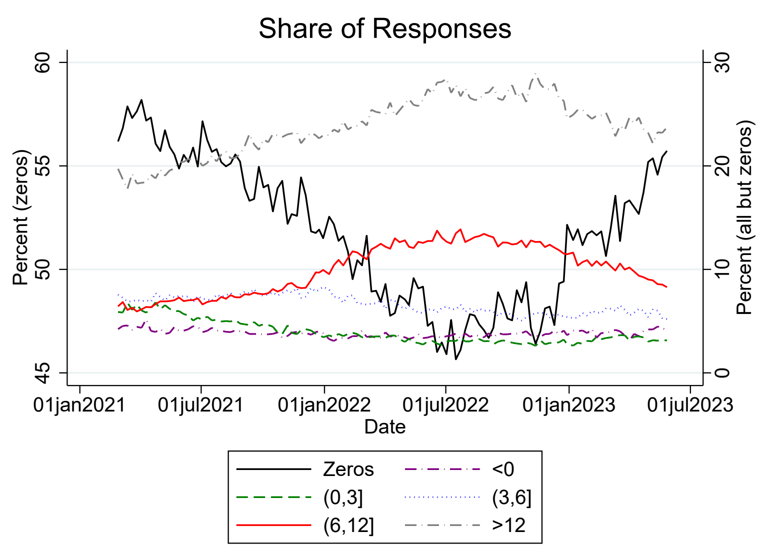

Notes: The figure shows the evolution of the weighted distribution of answers. We group the answers in the 6 groups: negative answers, zero answers, answers above 0 and less than or equal to 3 percent, answers above 3 percent and less than or equal to 6 percent, answers above 6 percent and less than or equal to 12 percent, and answers above 12 percent.

[7]Because the distribution is assimetric, the higher the trim, the lower will be the average, as more positive responses would be dropped. This would decrease the average measure, but for reasonable trims the cycles and trends remain the same

While the share of zeros was relatively stable until July 2021, it has gone down with higher realized inflation readings. Since approximately March 2022, the median response has been non-zero, as indicated by a share of zeros below 50 percent. In addition, the distribution shows a decline in the share of respondents reporting values below 3 percent. A shift of mass to the right tail of the distribution has compensated for this: the shares of respondents reporting values between 6 and 12 percent and above 12 percent have both increased. The latter has come close to 30 percent in recent weeks. While the large share of zero responses is somewhat puzzling and might reflect consumers' biases or rounding behavior - especially after a prolonged low inflationary environment during which inflation may not have been a big consideration for people- its evolution over time is consistent with the recent high inflation readings.

In order to explore the source of the zero responses, we first show that respondents' zero answers systematically correlate with demographic characteristics and whether respondents correctly answer our basic numeracy questions. Table 4 estimates a linear probability model, where the probability of a zero inflation expectations answer is a function of demographics and numeracy. As the table shows, respondents who answer the numerical questions correctly tend to be less likely to report zero answers. This can explain why this group of respondents has higher inflation expectations on average, as our evidence in Table 2 showed. We observe such a pattern for demographic characteristics more generally: whenever a demographic characteristic correlates with higher inflation expectations, it is also associated with a lower chance of a zero answer. This finding is especially true for older respondents who experienced high inflation during the 1970s. The exception to this pattern is the gender demographics: while having higher inflation expectations, female survey respondents also tend to be somewhat more likely to report zero answers.

Table 4: Numeracy, Demographic Characteristics and Zero Answers

	Pr(ICIE=0)					
	Q1		Q2		Q3	
	(1)	(2)	(3)	(4)	(5)	(6)
Numerate	-0.112***	-0.092***	-0.332***	-0.302***	-0.159***	-0.139***
	(0.014)	(0.014)	(0.013)	(0.014)	(0.012)	(0.012)
Income (50k-100k)		-0.013		-0.002		-0.019
		(0.014)		(0.015)		(0.013)
Income (Over 100k)		0.016		0.011		-0.004
		(0.018)		(0.019)		(0.018)
Age (35-44)		-0.080***		0.010		-0.018
		(0.019)		(0.021)		(0.019)
Age (45-64)		-0.146***		-0.040**		-0.158***
		(0.015)		(0.018)		(0.015)
Age (65 and more)		-0.190***		-0.067***		-0.191***
		(0.017)		(0.020)		(0.017)
Education (Bachelor's)		-0.109***		-0.038**		-0.071***
		(0.015)		(0.016)		(0.015)
Education (Post-grad)		-0.112***		-0.077***		-0.096***
		(0.018)		(0.018)		(0.018)
Gender (Female)		0.035***		0.023*		-0.013
		(0.012)		(0.013)		(0.012)
Constant	0.588***	0.719***	0.748***	0.771***	0.636***	0.780***
	(0.007)	(0.015)	(0.010)	(0.018)	(0.008)	(0.016)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,764	6,764	5,571	5,571	6,841	6,841
R-squared	0.011	0.048	0.098	0.105	0.026	0.060

Notes: The table shows the results of a linear regression where the dependent variable is 1 if respondents have an answer equal to zero and 0 if not. We use the same controls as in Table 2.

A second concern in the context of a large share of zero answers may lie in the choice of wording for the "stay about the same" answer option. In order to see how this affects the answer, the February 2022 survey supplement randomly assigned the option "stay the same" to a group of respondents to focus precisely on 0 changes. Given this alternative wording, we did not find a difference in the share of zero responses: 24.26 percent of respondents picked the alternative option of "stay the same" while 24.26 percent picked "stay about the same."

Regression analysis confirms that the choice of wording does not affect the probability of a zero answer. When we regress an indicator variable for a zero answer on an indicator variable for the wording choice, while including a full set of demographic covariates, the estimated coefficient on the wording indicator comes out statistically and economically insignificantly different from zero, either at -0.001 or 0.004 depending on the specification. The wording has no effect on the number of zero answers. Table 5 shows these results.

Table 5: Zero Answers over Different Wording

	(1)	(2)
"Remain about the same" (=1)	-0.001	0.004
	(0.021)	(0.021)
Constant	0.477***	0.475***
	(0.015)	(0.015)
Controls	No	Yes
Observations	2,209	2,209
R-squared	0.000	0.021

Notes: The table shows the results of a linear regression of a variable that is equal to 1 if respondents give a response or zero otherwise. The main independent variable is a dummy that is equal to one if the option was "Remain about the same" and 0 if the option was "Remain the same." Column (1) does not include demographic controls; Column (2) includes demographic controls. Regressions are computed using robust standard errors.

3.3 The Role of Experience

A few influential papers have argued that individual experiences may correlate with inflation expectations, such as Malmendier and Nagel (2016) in the case of specific cohort experiences, or D'Acunto, Malmendier, and Weber (2021) in the case of gender mediated by shopping experiences. Exploiting the rich national and international cross-sectional heterogeneity of experiences in our data, we show that the picture is more complicated. First, using the international data, we estimate the following regression specification:

$$\pi_{i,c,t}^{ICIE} = \alpha_t + \gamma_c + \theta X_i + \beta \gamma_c \times X_i + \varepsilon_{i,c,t} \tag{9}$$

where $\pi_{i,c,t}^{ICIE}$ denotes the ICIE for individual i in country c at time t; α_t denotes a time fixed effect; γ_c denotes a country fixed effect; X_i denotes an individual characteristic, such as an indicator variable for an age cohort or gender, and $\varepsilon_{i,c,t}$ captures the error term. Our main interest lies in the estimates of β, that is, the coefficients on the interactions of age and gender categories with country fixed effects. The coefficients capture the potentially different effects that age cohorts or gender have on inflation expectations across countries. Figure 4 plots the β coefficients for age interactions in the left panel and for gender interactions in the right panel.

A nuanced picture emerges: Age and gender interactions differ across countries. In terms of age, not all countries exhibit a monotonic relationship between the age interaction and inflation expectations. In general, younger cohorts have lower inflation expectations internationally. However, older individuals do not necessarily hold higher inflation expectations. For example, in China, the oldest cohort has inflation expectations similar to those of the younger cohort, and both have lower inflation expectations than the middle-age cohorts. In terms of the gender interaction, we also observe international differences. The pattern in the US where female respondents hold higher inflation expectations is not common internationally. Russia, Australia and Canada share this pattern, but in countries like China, Germany, India, Italy, Japan, Mexico, South Korea, and Spain the pattern is reversed. In the case of Brazil, France, and the UK, there appear to be no statistically significant gender differences.

Figure 4: Demographics and ICIE

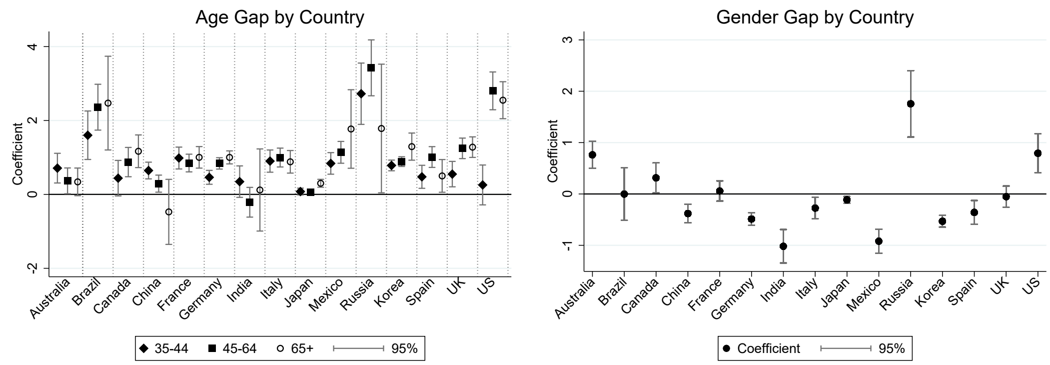

Notes: The figure shows how different demographics respond to the ICIE question in different countries. The left panel plots an age dummy interacted with a country-specific dummy, where the group between ages 18 and 34 is excluded. The right panel shows the results of a regression where a dummy that is equal to one if the respondents are female and zero otherwise is interacted with a country dummy. Regressions include a time and country FE.

Our international data allow us to further highlight the role that experience plays in inflation expectations. To do so, instead of relying on fixed effects as in the previous exercise, we use heterogeneous experience that is specific to a country and to a cohort to see how experience explains these patterns. Based on work by Malmendier and Nagel (2016), we focus on the average level of inflation inflation rate experienced by each cohort. We compute the average and since cohorts where 10 years old for the average year of the bracket. Then, we estimate the following regression specification:

$$\pi_{i,c,t}^{ICIE} = \alpha_t + \gamma_c + \delta Cohort + \beta \bar{\pi}_{i,c,t} + \varepsilon_{i,c,t} \tag{10}$$

where $\pi_{i,c,t}^{ICIE}$ denotes the ICIE for individual i in country c at time t; α_t is a time fixed effect; γ_c denotes a country fixed effect; $Cohort$ is a vector of age cohort dummies; $\bar{\pi}_{i,c,t}$ is the average inflation that an individual i experienced in country c at time t; and $\varepsilon_{i,c,t}$ captures the error term. Critically, this specification contains a cohort fixed effect that is common across countries. This common cohort effect is important because it filters out cohort averages and allows us to see whether high inflation expectations are coming from common cohort-specific characteristics (such as numerical bias, optimism, etc.) or from the historical experience, which is heterogeneous across country-cohort combinations.

Table 6: Cohorts' Inflation Experience and ICIE

	(1)	(1)	(2)	(3)
Mean Inflation Experienced	0.016***	0.016***	0.008***	0.008***
	(0.005)	(0.005)	(0.002)	(0.002)
Country FE	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	Yes
Cohort FE	No	No	Yes	Yes
Cohort-Time FE	No	No	No	Yes
Observations	280,347	280,347	280,347	280,347
R-squared	0.081	0.087	0.090	0.093

Notes: The table shows regressions using the panel of countries where we control for the mean inflation that an individual experienced in a given country (columns (1)-(4)). The experience is determined by the age groups used in the paper. Cohort FE is a common cohort fixed effect across countries. Standard errors are clustered at the country level.

The results highlight that inflation experience matters, as Malmendier and Nagel (2016) have shown: even after taking into account cohort fixed effects, the effect of inflation experiences survives, suggesting that the personal inflation experiences, identified across countries, are associated with inflation expectations for this panel of countries. The estimated coefficients for the level of experienced inflation are relatively small but significantly different from 0. The magnitude depends on the history of inflation of the countries. For example, the Japanese cohort between 35 and 44 years old has experienced an average inflation rate of 0.16 percent, while the oldest cohort of Brazilians has experienced an average inflation rate of 286 percent, given the hyper-inflation experience in Brazil. The standard deviation of the history of inflation is 51 percent, meaning that a one standard deviation increase in the history of inflation in this sample increases the ICIE by 0.42 percentage points. While the magnitude is relatively small, it is statistically significant and confirms the influence that cohort-specific experiences play. While in countries like Japan cohorts have experienced very similar histories of inflation, in countries like Brazil, Russia, and Mexico the experiences are different.

Exploiting the rich geographic variation in our data, we can firm up the result that local inflation environments seem to affect consumer inflation expectations. We now exploit variation across time and region to see if the local environment influences inflation expectations. Thanks to the granularity of the data, we can now add time fixed effects to focus on specific variation across locations. We use the local US as well as international data to make this point. In terms of the US data, the survey gives us information about the zipcode where respondents are located. In addition, the BLS publishes inflation data for 23 commuting zones, some every month and others every two months. Given the zipcodes that belong to each commuting zone, we associate each respondent with local CPI inflation information. In terms of the international data, we obtain CPI inflation from the OECD and associate that information with responses by country. We then take weighted averages in each location – country or commuting zone – and month.

A graphical correlation plot illustrates how geographic variation in price changes is indeed associated with variation in inflation expectations. Figure 5 shows a scatter plot for the US (left panel) and for the international surveys (right data).

Figure 5: Local Inflation and ICIE

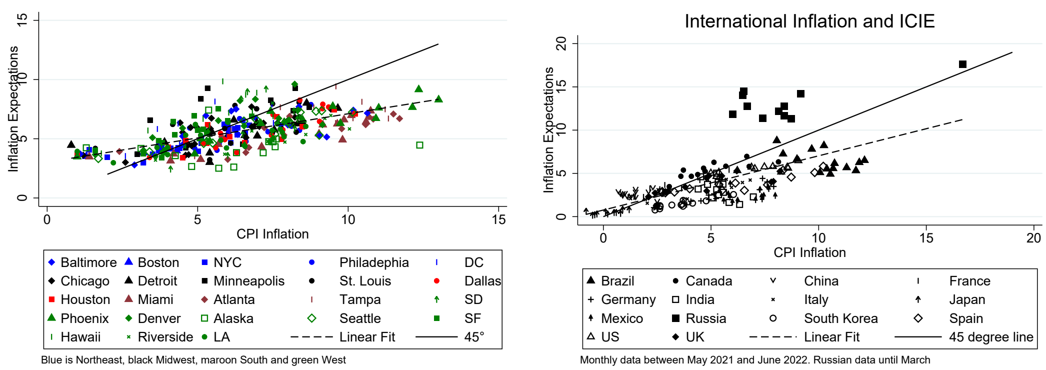

Notes: The figure shows a scatter plot of the monthly average of the ICIE in a given region (y-axis) and the CPI inflation (x-axis). The left panel shows the scatter plot for US cities and the right panel for different countries.

The figure shows a clear positive correlation between local inflation and local inflation expectations. The linear projections in both cases are flatter than the 45-degree line, implying that consumers generally tend to have lower inflation expectations than the actual experienced inflation rate. We formally test this by relying on the micro data and run a regression using geographic and time fixed effects. With this specification, we can differentiate region-specific (commuting zone or country) biases from the common time variation. We estimate the following regression specification:

$$\pi^{ICIE}_{i,c,t} = \alpha_t + \gamma_c + \delta\pi_{c,t} + \beta X_i + \varepsilon_{i,c,t} \tag{11}$$

where $\pi_{i,c,t}^{ICIE}$ denotes the ICIE for individual i in country or city c at time t; α_t denotes a time fixed effect; γ_c denotes a country or city fixed effect; $\pi_{c,t}$ is CPI inflation in country or city c at time t over the last year; X_i is a vector of demographic characteristics that includes age and gender; and $\varepsilon_{i,c,t}$ captures the error term.

Table 7: Regression of ICIE on Local Inflation

	US			International		
	(1)	(2)	(3)	(4)	(5)	(6)
CPI	0.511***	0.135**	0.139**	0.390***	0.426***	0.420***
	(0.043)	(0.054)	(0.061)	(0.045)	(0.093)	(0.108)
Region FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	No	Yes	Yes
Dem FE	No	No	Yes	No	No	Yes
Observations	392,793	392,793	387,390	267,429	267,429	189,866
R-squared	0.017	0.024	0.042	0.110	0.111	0.100

Notes: The table shows a regression of ICIE on local CPI inflation. Columns (1), (2) and (3) use a panel of US cities and columns (4), (5) and (6) use a panel of countries. Dem FE are demographic controls. We use age and gender as demographic controls. Standard errors are clustered at the time and region (country or city) level.

Table 7 shows the results: there continues to be a clear and statistically significant positive relationship between both inflation expectations and local inflation experience, even after including time fixed effects. The coefficient for the US ranges from 0.20 to 0.48, meaning that a 1 percentage point increase in the commuting-zone-specific current inflation rate is associated with an increase in the ICIE between 0.2 to 0.48 percentage points. For the international data, the estimated coefficient is higher: 1 percentage point increase in a country-specific inflation rate is associated with an increase in inflation expectations of 0.59 percentage points. These results suggest that local inflation environments may be affecting consumer inflation expectations.[8] Moreover, after controlling for what happens in the aggregate economy in the US, region-specific price changes influence inflation expectations. The next section explores in more detail the role of local shocks in ICIE using micro data.

[8]However, we do not test the direction of causality, so it is possible that time variation in local consumer inflation expectations is helping to drive some portion of the fluctuations in local inflation outcomes. But the timing – in which inflation is related to expectations of future inflation– suggests the causality goes in the other direction

3.4 Gasoline Price Shocks and Indirect Consumer Inflation Expectations

D'Acunto et al. (2021) show that shopping experiences influence inflation expectations. In that context, many economists have focused on the price of gas, as an important good for consumers and a salient price. For example, Coibion and Gorodnichenko (2015b) show that inflation expectations are highly correlated with changes in gasoline prices. Using local variation, Binder and Makridis (2022) show that gasoline prices influence economic sentiment. In this section, we use the cross-sectional variation in our survey to understand how inflation expectations are affected by gasoline prices. In particular, taking advantage of aggregate variation in gasoline prices and local variation in the use of gasoline based on commuting patterns, we show that consumers' inflation expectations react differently to changes in gasoline prices depending on the intensity of their gasoline use. Moreover, our findings suggest a slight overreaction of inflation expectations to changes in gasoline prices.

We first compute the percentage changes in gasoline prices on a weekly basis compared with one year ago. To measure the share of households that use gasoline more intensively for commuting, we rely on the American Community Survey of 2019 (Ruggles et al. (2019)). In this survey, respondents are asked whether they use their own vehicle to commute. We then compute the share of households that use their own vehicle in a given county, and run the following regression:

$$\pi_{i,t}^{ICIE} = \alpha_{c(i)} + \gamma_t + \beta \log P_{gas,t} \times Comm_{c(i)} + \varepsilon_{i,t} \tag{12}$$

where $\pi_{i,t}^{ICIE}$ denotes the ICIE of respondent i in week t; $\alpha_{c(i)}$ denotes the county c fixed effect where respondent i answered the survey; γ_t is a weekly fixed effect; $\log P_{gas,t}$ is the log of the price of average price of gasoline in the US for week t; $Comm_{c(i)}$ is the share of people who use their own car to commute in county c where respondents i answered the survey, and $\varepsilon_{i,t}$ is a regression error.

This approach has advantages compared to other strategies. In particular, the price of gasoline influences other prices in the economy, so the positive correlation between inflation and gasoline prices should be higher than the exact share of gasoline in the CPI. Therefore, reactions of inflation expectations to changes in gasoline prices could be influenced by the pass-through of gasoline prices to other prices. Including a time fixed effect in (12) helps to isolate the effects of such a pass-through. Moreover, the share of people using their own car for commuting is related to time-invariant local county characteristics, such as public transportation, that are not necessarily related to the pass-through of gasoline prices to other prices and that would be captured by the county fixed effect. For example, Philadelphia, Pennsylvania, is a city with a relatively low use of own cars to commute, compared to Jackson, Mississippi. Differences in car usage generate differences in the direct exposure of consumers to gasoline prices, but not necessarily to the pass-through from gasoline prices to other prices. For example, the price of tradable goods across counties should be similarly affected by changes in the price of an input such as gasoline. While some differences might still be there, as long as they are not correlated with the commuting data, the time fixed effect will control for them. After controlling for time fixed effects, the estimate of β in (12) shows how gasoline prices affect own car commuters' inflation expectations directly, based on common gasoline price changes. Table 8 shows the results.

Table 8: Gasoline Price , Car Use and ICIE

	(1)	(2)	(3)	(4)
$\Delta P_{gas,t}$	7.401***	3.005***		
	(0.489)	(0.395)		
$\log P_{gas,t} \times Comm_{c(i)}$		5.172***	5.734***	4.618***
		(0.596)	(0.674)	(0.998)
County FE	Yes	Yes	Yes	Yes
Time FE	No	No	Yes	Yes
State-Time FE	No	No	No	Yes
Observations	1,860,494	1,193,790	1,193,790	1,193,781
R-squared	0.038	0.024	0.031	0.032

Notes: The table shows the results of regression (12). $\log P_{gas,t}$ is log of the national average price of gasoline and $Comm_{c(i)}$ is the share of households that use their own vehicle to commute according to the ACS 2019. Standard errors are clustered at the county and city levels.

We find that changes in gasoline prices have a larger positive effect on our measure of inflation expectations in areas where the share of commuters is higher. In particular, after adding time fixed effects, a 1 percent increase in gasoline prices is associated with an increase in the ICIE of 0.046 percentage points in an area where everybody uses gasoline to commute. This result is robust to controlling for demographic characteristics.

The next challenge is to understand what the estimate of β means for the reaction of our respondents' inflation expectations to gasoline price changes. In order to do so, we use disaggre-

Figure 6: Share of Gasoline in CPI and Car Use

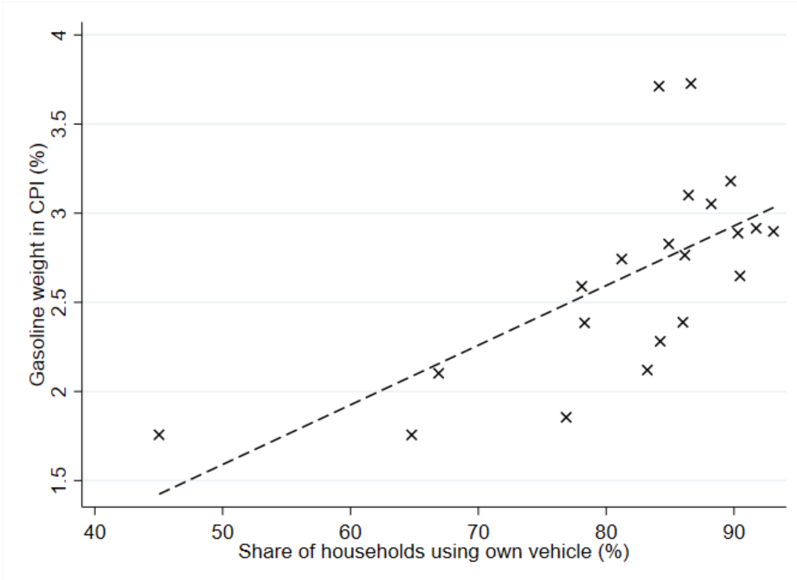

Notes: The figure shows a plot between the expenditure share of gasoline in the CPI (y-axis) and the share of households that use their own vehicle to commute (x-axis). The share of gasoline in the CPI is available for commuting zones, so we add the share of households that use their own vehicle to commute within commuting zones to plot the scatter plot. The dashed line represents a linear projection of the data.

gated data from the BLS on the share of gasoline spending in the CPI for 21 commuting zones. We obtain the share of gasoline measured for 21 commuting zones. We match the counties that correspond to those commuting zones to obtain a measure of expenditure shares on gasoline that varies by commuting zone. This allows us to translate our gasoline usage measure based on the share of households using their own vehicle to commute into direct usage of gasoline based on expenditures. Figure 6 shows a positive and statistically significant correlation between both variables. A 1 percentage point increase in the car use measure increases the gasoline weight in the CPI by 0.033 percentage point. The average share of households that use a car is 82 percent. For the average household, this implies that a 1 percent increase in the price of gasoline increases the ICIE by 0.038 percentage point. By contrast, the gasoline weight in the CPI was 2.7 percent on average across consumers in our sample, implying that a 1 percent increase in gasoline prices would raise CPI inflation.[9]

This outcome suggests that consumers are slightly overreacting to gasoline price changes in terms of inflation expectations. But more importantly, this further implies that consumers perceive changes in gasoline prices to be highly persistent. Our respondents are asked about prices

[9]Our sample is bigger than the CPI shares, so we use the average share of households that use a car to commute and then multiply that average by the slope of the linear relationship from Figure 6 to obtain an average gas CPI weight.

one year ahead, and realized changes in gasoline prices usually are not very persistent. In other words, consumers are biased and perceive changes in current gasoline price to be long-lasting, or at least more persistent than what history would suggest.

This result shows evidence of a perceived high persistence of a gasoline shock, or over-extrapolation. As explained in Angeletos, Huo, and Sastry (2021), this is common in expectations data and it is a way to differentiate among models of expectations formation processes.

4 Longer Horizon Indirect Consumer Inflation Expectations

We have also applied our indirect utility approach to asking respondents about their longer-horizon inflation expectations. Specifically, in April and June 2022, we asked respondents what they thought annual inflation would be, using the indirect approach, on average over the next 3 years, 5 years, and 10 years. The question asked respondents about the 3-year-ahead inflation expectations is:[10]

Now we would like to ask you to think about changes in prices over a longer time in relation to your income. Given your expectations about developments in prices of goods and services during the next 3 years, how would your income have to change each year, on average, to make you equally well-off relative to your current situation, such that you can buy the same amount of goods and services as today?(For example, if you consider prices will fall by 2% on average each year during the next 3 years, you may still be able to buy the same goods and services if your income also decreases by 2% on average per year.) To make me equally well off, on average my income would have to. . .

As before, respondents then select from three options, filling in the percentages if they select (1) or (3):

1. *Increase by __%;*
2. *Stay about the same;*
3. *Decrease by __%.*

To understand how longer-run ICIE are related to the 1-year-ahead ICIE, we run the following regression:

$$\pi^{ICIE}_{i,t,h} = \alpha + \gamma_t + \beta\pi^{ICIE}_{i,t} + \varepsilon_{i,t} \tag{13}$$

where $\pi^{ICIE}_{i,t,h}$ denotes the $h-$year-ahead ICIE of respondent i; α is a constant; and γ_t is a week fixed effect.

[10] The same question is asked for the 5- and 10-year-ahead inflation expectations.

Table 9: 1-Year Ahead and Longer-Run ICIE

	3y ICIE		5y ICIE		10y ICIE	
	(1)	(2)	(3)	(4)	(5)	(6)
1y ICIE	0.835***	0.828***	0.856***	0.850***	0.846***	0.842***
	(0.023)	(0.023)	(0.022)	(0.022)	(0.028)	(0.027)
Constant	1.730***	1.790***	1.687***	1.733***	1.955***	1.989***
	(0.178)	(0.181)	(0.169)	(0.171)	(0.192)	(0.192)
Time FE	No	Yes	No	Yes	No	Yes
Observations	2,202	2,202	2,229	2,229	2,237	2,237
R-squared	0.648	0.651	0.628	0.630	0.578	0.579

Notes: The table shows results of a regression where the dependent variable is the ICIE at 3-year, 5-year or 10-year horizon and the independent variable is the 1-year ICIE. We use robust standard errors.

Table 9 presents the results. We find that short-run inflation expectations are highly correlated with longer-run expectations: consumers who believe that inflation will be higher over the next 1 year generally expect it to remain high in the following years as well. Our finding is consistent with the theory developed in Goldstein and Gorodnichenko (2022) that the persistence between inflation expectations of different horizons increases with the persistence of current inflation. However, the high correlation could also be due to high *perceived* persistence of inflation expectations. Other work, using different surveys, also finds a high correlation between short- and long-run inflation expectations (e.g., Coibion et al. (2020)).

5 Conclusion

Based on indirect utility theory, we have introduced a novel methodology of measuring inflation expectations indirectly. This methodology starts at the individual level, asking consumers about the change in income required to buy the same amounts of goods and services one year ahead. Analytically, our methodology possesses smaller ex-post aggregate inflation forecast errors relative to forecasts based on conventional survey questions. Using data from a large-scale, high-frequency survey implementation in the US and 14 countries, we show that indirect consumer inflation expectations share some similarities with other measures of inflation expectations, and some notable differences.

Exploiting the geographically detailed, high-frequency variation in the data, we then show that individual experiences matter for inflation expectations, in a nuanced way. For example, age and gender have different effects internationally, while individual inflation and local experiences are generally highly relevant. In an application to gasoline price changes, our analysis has identified large effects of experienced gasoline price changes on inflation expectations, characterized by both overreaction and persistence.

Finally, we find that longer-term inflation expectations measured by the ICIE are strongly correlated with the 1-year-ahead ICIE, suggesting that these longer-term expectations were elevated at the end of our sample.

Appendix

A Additional Empirical Results

Table 10: Numeracy and Demographic Characteristics

	Q1	Q2	Q3
Income (50k-100k)	0.046***	0.027**	0.042***
	(0.011)	(0.012)	(0.012)
Income (Over 100k)	0.043***	0.007	0.085***
	(0.015)	(0.016)	(0.017)
Age (35-44)	-0.066***	0.076***	-0.027
	(0.015)	(0.020)	(0.017)
Age (45-64)	-0.069***	0.335***	-0.014
	(0.012)	(0.015)	(0.014)
Age (65 and more)	-0.062***	0.414***	0.028*
	(0.013)	(0.016)	(0.016)
Education (Bachelor's)	0.118***	0.074***	0.112***
	(0.012)	(0.013)	(0.014)
Education (Post-grad)	0.139***	0.119***	0.139***
	(0.015)	(0.014)	(0.016)
Gender (Female)	-0.057***	-0.041***	-0.081***
	(0.009)	(0.011)	(0.011)
Constant	0.238***	0.389***	0.463***
	(0.012)	(0.015)	(0.014)
Observations	8,382	6,802	8,537
R-squared	0.040	0.144	0.036

Figure 7: ICIE across Demographics

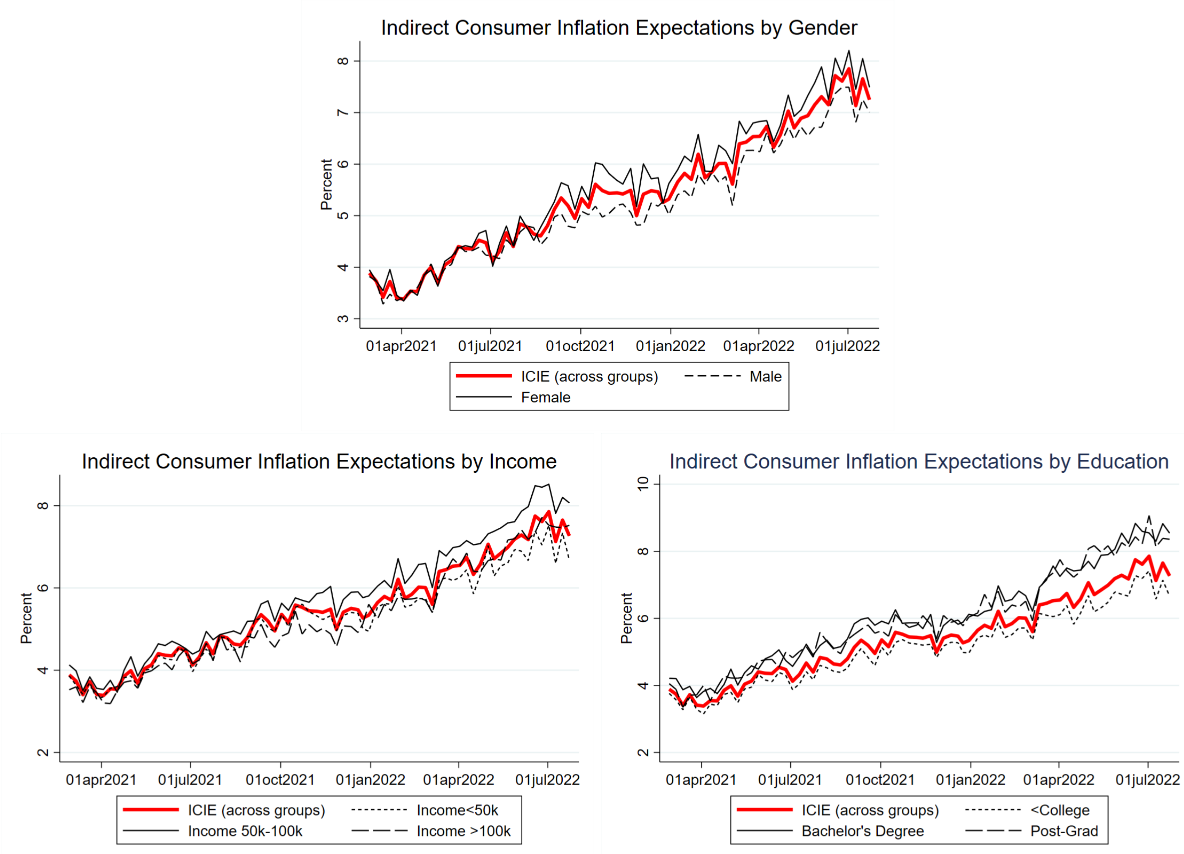

B Survey Instrument

B.1 Supplementary Questions

B.1.1 Numeracy Questions

The survey included 3 questions related to respondent numeracy:

Q1 Imagine there are white and black balls in a ballot box. You draw a ball 70 times. 56 times, you have drawn a white ball, 14 times a black ball. Given this record, what would you say is the probability of drawing a black ball the next time? The probability is ___ percent. (N=8,382)

Q2 Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, how much would you be able to buy with the money in this account? (N=6,802)

Q3 Imagine that we roll a fair, six-sided dice 1,000 times. Out of 1,000 rolls, how many times do you think the dice would come up even (2, 4, or 6)? (N= 8,537)

C Derivation of $\pi^e_{t,aggregate}$

Consider the law of motion for the signal and inflation:

$$\pi_t = \rho\pi_{t-1} + \varepsilon_t \tag{C.1}$$

$$s_{it} = \pi_t + u_{it} \tag{C.2}$$

Given (C.1) and (C.2), we have that

$$\mathbb{E}_{it}\pi_{t+1} = \rho\mathbb{E}_{it}\pi_t$$

where

$$\mathbb{E}_{it}\pi_t = \rho(1-\kappa)\mathbb{E}_{i,t-1}\pi_{t-1} + \kappa s_{it} \tag{C.3}$$

is the weighted average between the signal s_{it} and the prior $\rho\mathbb{E}_{i,t-1}\pi_{t-1}$ with Kalman gain κ (to be derived below). From (C.3), it follows that $\mathbb{E}_{it}\pi_t = \frac{\kappa}{1-\rho(1-\kappa)\mathbb{L}}s_{it}$, and therefore that

$$\mathbb{E}_{it}\pi_{t+1} = \frac{\rho\kappa}{1-\rho(1-\kappa)\mathbb{L}}s_{it}$$

To compute the Kalman gain, let (5) be the state equation and (C.2) be the measurement equation. Moreover, let $\psi = \mathbb{E}\left((\pi_t - \mathbb{E}_{i,t-1}\pi_t)^2\right)$. Then, the Kalman filter grants the following relationship

$$\psi = \mathbb{E}\left((\pi_t - \mathbb{E}_{i,t-1}\pi_t)^2\right) = \rho^2\psi(1-(\psi+\tau^{-1})) + 1$$

Solving the quadratic equation above for ψ (such that $\psi > 0$), we have that

$$\psi = \frac{\rho^2 + \tau - 1 + \sqrt{(1-\rho^2-\tau)^2 + 4\tau}}{2\tau}$$

Finally, the Kalman gain is given by

$$\kappa = \frac{\psi\tau}{\psi\tau+1} = \frac{\rho^2+\tau-1+\sqrt{(1-\rho^2-\tau)^2+4\tau}}{\rho^2+\tau+1+\sqrt{(1-\rho^2-\tau)^2+4\tau}} = 1 - \frac{2}{\rho^2+\tau+1+\sqrt{(1-\rho^2-\tau)^2+4\tau}}$$

References

Andrade, Philippe, Erwan Gautier, and Eric Mengus. 2020. "What Matters in Households' Inflation Expectations?" *Federal Reserve Bank of Boston Research Department Working Papers No. 20-18.* DOI https://doi.org/10.29412/res.wp.2020.18.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry. 2021. "Imperfect Macroeconomic Expectations: Evidence and Theory." *NBER Macroeconomics Annual* 35:1–86. DOI https://doi.org/10.3386/w27308.

Binder, Carola and Christos Makridis. 2022. "Stuck in the seventies: Gas prices and consumer sentiment." *Review of Economics and Statistics* 104 (2):293–305. DOI https://doi.org/10.1162/rest_a_00944.

Bordalo, Pedro, Giovanni Burro, Katherine B Coffman, Nicola Gennaioli, and Andrei Shleifer. 2022a. "Imagining the Future: Memory, Simulation and Beliefs about Covid." Working Paper 30353, National Bureau of Economic Research. DOI https://doi.org/10.3386/w30353.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer. 2022b. "Memory and Probability." *The Quarterly Journal of Economics* DOI https://doi.org/10.1093/qje/qjac031. Qjac031.

Bordalo, Pedro, Nicola Gennaioli, Ma Yueran, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review* 110 (9):2748–82. DOI https://doi.org/10.1257/aer.20181219.

Bruine de Bruin, Wändi, Simon Potter, Robert W Rich, Giorgio Topa, and Wilbert Van der Klaauw. 2010a. "Improving survey measures of household inflation expectations." *Current Issues in Economics and Finance* 16 (7).

Bruine de Bruin, Wändi, Wilbert Van der Klaauw, Julie S Downs, Baruch Fischhoff, Giorgio Topa, and Olivier Armantier. 2010b. "Expectations of inflation: The role of demographic variables, expectation formation, and financial literacy." *Journal of Consumer Affairs* 44 (2):381–402. DOI https://doi.org/10.1111/j.1745-6606.2010.01174.x.

Coibion, Olivier and Yuriy Gorodnichenko. 2015a. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8):2644–2678. DOI https://doi.org/10.1257/aer.20110306.

———. 2015b. "Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation." *American Economic Journal: Macroeconomics* 7 (1):197–232. DOI https://doi.org/10.1257/mac.20130306.

Coibion, Olivier, Yuriy Gorodnichenko, Edward S Knotek II, and Raphael Schoenle. 2020. "Average inflation targeting and household expectations." Tech. rep., National Bureau of Economic Research. DOI https://doi.org/10.3386/w27836.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber. 2021. "Exposure to grocery prices and inflation expectations." *Journal of Political Economy* 129 (5):1615–1639. DOI https://doi.org/10.1086/713192.

D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber. 2021. "Gender roles produce divergent economic expectations." *Proceedings of the National Academy of Sciences* 118 (21):e2008534118. DOI https://doi.org/10.1073/pnas.2008534118.

Gabaix, Xavier. 2020. "A Behavioral New Keynesian Model." *American Economic Review* 110 (8):2271–2327. DOI https://doi.org/10.1257/aer.20162005.

Goldstein, Nathan and Yuriy Gorodnichenko. 2022. "Expectations Formation and Forward Information." Working Paper 29711, National Bureau of Economic Research. DOI https://doi.org/10.3386/w29711.

Gorodnichenko, Yuriy and Dmitriy Sergeyev. 2021. "Zero Lower Bound on Inflation Expectations." Working Paper 29496, National Bureau of Economic Research. DOI https://doi.org/10.3386/w29496.

Hajdini, Ina, Edward S Knotek, John Leer, Mathieu Pedemonte, Robert W Rich, and Raphael Schoenle. 2022a. "Low Passthrough from Inflation Expectations to Income Growth Expectations: Why People Dislike Inflation." DOI https://doi.org/10.26509/frbc-wp-202221.

Hajdini, Ina, Edward S Knotek II, Mathieu Pedemonte, Robert Rich, John Leer, and Raphael Schoenle. 2022b. "Indirect Consumer Inflation Expectations." *Economic Commentary* (2022-03). DOI https://doi.org/10.26509/frbc-ec-202203.

Kohlhas, Alexandre N and Ansgar Walther. 2021. "Asymmetric attention." *American Economic Review* 111 (9):2879–2925. DOI https://doi.org/10.1257/aer.20191432.

Krohn, Mason, Robert W Rich, and Joseph Tracy. 2022. "More Workers Find Their Wages Falling Even Further Behind Inflation." DOI https://www.dallasfed.org/research/economics/2022/1004.

Malmendier, Ulrike and Stefan Nagel. 2016. "Learning from inflation experiences." *The Quarterly Journal of Economics* 131 (1):53–87. DOI https://doi.org/10.1093/qje/qjv037.

Mankiw, N Gregory and Ricardo Reis. 2002. "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve." *The Quarterly Journal of Economics* 117 (4):1295–1328. DOI https://doi.org/10.1162/003355302320935034.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek. 2019. "Integrated Public Use Microdata Series: Version 12.0 [dataset]." *Minneapolis: University of Minnesota* DOI https://doi.org/10.18128/D010.V12.0.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3):665–690. DOI https://doi.org/10.1016/S0304-3932(03)00029-1.

———. 2010. "Rational Inattention and Monetary Economics." *Handbook of Monetary Economics* 3:155–181. DOI https://doi.org/10.1016/B978-0-444-53238-1.00004-1.

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press.