Learning-through-Survey in Inflation Expectations

By Gwangmin Kim and Carola Binder*

When surveys rely on repeat participants, this raises the possibility that survey participation may affect future responses, perhaps by prompting information acquisition between survey waves. We show that these "learning-through-survey" effects are large for household inflation expectations. Repeat survey participants generally have lower inflation expectations and uncertainty, particularly if their initial uncertainty was high. Consequently, repeat participants may be more informed about or attentive to inflation. This has important implications: for example, inflation expectations of new participants are more influenced by oil prices, and estimates of the elasticity of intertemporal substitution are lower for new participants.
JEL: D83, D84, E31

Inflation expectations are believed to play a central role in economic dynamics. Federal Reserve Chair Jerome Powell testified to Congress in February 2019 that "Inflation expectations are the most important driver in actual inflation" (Powell, 2019). In addition, survey-based inflation expectation measures are increasingly used in economic research in various ways: for estimating the intertemporal elasticity of substitution (Crump et al., 2015), studying inflation expectations of firms (Coibion, Gorodnichenko and Kumar, 2018), and estimating expectations-augmented Phillips curve (Coibion, Gorodnichenko and Ulate, 2019).

Therefore, accurately measuring inflation expectations is crucial for monetary policymaking and economic research. For this reason, the Federal Reserve Bank of New York (FRBNY) began conducting the Survey of Consumer Expectations (SCE) monthly in 2013. Other central banks, like the European Central Bank (ECB) and the Bank of Canada, are also introducing new household surveys. For example, the ECB launched the Consumer Expectations Survey in January 2020.[1] Globally, dozens of countries run household inflation expectation surveys on a regular basis.[2]

* Kim: University of Texas at Austin, 2225 Speedway BRB. 4.118, Austin, TX 78712, USA (e-mail: gmkim@utexas.edu); Binder: Haverford College, 370 Lancaster Avenue, Haverford, PA 19041-1392 (e-mail: cbinder1@haverford.edu). This paper is a part of Ph.D. dissertation chapters of Gwangmin Kim at UT Austin. We are grateful to Olivier Coibion, Stefano Eusepi, Saroj Bhattarai, and seminar participants at the University of Texas at Austin for helpful comments and suggestions.

[1]The ECB Consumer Expectations Survey is in pilot phase in six countries: Belgium, France, Germany, Italy, the Netherlands, and Spain. More countries may be added if the pilot proves successful. `https://www.ecb.europa.eu/stats/ecb_surveys/consumer_exp_survey/html/index.en.html` provides further details about this new survey.

[2]See Arioli et al. (2017) and Appendix Table 1 of Coibion, Gorodnichenko and Ulate (2019) for a list of countries running inflation expectation surveys targeting households. Norway, which is not on their list, also has run an inflation expectation survey since 2002.

FRBNY SCE respondents can participate in the survey for up to twelve months in a row. A long panel dimension is usually thought to be a desirable feature for a survey, since measuring the same person over time allows researchers to control for unobservable individual-specific characteristics. However, reliance on repeat participants—the SCE includes about 150 new participants out of 1300 in each wave—could pose problems if the act of participating in the survey affects the subsequent responses of these participants. These so-called *learning-through-survey* or *panel conditioning* effects are small in some surveys.[3]

However, we show that this is decidedly *not* the case in surveys of household inflation expectations. After being asked about their inflation expectations, individuals significantly (and predictably) revise their expectations in subsequent surveys. For example, after participating twelve consecutive times in the SCE, respondents end up with a 2.6 percentage point lower inflation forecast and 34% lower inflation uncertainty on average than in the first interview, with most of the decline happening in the first two months of participation. Results are similar for longer-run inflation expectations. These effects are so large that repeat participants can no longer be considered representative of the general population.

The extensive panel component of the SCE makes it an ideal setting for studying panel conditioning effects. We note that Armantier et al. (2017) briefly examine panel conditioning effects in the SCE by comparing the median absolute change in the density mean of respondents' inflation expectations across different tenure groups.[4] They find that after the first month of participation, the density mean of a respondent's inflation forecast remains relatively stable.

We provide evidence, however, that the learning-through-survey effects are larger and more economically meaningful than previously recognized. Our approach is to use panel regressions with time and respondent fixed effects and tenure dummy variables to detect conditioning effects that may occur over multiple survey waves, without imposing parametric assumptions on how effects depend on tenure. This is a novel methodological contribution to a relatively large literature on panel conditioning.[5] The coefficients on the tenure dummies provide non-parametric estimates of how inflation expectations change with survey tenure. Since consumers generally overestimate future inflation,[6] they lower their

[3]For example, Halpern-Manners, Warren and Torche (2017) find that only about 12% of the selected core items of the General Social Survey display panel conditioning effects at a 5% significance level. They report that the responses of different survey cohorts do not appear to differ in predictable or meaningful ways in most cases.

[4]Throughout this paper, "tenure" refers to the total number of past survey experience of respondents, including the current survey wave. For example, a SCE respondent surveyed each month starting in January will have a tenure of 3 in March.

[5]Previous studies conduct t-tests for the difference in mean responses between two cohorts of respondents who first entered the survey sample at two different dates. For example, Halpern-Manners, Warren and Torche (2017) compare responses on the 2008 General Social Survey for respondents who took the survey in 2006 and 2008 versus respondents who took the survey in 2008 and 2010. We pool information from *all* dates of the SCE rather than from a single survey date, and observe how responses change not only from a respondent's first to second round of participation, but also from her second to third round of participation and so on."

[6]Although Consumer Price Index inflation has been recently low and stable—at around 1.5% from

forecasts and make smaller forecast errors as they remain in the survey and acquire more information.[7]

Furthermore, we characterize which individuals are most sensitive to having their beliefs change through participation in the survey. Respondents who report higher uncertainty about inflation at the time of their first survey tend to make larger revisions to their inflation expectations in subsequent surveys. In addition, more educated and higher-income individuals and retirees, who are generally more informed about inflation prior to the survey, display significantly smaller learning effects throughout the survey waves.

These heterogeneity results are consistent with models that emphasize the endogenous nature of information rigidities. Under the rational inattention model, economic agents have a limited cognitive ability to process information, and must choose how to allocate attention.[8] Though the inflation rate is an important aggregate variable, it may be optimal for households to pay greater attention to tracking other variables, like their own income, that are more relevant to their consumption decisions (Carroll et al., 2020). Indeed, Karahan, Mihaljevich and Pilossoph (2017) find that the income expectations of consumers tend to be accurate. Similarly, when firm-specific conditions are more important than aggregate conditions, firms pay more attention to idiosyncratic variables and devote few resources to collecting and processing information about inflation (Mackowiak and Wiederholt, 2009). Consumers' relatively greater attentiveness to their own income than to inflation is consistent with our result that the learning effect is smaller for personal earnings and household income growth expectations than for inflation expectations.

As we demonstrate, our results should be kept in mind by users and developers of new household surveys, as they affect empirical estimates and interpretations using the survey data in some contexts. We illustrate this point using two application cases: the oil price collapse in 2014 and the estimation of elasticity of intertemporal substitution by Crump et al. (2015). Household inflation expectations are known to be sensitive to gas prices (Coibion and Gorodnichenko, 2015*a*). However, this stylized fact tends to be significantly weaker for repeat participants of the SCE. During a period of sharp decline in oil and gas prices, we find that inflation expectations of new participants are more influenced by gas prices when compared to repeat survey participants. Also, we show that estimates of the elasticity of intertemporal substitution following the methodology of Crump et al. (2015) are lower for new participants. We suggest that users of survey microdata should check whether their estimates are robust to using subsamples of shorter-

2013 to 2018—the inflation expectation of consumers was consistently above 2.5% during the same period.

[7]One difference between our analysis and that of Armantier et al. (2017) is that we identify the average change, not the average absolute change, in inflation forecasts. This is appropriate in this context because forecasts have positive bias, so errors have non-zero mean.

[8]Related models allow for different types of information rigidities such as infrequent updating (Mankiw and Reis, 2002; Reis, 2006), signal-extraction problems (Sims, 2003), or model complexity restrictions (Gabaix, 2014).

tenured and longer-tenured respondents when they use panel data.

Finally, we provide evidence of learning-through-survey effects in other contexts using the Michigan Survey of Consumers (MSC) and a new firm survey to show that learning-through-survey effects are not confined to a specific period or only to household surveys. Learning-through-survey effects tend to be smaller when there is a longer period between baseline and follow-up surveys, as in the MSC. This difference in size of learning effects is consistent with recent evidence from randomized information treatments that finds providing information about inflation to households has large contemporaneous effects on their expectations but that these effects fade very rapidly (Coibion, Gorodnichenko and Weber, 2020).

The paper is structured as follows. Section I provides information about the dataset and presents our estimates the learning-through-survey effects in the FRBNY SCE. Section II provides implications of the learning-through-survey effect for interpretation of survey data using two application cases: the oil price collapse in 2014 and estimation of elasticity of intertemporal substitution. Section III documents the effect in two other survey datasets. Section IV discusses an alternative explanation for the tenure effects and implications for future research.

I. Learning Effects in the Survey of Consumer Expectations

A. Data

The Federal Reserve Bank of New York's (FRBNY) Survey of Consumer Expectations (SCE) is an online survey that began 2013. The SCE is monthly and nationally-representative with a rotating panel structure, tracking each respondent up to 12 times consecutively. Each month, the SCE has a sample size of approximately 1,300, and the number of new participants is about 150.

In addition to inflation point forecasts, the FRBNY elicits the respondent's histogram or density forecasts for inflation by asking the respondent to assign probabilities that future inflation will fall into various bins, summing to 100%. Hence, for inflation uncertainty, we use the interquartile range (IQR) estimated from each individual's probabilistic forecast.[9] The exact phrasing of survey questions is available in the appendix.

B. Identification of Tenure Effects

We begin by documenting the presence of tenure effects in the FRBNY SCE data, using linear panel fixed effects regressions of the form:

$$y_{its} = \sum_{s=2}^{12} \beta_s \tau_s + \alpha_i + \gamma_t + \varepsilon_{it}, \tag{1}$$

[9]The FRBNY provides estimates of the mean, median, and IQR of each density forecast. These estimates are obtained by fitting parametric (beta) distributions to the density forecasts. See FRBNY SCE documentation for details.

where the dependent variable y_{its} is the inflation expectation or inflation uncertainty of respondent i with survey experience (or tenure) s at time t, τ_s is an indicator variable for tenure s, α_i and γ_t are individual- and time- fixed effects to control for unobserved heterogeneity, and ε_{it} is an error term. The regression coefficients on the tenure dummies, $\{\beta_s\}_{s=2}^{12}$, measure the average learning-through-survey effects on the dependent variable. To make the regression coefficients more robust to outliers, we remove the top and bottom 5% of each dependent variable for each tenure group and period.[10]

One identification issue is that sample selection may occur due to panel attrition. For example, more educated and higher income respondents tend to stay in a survey for more waves, and attrition may also depend on unobservable characteristics. To prevent confounding panel conditioning effects with attrition effects, we restrict our sample to consist of "non-attriters," or respondents who eventually participate in the survey for the maximum number of times, following Halpern-Manners, Warren and Torche (2017).

Panel A of Figure 1 shows that the estimated average learning-through-survey effect is large and statistically significant for both one-year-ahead and three-year-ahead inflation expectations. Respondents revise their one-year-ahead inflation expectations downward by 1.2 percentage points immediately after the first interview. Respondents with tenure 12 have expectations that are 2.6 percentage points lower than those of new respondents. Three-year-ahead inflation expectations display similar, if slightly smaller, tenure effects. Given this similarity, we primarily focus on one-year-ahead inflation expectations for the remainder of the paper. A full regression table is in Appendix Table A1.

While the results in Panel A correspond to respondents' point forecasts of inflation, the FRBNY frequently reports on the density mean forecasts. Panel conditioning effects for the density mean forecasts are reported in Columns 3 and 4 of Table A1. Though statistically significant, they are smaller in magnitude than the effects for the point forecasts: repeat respondents have density means that are around half a percentage point lower than those of new respondents. These smaller tenure effects may be related to the guidance provided by the bin intervals when density forecasts are solicited.

Recall that after the respondent provides a point forecast, her density forecast is solicited, with upper and lower bins corresponding to inflation above 12% and deflation below -12%. The bins near zero are narrower than those above 4% or below -4%. From these bins, the respondent may infer that most of the probability should be placed in $[-4\%, 4\%]$, or at least in $[-12\%, 12\%]$. (Almost a third of first-time respondents provide point forecasts outside of $[-12\%, 12\%]$ and over half provide point forecasts outside of $[-4\%, 4\%]$.) Any learning that occurs from the bin endpoints constitutes "learning-through-survey" in the very first round of the survey, resulting in smaller observed tenure effects in subsequent rounds.

[10]We obtain qualitatively and quantitatively similar results for different thresholds. Appendix Figure A1 reproduces the results in Figure 1 for lower and higher thresholds.

We can see some evidence that this occurs, because density forecast means are notably lower than point forecasts. Moreover, new respondents have density forecast means that are significantly further from their point forecasts compared to respondents of higher tenure.

Even though the tenure effects are smaller in magnitude for density mean forecasts than for point forecasts, sizeable tenure effects appear for other features of the density forecasts, such as the interquartile range (IQR). The IQR decreases by about 0.7 percentage points after the first round of the survey, as shown in Panel B of Figure 1.

As consumers' uncertainty declines with tenure, their forecast errors also decline, as Appendix Table A2 shows. Since consumer inflation expectations are typically biased upward, the downward revisions in expectations improve forecast accuracy. The mean absolute forecast error for respondents of tenure 2 is 2.0 percentage points lower, and for respondents of tenure 12 is 4.3 percentage points lower, than that of new participants. The same table shows that higher-tenure respondents make much less frequent forecast revisions.[11] Dräger and Lamla (2017) similarly document that the probability of updating inflation expectations increases when individuals had higher forecast errors in the past. That is, repeat participants achieve lower forecast errors and are less likely to replace their current forecasts, saving cognitive effort in processing information.

Panels C and D of Figure 1 show time series plots of the mean inflation point forecasts and density IQR for new respondents compared to all respondents. Here we include both attriters and non-attriters, since policymakers typically monitor the aggregate time series without imposing a non-attrition restriction. The time series corresponding to new respondents are higher and more volatile. These graphs also reveal that tenure effects vary over time. Notably, in March 2020, at the start of the Covid-19 pandemic, new respondents had a mean inflation forecast of 11.5%, compared to 3.6% for the average respondent of any tenure and 2.4% for respondents with tenure greater than one. The relatively less-informed new respondents may have assumed the pandemic would be inflationary (see Binder (2020)).

Our time sample also includes a disinflationary episode in 2015. We estimate the tenure effects for 2015 only, and find that the effects are larger in magnitude during this episode (see Column 7 of Table A1). This is consistent with the learning-through-survey hypothesis, if repeat respondents are more aware of declining inflation than are new respondents. We also construct a time series of the number of articles containing the word "inflation" in the New York Times each month. This series is positively correlated with the inflation expectations of repeat respondents (with correlation coefficient of 0.55), but much less correlated with the inflation expectations of new respondents (with correlation coefficient of 0.38). This could indicate that repeat respondents, primed by their earlier survey

[11]Note that the frequency of forecast revisions is often used as a proxy for the information rigidity parameter in sticky information models (Coibion and Gorodnichenko, 2015*a*; Binder, 2017*a*)

participation, are more likely to notice news coverage of inflation and incorporate it into their inflation expectations.

C. Robustness to Estimation Strategy

A potential concern to our non-attrition restriction is that the tenure effect we identify may only exist for "ready-to-learn" survey respondents—that is, for respondents who are committed to the survey and thereby have more willingness to learn about the economy. However, quantitatively and qualitatively similar results are obtained for different sampling rules. Focusing on one-year-ahead inflation expectations, Panel A of Figure 2 reproduces our baseline results from Panel A of Figure 1 under various sampling rules. Even if we only include respondents who skip a survey at least one time ("skippers"), who participate in the SCE less than six times in total ("half-participants"), or include the full sample ("full sample"), the estimated tenure effects are similar to those from our "non-attriters" sample. Therefore, throughout this paper, we keep "non-attriters" as our baseline sample.

Another potential identification issue is known as the Age-Period-Cohort (APC) problem, which in its original formulation refers to the problem of separating the independent effects of age, time period, and cohort due to exact linear dependence (Hobcraft, Menken and Preston, 1985; Deaton and Paxson, 1994). In our context, survey experience dummy variables, monthly time-fixed effects, and individual-fixed effects correspond to age, period, and cohort, respectively.[12] One simple solution to this APC problem is to replace monthly time-fixed effects with quarterly ones, which we do for the remainder of the paper. Other solutions may include normalization of the parameters (Deaton and Paxson, 1994), replacing the time-fixed effects with aggregate variables (Heckman and Robb, 1985), or omitting time-fixed effects altogether. Panel B of Figure 2 reproduces our baseline results in Panel A of Figure 1 with each of these alternatives. Our results remain quantitatively and qualitatively similar.[13]

Since central banks often monitor median rather than mean inflation expectations, we also estimate tenure effects on medians using a fixed effects panel quantile regression (Machado and Silva, 2019). The estimated median effects are very close to the mean effects from our baseline regression.

[12]This correspondence is not obvious. Note that first, there is a correspondence between a cohort dummy variable and (a sum of) individual-fixed effects. To see this correspondence, imagine that respondents are "born" when they enter into a survey. Then, a sum of individual dummy variables of respondents who are "born" in period t will be identical to a cohort dummy variable for the respondents who are "born" in period t. Second, note that when there is no panel attrition, # of Survey Experiences (Age) = Current Period (Period) - Survey Entrance Period (Cohort, the date of birth) holds. That is, survey experience dummies (Age), time dummies (Period), and individual dummies (Cohort) are going to be co-linear if they are all used in same time frequency in a linear panel regression.

[13]One possible reason for this robustness could be the fact that inflation rates have been very stable in recent years. Thereby, the effects of time-fixed effects could have been weak during our sample periods; the overall R^2 is virtually identical when we drop the quarterly time-fixed effects entirely.

D. Tenure Effects for Other Survey Measures

Since the SCE includes a variety of other questions about expectations, we estimate analogous regressions for additional outcome variables. In the first column of Table A3, the dependent variable is the respondent's reported percent change that the unemployment rate will be higher in 12 months. Since unemployment fell steadily from June 2013 to February 2020, we restrict the sample to June 2013 to February 2019, so that lower responses are more accurate. The coefficient estimates indicate that responses indeed become more accurate with higher survey tenure.

In columns 2 through 7, the dependent variables are expected price changes for gas, food, medical expenses, college education, rent, and gold over the next 12 months. These questions are only asked of respondents with tenure of at least 2. For all price categories except gold, expectations are 1.2 to 2.7 percentage points lower for respondents with tenure 12. The tenure effects for gold price expectations are smaller, with a coefficient estimate of 0.9 percentage points for tenure 12. Respondents of tenure 2 or greater are also asked for a density forecast of national house price growth over the next 12 months. Column 8 of Table A3 shows that the density forecast interquartile range shrinks with tenure, and is about 1.3 percentage points smaller for tenure 12 respondents compared to tenure 2 respondents.

We also estimate analogous regressions for respondents' nominal personal earnings growth expectations and household income expectations. The estimated effects for personal earnings and income expectations are much smaller when compared to those for inflation expectations. See Columns 9 and 10 of Table A3.[14]

The larger tenure effects for inflation expectations compared to income expectations are in line with rational inattention theory, which suggests that households may selectively pay attention to economic variables. Households should be highly attentive to their own income process even prior to participating in the survey, because it is so relevant to their consumption decisions. However, especially in low-inflation environments, consumers with limited information-processing capacity may pay little attention to inflation. Thus, households may not have a good understanding of the nation-wide average price process before taking the survey. For example, Carroll et al. (2020) show that consumers tend to underreact to aggregate macroeconomic shocks. Consumers may neglect aggregate variables in their consumption decisions because aggregate shocks consists only a small proportion of the uncertainty that consumers face, compared to highly idiosyncratic

[14]One may point out that inflation rates could be naturally harder to forecast than respondents' own income path. That is, for example, one can argue that 0.5 percentage points of learning effects in income expectations should not be treated equally to the same magnitude in inflation expectations. Reflecting this argument that different forecasts may have different scales, we normalize the estimated learning effects by standard deviation or mean of each forecast. However, we still find that the estimated learning effects for inflation expectations are more than 50 percent larger than those for earnings and income expectations.

variables like their own income.[15] Therefore, questions that ask for the respondents' beliefs about inflation are more likely to prompt additional attention to inflation, prompting participants to collect more information about inflation.

E. Heterogeneity in Tenure Effects

While the equation (1) estimates tenure effects for the average respondent, these effects may be heterogeneous depending on households' initial expectations and uncertainty. For example, households who enter the survey with high uncertainty may be more susceptible to learning-through-survey effects since their priors are weaker. To allow for such heterogeneity, we extend equation (1) by including interaction terms of tenure dummies with initial inflation uncertainty of respondents from their first survey:

$$\pi^e_{its} = \sum_{s=2}^{12} \left\{ \beta_{1,s} + \beta_{2,s} IQR_i + \beta_{3,s} IQR_i^2 \right\} \tau_s + \alpha_i + \gamma_t + \varepsilon_{it}, \tag{2}$$

where π^e_{its} denotes one-year-ahead expected inflation of individual i wof tenure s in period t, and IQR_i is the interviewee's *initial* IQR reported in the first survey. We include the squared term IQR_i^2 to control for possible non-linearity. All other terms are defined as in regression (1). The coefficients on the interaction terms are highly statistically significant.

In Figure 3, we plot the estimated tenure effects for values of IQR_i corresponding to the 25th, 50th, and 75th percentiles, which we denote U_L, U_M, and U_H. These values are 1.5%, 3.0%, and 6.4%, respectively. Specifically, using regression (2), assuming $\alpha_i = 0, \gamma_t = 0$, we plot

$$\{\frac{\partial \pi^e_{its}}{\partial \tau_s} \mid \; IQR_i \in \{U_H, U_M, U_L\}\}_{s=2}^{12}$$

Respondents who initially entered the survey with a high level of inflation uncertainty learn more about inflation. Figure 3 shows that respondents whose initial uncertainty over the future inflation rate was in the bottom quartile of the distribution display a small learning effect. However, if the respondents were initially in the top quartile of inflation uncertainty, then the effect is large: right after the first survey, inflation expectations decrease by 2.2 percentage points on average.

We find further evidence of the heterogeneity of the learning effect when including demographic variables and measures of respondents' understanding of

[15] In section II.B of this paper, we provide more direct evidence consistent with Carroll et al. (2020); consumption expectations of new survey participants respond more sluggishly to inflation expectations, while repeat participants more promptly reflect inflation expectations to consumption expectations.

inflation as interactions. For a categorical variable D_i describing a characteristic of respondent i, we estimate the below regression and calculate $\beta_{1,s} + \beta_{2,s} D_i$.

$$\pi^e_{its} = \sum_{s=2}^{12} \left\{ \beta_{1,s} + \beta_{2,s} D_i \right\} \tau_s + \alpha_i + \gamma_t + \varepsilon_{it} \tag{3}$$

We use three demographic variables: income level (less than \$50k, \$50k to \$100k, more than \$100k), education level (college, some college, or high school), and retiree status. In addition, in the SCE, new respondents are required to answer a set of questions measuring their numeracy and financial literacy. A question designed to measure understanding of inflation asks, "Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After one year, how much would you be able to buy with the money in this account?" Respondents can choose "More than today," "Exactly the same," or "Less than today." We measure respondents' understanding of inflation by whether respondents gave a correct answer to this question. Only 52% answered this question correctly.

Consistent with the previous results from Figure 3, Figure 4 shows that respondents who are generally more informed about inflation prior to the survey display significantly smaller learning effects throughout the survey waves. Panels A and B show that the estimated learning effects are substantially smaller for higher-income and more educated individuals, while Panels C and D show that retirees and respondents who gave a correct answer the question measuring understanding of inflation display relatively smaller learning effects.[16] The coefficients on the interaction terms for income, education, retiree status, and inflation understanding are all highly statistically significant. We also run the same regression jointly including all indicator variables for demographics and inflation understanding; the overall results are similar to those in Figure 4, though the interaction terms with inflation understanding loses statistical significance..

In summary, tenure effects are a robust feature of the SCE data. On average, consumers with more past survey experience have inflation expectations that are lower and more accurate. They also tend to have lower uncertainty about inflation and are less likely to update their forecasts in subsequent surveys. Further, survey participants who are generally more informed about inflation prior to the survey display smaller tenure effects.

II. Implications for Interpretation of Survey Data

Inflation expectation surveys conducted by the central banks are generally intended to be used for two major purposes: i) monitoring inflation expectations

[16] Aguiar and Hurst (2007) find that older households invest more in shopping time and pay the lowest prices compared to other households.

through an aggregate index and ii) researching consumer expectations and behavior using the underlying micro data. The results from the previous section imply that aggregate measures of inflation expectations and of inflation uncertainty would be higher if only new participants—those not subject to tenure effects—were included. This was shown in Panels C and D of Figure 1.

Using an episode of oil price collapse in 2014, we show how the tenure effects can potentially impede central banks' monitoring of inflation expectations. Then, by revisiting estimation of elasticity of intertemporal substitution by Crump et al. (2015), we show how the learning effect can influence micro estimates and provide useful insights for studies using survey micro data.

A. Oil Price Decline in 2014

In order to monitor the inflation expectations of U.S. consumers, the FRBNY conducts the SCE data each month and reports the sample median of the density mean inflation expectation. However, the tenure effects we documented suggest that repeat participants' prior survey participation may have prompted them to seek information about or otherwise reflect on inflation. Exploiting the episode of a sharp drop in oil prices during 2014, we show that the dynamics of inflation expectations of new participants can be significantly different than those of the repeat survey participants.

In particular, household inflation expectations are sensitive to oil and gas prices. Coibion and Gorodnichenko (2015*b*) find that the increase in inflation expectations of households during the Great Recession can be attributed to the rise in the oil price, since the price of gasoline is one of the most salient prices for consumers. However, this stylized fact tends to be significantly weaker for repeat participants of the FRBNY SCE compared to new participants.

First, Figure 5 shows that crude oil prices plunged by half in only six months in 2014, from $103.59 per barrel in July 2014 to $50.58 in February 2015, as innovation in Hydraulic Fracturing technology boosted oil production in the U.S. During this period, other macroeconomic conditions were fairly stable; the seasonally-adjusted industrial production index decreased by 0.17 percent and the unemployment rate decreased by 0.7 percentage points.[17]

Figure 5 also compares the median density mean inflation expectations of new SCE participants with those of repeat participants, whose expectations have been subject to the tenure effects. The inflation expectations of repeat participants are relatively steady from June 2013 to February 2015, only declining by 0.27 percentage points from July 2014 to February 2015 as oil prices fell. By contrast, the inflation expectations of new participants generally track high-frequency fluctuations in oil prices, which is consistent with what Coibion and Gorodnichenko (2015*b*) have found. Inflation expectations from the Michigan Survey of Con-

[17] See Baffes et al. (2015) for more discussion on the causes of the oil price decline, including weakening global demand, a significant shift in OPEC policy, geopolitical shifts, and U.S. dollar appreciation.

sumers (MSC) show a similar pattern to the inflation expectations of new participants in the SCE.[18]

We quantitatively evaluate the differences in responses to gas prices between new and repeat participants using the following panel linear regression:

$$\pi^e_{its} = \sum_{s=1}^{12} \beta_s(\tau_s \times log(Gas_t)) + \alpha_i + \gamma_t + \varepsilon_{it} \tag{4}$$

where $log(Gas_t)$ is the log of the monthly gas price, and τ_s is a tenure dummy variable for s number of total survey experience. π^e_{its} denotes one-year-ahead density mean inflation expectations of an individual i whose total number of survey experience is s at period t. α_i and γ_t are individual- and quarterly time-fixed effects. ε_{it} is an error term. The sample period is from July 2014 to February 2015, the period when oil and gas prices plunged. We restrict samples to respondents who participate in the survey for the maximum number of times, as we did in our main result section.

Figure 6 visually shows the estimated regression coefficients $\{\beta_s\}_{s=1}^{12}$ by tenure group s. Clearly, new survey participants display the largest regression coefficients indicating the strongest response of inflation expectations to gas prices. For our benchmark regression, we find that the inflation expectations of new survey participants respond about 50 percent more strongly to gas prices on average when compared to the most experienced participants.[19] This is consistent with our previous finding from the aggregate times-series data. Table A4 presents various regression specifications, including those with truncation of extreme expectations, full sample periods, and point inflation expectations. Qualitative features of our results are not changed. Rather, our benchmark regression specification tends to be conservative when compared to the results from other specifications. When we truncate 10 percent of extreme expectations, new participants respond almost twice as much as repeat participants to gas prices.

Our results are also consistent with those of Verbrugge and Binder (2016), who partitioned the MSC respondents into those with low and high inflation uncertainty, using the methodology in Binder (2017*b*). They show that the inflation expectations of less-uncertain consumers are more stable than those of more-uncertain consumers. In particular, the expectations of less-uncertain consumers did not respond strongly to the oil price decline in 2014.

This evidence suggests that the tenure effects we documented are not constant over time, and thus cannot be removed simply by taking a first difference. The

[18] MSC has a rotating panel component, but respondents are surveyed at most twice, with six months between interviews. Tenure effects on the MSC are discussed in a later section.

[19] We have conducted an F-test for H0: $\beta = \beta_s \ \forall s$ over extended sample periods (from March 2014 to June 2015) and full sample periods (from June 2013 to October 2020) with Driscoll-Kraay standard errors of lag one. For both cases, we could reject the null hypothesis that the regression coefficients across survey tenure are equal to each other at 1% significance level.

expectations of repeat and new participants can exhibit different dynamics in response to economic shocks. In such a case, repeat participants cannot be viewed as representative of the broader population who potentially lack any past survey experience. In this example, if the central bank were only given inflation expectations of repeat participants, they would conclude that the inflation expectations of consumers do not respond to the plunging oil prices and miss some of the timely high-frequency information from survey expectations.

B. Estimating Elasticity of Intertemporal Substitution

The survey micro data on consumer expectations has begun to be used in a variety of applications, and holds great potential for use in many more. The tenure effects that we have documented may affect the estimates and interpretation of such studies. For example, Crump et al. (2015) use the SCE data to estimate the elasticity of intertemporal substitution (EIS). More precisely, they estimate the response of expected consumption growth to changes in expected inflation rates. We revisit this analysis by allowing estimates to vary by respondents' survey experience. Among regression specifications of Crump et al. (2015), for simplicity we focus on the following panel linear regression model with fixed effects:[20]

$$ExpCG^i_{t,t+12} = -\sigma ExpInf^i_{t,t+12} + \gamma ExpIG^i_{t,t+12} + \alpha_i + \beta_t + \varepsilon_{i,t}, \tag{5}$$

where $ExpInf^i_{t,t+12}$ is a 12-month ahead density-implied mean inflation expectation of household i at period t, and $ExpCG^i_{t,t+12}$ is expected real consumption growth over the next 12 months by household i at period t, which is calculated as, $ExpCG^i_{t,t+12} \equiv ExpSG^i_{t,t+12} - ExpInf^i_{t,t+12}$, when $ExpSG^i_{t,t+12}$ is a point forecast for nominal spending growth of the household over the next 12 months. Similarly to the calculation of $ExpCG^i_{t,t+12}$, expected real household income growth, $ExpIG^i_{t,t+12}$, is the difference between point forecast for household nominal income growth and $ExpInf^i_{t,t+12}$. α_i and β_t are individual- and time-fixed effects.

The above expression represents the first-order approximation of a usual consumption Euler equation where σ is the elasticity of intertemporal substitution (EIS) and γ measures "excess sensitivity" of consumption growth to anticipated income changes. The literature commonly finds that expected/predictable income growth has a significant effect on consumption growth. Inclusion of γ in the regression model therefore reflects a possible deviation from the permanent income hypothesis.

First, we estimate the above regression model as-is and find that our estimates

[20] The most recent version of Crump et al. (2015) uses a panel linear regression model *without* fixed effects as their baseline since the SCE data allows many control variables. However, Crump et al. (2015) also show results based on a model with fixed effects and emphasize that their main results remain similar. See section 6.4 and table 9 of Crump et al. (2015).

on σ and γ are indeed very similar to those of Crump et al. (2015); $\hat{\sigma}$ of Crump et al. (2015) under fixed effects is 0.71 and $\hat{\gamma}$ is 0.20 while our estimates are 0.70 and 0.24.[21] Next, we replace σ and γ with $\sum_{s=1}^{12} \sigma_s \tau_s$ and $\sum_{s=1}^{12} \gamma_s \tau_s$ and re-estimate the regression, where τ_s is an indicator variable for tenure s. This modification allows the regression coefficients σ and γ to vary by survey experience of respondents non-parametrically. Also following Crump et al. (2015), we use an instrumental variable (IV) strategy that uses the point inflation expectation as an instrument of density-implied mean inflation expectation, and again allow coefficients to vary by tenure. A full regression table is available in Table A5 of the appendix.

Figure 7 shows that the estimated EIS, $\hat{\sigma}$, increases with survey experience, using either the OLS or IV estimates. That is, more experienced survey participants tend to more actively reflect changes in inflation expectations in their consumption expectations. However, $\hat{\gamma}$ is similar for all tenure groups. Note that the range of EIS estimates in Crump et al. (2015) (around 0.5 to 0.8) are near the lower end of the range of micro estimates from prior literature. Our results indicate that the estimated EIS tends to moderately increase with survey experience. Thus the EIS of the general population who lack any prior survey experience is likely to be even lower than the original estimates of Crump et al. (2015). Survey participation that induces learning about inflation (or equivalently, greater attention to inflation) may result in larger responsiveness to reported inflation expectations by survey respondents. In this example, it results in higher estimate of the EIS.

Why does $\hat{\sigma}$ tend to be larger for experienced survey participants, but not $\hat{\gamma}$? Consistent with our findings throughout the paper, consumers' imperfect attention to aggregate shocks can account for this otherwise puzzling phenomenon. When it comes to spending decisions, consumers tend to focus on their income, but may not pay careful attention to general inflation rates. Therefore, they have sluggish responses to aggregate shocks (Carroll et al., 2020). In other words, if consumers become more attentive to inflation rates because survey experience, their consumption expectations may more quickly respond to change in future inflation (larger $\hat{\sigma}$ with survey experience). In contrast, how households' reported consumption plans respond to expected future income may not change after taking more surveys. Prior to taking a survey, they may already understand their own future income path well and have an established rule for how to adjust their consumption plan with future income changes.

This exercise shows how tenure effects can influence micro estimates. Awareness

[21] Our estimates are slightly different from those of Crump et al. (2015) since we use the same sampling rule as in the rest of this paper, truncating the top and bottom 5% of all point forecasts for each tenure group and period. We restrict samples to respondents who participate in the survey for the maximum number of times in order to minimize the effects of panel attrition. Finally, we use quarterly time-fixed effects instead of monthly time-fixed effects. While our sampling rule is different from that of Crump et al. (2015), as mentioned in the main text, our baseline estimates are very similar to their results, suggesting sampling rule did not drive our results in this section.

of such learning effects is useful for interpretation of analysis using survey micro data. We suggest that it would be good practice for users of survey micro data to check whether their estimates are robust using subsamples of shorter-tenured and longer-tenured respondents.

III. Other Surveys

One potential concern is that the tenure effects might arise from a particular feature of the SCE, such as the short time period of relatively low and stable inflation in which the SCE has been conducted. This section uses the Michigan Survey of Consumers and a survey of U.S. firms to provide evidence of tenure effects in other contexts.

A. Michigan Survey of Consumers

The Michigan Survey of Consumers (MSC), like the SCE, is a monthly survey of consumer expectations. However, whereas the SCE consecutively tracks respondents up to twelve times, the MSC only allows respondents to participate in a maximum of two interviews, with a six-month gap between interviews. A longer gap between surveys tends to reduce the size of tenure effects (Warren and Halpern-Manners, 2012). Despite its more limited panel structure and longer gap between surveys, the MSC does have the advantage of beginning in 1978, rather than 2013, allowing us to check how tenure effects have varied over time and to confirm that they are not only a feature of recent data.

First, Panel A of Figure 8 shows the mean inflation expectations of new and repeat respondents. As before, we apply the non-attrition restriction, so our samples consist of respondents who eventually participated in a follow-up survey.[22] The mean expectations of new respondents are typically slightly higher than those of repeat respondents—on average, the gap is 0.3 percentage points. As expected, this effect is smaller than the effect found in the SCE data. This smaller learning effect in the MSC compared to that in the SCE is consistent with recent evidence from a randomized information treatment experiment that shows providing information about inflation to households has large contemporaneous effects on their expectations but that these effects rapidly diminish over time (Coibion, Gorodnichenko and Weber, 2020).

Panel A also shows that the size of the gap between repeat and new respondents' expectations can vary over time. In order to construct a time-series of tenure effects, we use a regression equation analogous to equation (1), but with fixed effects replaced by demographic control variables to account for the limitations of the MSC dataset. We estimate the following equation for each year t separately—since estimates at the monthly frequency are quite noisy—and obtain a sequence of yearly tenure effects, $\{\hat{\delta}_t\}_{t=1}^{T}$:

[22]Summary statistics of expectations by tenure without the non-attrition restriction are in Appendix Tables A6 and A7 (for the SCE) and Table A8 (for the MSC).

$$\pi^e_{its} = \alpha_t + \delta_t \tau_2 + \beta_t X_{it} + \varepsilon_{it} \tag{6}$$

where π^e_{its} denotes the one-year-ahead point inflation forecast of individual i in year t with tenure $s \in \{1, 2\}$, τ_2 is an indicator variable for tenure 2, X_{it} is a vector of control variables including sex, education, region, the number of kids, marital status, log of nominal household income, age, and age squared, and ε_{it} is an error term.

Panel B of Figure 8 shows resulting estimates of the yearly tenure effects, along with CPI inflation and shaded bars indicating recessions. These tenure effects δ_t are nearly identical to the difference between the mean expectations of repeat and new respondents in year t, and have mean 0.3 and standard deviation 0.2.

We see that tenure effects vary over time, and tend to be larger in magnitude during recessions, when economic uncertainty is high (the exception is the early 2000s recession). This suggests that households form inflation expectations in a Bayesian manner, putting more weight on new information when they are more uncertain in their beliefs. The magnitude of the tenure effects are also positively correlated with inflation uncertainty, disagreement, and volatility, with correlation coefficients of 0.2, 0.6, and 0.3, respectively.[23]

Tenure effects also vary with inflation. In particular, the largest negative values of δ_t occur when inflation is falling or about to fall, and the near-zero or positive values occur when inflation is rising or about to rise. Most notably, the most negative value of δ_t (-0.82) occurs in 1982. Inflation fell from 10% in 1981 to 6% in 1982 and would fall to 3% in 1983. New respondents, less informed about falling inflation, reported much higher inflation expectations than repeat respondents. A similar pattern occurs in 1990 and 1991: new respondents were less aware than repeat respondents that inflation was beginning to decline. Another striking example occurs in 2008. Though inflation fell during the Great Recession, many consumers expected the Recession to be inflationary, so expectations rose sharply, especially for new respondents. Repeat participants' expectations were 0.6 percentage points lower than those of new respondents. Conversely, the most positive value of δ_t (0.37) occurs in 2021. In most years, consumer inflation expectations are higher than realized inflation, so consumers revise their forecasts downward as they become more informed. Recently, as inflation has risen sharply, more informed consumers revise their expectations upward.

To summarize, the results from the MSC show that during most periods, repeat participants generally report lower inflation expectations than new participants, though the degree of the learning-through-survey effects changes over time. In addition, the learning-through-survey effect tends to be smaller for the MSC than

[23]Inflation uncertainty is the updated inflation uncertainty index from Binder (2017*b*). Inflation disagreement is the cross-sectional standard deviation of inflation expectations from the MSC. Inflation volatility is the rolling five-year standard deviation of annual CPI inflation.

those of the SCE, likely because there is a longer period of time between baseline and follow-up surveys.

B. Inflation Expectations of Firms

While our focus so far has been household surveys, Coibion, Gorodnichenko and Kumar (2018) show that the inflation expectations of firm managers tend to resemble those of households, which suggests that tenure effects may exist in firm surveys as well. To study whether this is the case, we use a new firm expectation survey targeting businesses in the U.S. (Candia, Coibion and Gorodnichenko, 2021)

This firm survey, the new Survey of Firms' Inflation Expectations (SoFIE), is collected by a business intelligence company that has been collecting CEOs' and top executives' perceptions and expectations for various firm-specific economic outcomes. The panel is intended to be representative of "the underlying structure of each sector in the economy according to its contribution to the gross value added."[24] The survey covers the U.S. firms in manufacturing and services sectors. About 300 to 600 firms participate in this survey each wave and stay in the panel for about three waves on average. A question asking about one-year-ahead CPI inflation rates was added its quarterly survey in 2018. We use this data, which was collected from April 2018 to April 2020 at quarterly frequency.

We estimate the learning-through-survey effects for firms in regressions analogous to equation (1). However, since only nine survey waves are available, we relax our "non-attrition" restriction to save observations; we restrict our samples to consist of firms who participated in the survey more than three times. As before, we either use Deaton's normalization method or aggregate control variables to avert the APC problem[25], and winsorize the top and bottom 5% of data.

Panel C of Figure 8 shows the resulting estimates of the tenure effects from the firm survey. Inflation expectations of repeat survey participants decrease with survey experience on average, consistent with what we found in the consumer surveys. This provides additional evidence that firm executives, who are likely the price-setters in the economy, typically face information constraints that may influence their expectations of aggregate inflation. For example, Mackowiak and Wiederholt (2009) show that firms pay more attention to idiosyncratic variables because firm-specific conditions are generally more important in their decision-making than aggregate conditions.

Again, tenure effects in the firm survey are smaller than those in the SCE, possibly because the time between surveys is longer. Firm respondents of tenure 3 (6 months after first participation) have about 0.20 to 0.33 percentage points lower inflation expectations than new participants. In the SCE, after 6 months of

[24]Candia, Coibion and Gorodnichenko (2020) and `http://firm-expectations.org` provide more detailed information about this firm survey. In particular, information about survey representativeness is at `http://firm-expectations.org/weightsbuilding.html`.

[25]We cannot use the quarterly time fixed effects method for this case since our data is quarterly.

participation, repeat survey participants have about 2.2 percentage points lower inflation expectations than those of new participants. Another reason why the tenure effect in a firm survey is smaller than that of household surveys could be that the initial inflation expectations of firms are generally more accurate than those of households; after the winsorization, the average inflation expectation of firms is 2.8 percent, with standard deviation 1.86.

While the dataset is limited, the firm survey results confirm that the tenure effects are not exclusive to household surveys. It also shows that the degree of panel conditioning may depend on time between surveys and respondents' prior level of knowledge on the subjects being asked.

IV. Discussion and Conclusion

We have shown that the reported beliefs of survey respondents significantly change over their tenure in the survey. We have documented the prevalence of these tenure effects across two surveys of consumers and a survey of firms. The size of the effects vary depending on the type of survey question, survey frequency, and respondent characteristics. The effects are particularly large and robust for inflation expectations reported on the FRBNY SCE: consumer inflation expectations and uncertainty decrease notably with survey tenure, especially after the first round of survey participation.

We believe that our results have implications both for our understanding of the expectations formation process and information rigidities, and for the design and interpretation of expectations surveys in the future.

In the literature on consumer expectations formation, two major questions that arise are (1) How do consumers allocate their limited attention and cognitive efforts?, and (2) Why do consumers disagree so much with each other and with professional forecasters? In some of this literature, households allocate their attention and cognitive efforts based on the cost and benefit of information acquisition, and as they acquire new information, they update beliefs in a Bayesian manner. For example, when households' priors are less certain, they put more weight on new information. Then, differences in beliefs can arise from different information processing constraints and different perceived benefits in acquiring certain types of information.

Households may pay little attention to inflation, either because it is costly for them to do so, or because they find it more worthwhile to devote their attention to other things. Our primary interpretation of our main results is that when a survey asks respondents about their inflation expectations, it may prompt them to devote more attention to inflation before retaking the survey (such as by looking up official inflation statistics or media reports, talking to acquaintances, or reflecting on their own experiences and the prices they have observed). Indeed, respondents significantly and predictably revise their expectations in subsequent surveys. The revisions are consistent with Bayes' rule, in the sense that forecast errors shrink and that the impact of past survey experience is larger for households who are generally less informed about inflation prior to the survey. The tenure effects are smaller in the firm survey, as respondents may be more informed about inflation prior to their first survey experience.

Moreover, the panel conditioning effects are minimal for questions about respondents' own earnings or income. Households are likely more attentive to these than to aggregate inflation, so the act of taking a survey does not much change their attention to these variables.

The implications of our findings for the design and interpretation of expectations surveys depends on whether our explanation for the panel conditioning effects is correct, and on details of how the survey data is being used. We have suggested that the panel conditioning effects arise because respondents pay more attention to inflation after being asked about inflation in one or more rounds of a survey ("cognitive stimulus" or "learning-through-survey.") One alternative explanation is a "reporting error" hypothesis.[26] Under this hypothesis, survey respondents have an underlying belief distribution about inflation that is *not* affected by survey participation. However, respondents with lower survey tenure answer questions with extra reporting error, because they must expend cognitive effort to formalize, retrieve and report their underlying beliefs accurately.

More formally, under the reporting error hypothesis, respondent i of tenure s reports an inflation expectation $r_{its} = \pi^e_{its} + \varepsilon_{its}$ where π^e_{its} is the respondent's true underlying inflation expectation in period t and ε_{its} is a reporting error. Suppose that as a respondent gains more experience in answering survey questions, then he becomes better in expressing his beliefs more accurately. That is, the distribution of ε_{its} becomes tighter as s increases. As we and others have documented, consumer inflation expectations are upward-biased, so ε_{its} is not zero in expectation, but rather positive. Suppose, for example, it has a log normal distribution, $\ln(\varepsilon_{its}) \sim N(0, \sigma_s^2)$, and σ_s^2 decreases with survey experience. This simple model can generate some of the effects documented in our paper: reported inflation expectation levels and uncertainty decrease with survey experience.

This reporting error would need to depend on demographic characteristics such as education and income in order to explain the heterogeneity in learning effects we found. If the reporting error is simply a type of measurement error, with σ_s^2 constant over time, then this hypothesis does not explain the time-varying nature of the tenure effects, and their correlation with economic shocks. In our gas prices example in section II.A, for example, we showed that the tenure effects vary with gas prices. We also showed that the difference between new and repeat respondents' expectations was especially large at the start of the Covid-19 pandemic. So for the reporting error hypothesis to explain our results, the error or σ_s^2 would need to be allowed to depend on economic conditions. While the learning hypothesis seems more likely to us, we acknowledge that the reporting error hypothesis cannot be entirely ruled out. It is also possible that learning effects and reporting error effects are both present.

[26]We thank an anonymous referee for providing constructive comments related to this issue.

Future research to help disentangle these two hypotheses—perhaps by using focus groups and cognitive interviews of survey participants, or by adding special questionnaires that ask about behavior between survey rounds—will be quite important, because these hypotheses have different implications. The central issue is whether the change in responses throughout the survey waves reflect changing beliefs about the economy, or indicate that respondents are improving when expressing beliefs. Under the reporting error hypothesis, reported beliefs of respondents with longer tenure become *more* representative of the "true" underlying beliefs of the population. Under the learning hypothesis, in contrast, reported beliefs of respondents with longer tenure become *less* representative of population beliefs; instead they reflect the beliefs of people who are more attentive to inflation. For some applications, these beliefs may be of higher interest to researchers, if they more resemble the beliefs of price-setters or of consumers who are planning to make a major financial decision, for example.

Either way, central banks running household surveys to measure inflation expectations, and researchers working with the underlying microdata, should take note of our evidence of tenure effects. In a discussion of the SCE, Armantier et al. (2017, p. 64) argue that "the design of the panel, with a constant in- and outflow of respondents each month, ensures a stable survey tenure distribution, so the extent of learning and experience (and any associated impact on responses) is constant over time. As a result, month-to-month changes in median responses should capture real changes in population beliefs." But because of the time-varying nature of the tenure effects, month-to-month changes in median responses may not always capture changes in population beliefs in such a straightforward manner (as we showed in our gas prices example).

This is not to say that the panel component of the survey should be removed. The panel component has the clear benefit of allowing researchers to control for unobservable individual characteristics. If central banks wish to minimize the panel conditioning effects, one option is to increase the time length between surveys to minimize learning effects. Another option is to increase the size of the sample of new participants in each wave, but only invite some fraction of them to become repeat participants. This would allow researchers to conduct analysis on the full panel, on new participants only, or on repeat participants only, as appropriate to the situation. In addition, it would be good practice for users of survey microdata to check whether their estimates are robust to using subsamples of shorter-tenured and longer-tenured respondents.

REFERENCES

Aguiar, Mark, and Erik Hurst. 2007. "Life-cycle prices and production." *American Economic Review*, 97(5): 1533–1559.

Arioli, Rodolfo, Colm Bates, Heinz Dieden, Ioana Duca, Roberta Friz, Christian Gayer, Geoff Kenny, Aidan Meyler, and Iskra Pavlova.

2017. "EU consumers' quantitative inflation perceptions and expectations: An evaluation." ECB Occasional Paper.

Armantier, Olivier, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2017. "An Overview of the Survey of Consumer Expectations." *Federal Reserve Bank of New York Economic Policy Review*, 23(2): 51.

Baffes, John, M Ayhan Kose, Franziska Ohnsorge, and Marc Stocker. 2015. "The great plunge in oil prices: Causes, consequences, and policy responses." *Consequences, and Policy Responses (June 2015).*

Binder, Carola. 2017*a*. "Consumer forecast revisions: Is information really so sticky?" *Economics Letters*, 161: 112–115.

Binder, Carola C. 2017*b*. "Measuring uncertainty based on rounding: New method and application to inflation expectations." *Journal of Monetary Economics*, 90: 1–12.

Binder, Carola Conces. 2020. "Coronavirus Fears and Macroeconomic Expectations." *Review of Economics and Statistics*, 102(4): 721–730.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko. 2020. "Communication and the Beliefs of Economic Agents." National Bureau of Economic Research.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko. 2021. "The Inflation Expectations of US Firms: Evidence from a new survey." National Bureau of Economic Research.

Carroll, Christopher D, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White. 2020. "Sticky expectations and consumption dynamics." *American economic journal: macroeconomics*, 12(3): 40–76.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015*a*. "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105(8): 2644–78.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015*b*. "Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation." *American Economic Journal: Macroeconomics*, 7(1): 197–232.

Coibion, Olivier, Yuriy Gorodnichenko, and Mauricio Ulate. 2019. "Is Inflation Just Around the Corner? The Phillips Curve and Global Inflationary Pressures." *AEA Papers and Proceedings*, 109: 465–69.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2020. "Monetary policy communications and their effects on household inflation expectations." National Bureau of Economic Research.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How do firms form their expectations? new survey evidence." *American Economic Review*, 108(9): 2671–2713.

Crump, Richard K, Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa. 2015. "Subjective intertemporal substitution." *FRB of New York Staff Report*, , (734).

Deaton, Angus S, and Christina Paxson. 1994. "Saving, growth, and aging in Taiwan." In *Studies in the Economics of Aging*. 331–362. University of Chicago Press.

Dräger, Lena, and Michael J Lamla. 2017. "Imperfect information and consumer inflation expectations: Evidence from microdata." *Oxford Bulletin of Economics and Statistics*, 79(6): 933–968.

Gabaix, Xavier. 2014. "A sparsity-based model of bounded rationality." *The Quarterly Journal of Economics*, 129(4): 1661–1710.

Halpern-Manners, Andrew, John Robert Warren, and Florencia Torche. 2017. "Panel Conditioning in the General Social Survey." *Sociological methods & research*, 46(1): 103–124.

Heckman, James, and Richard Robb. 1985. "Using longitudinal data to estimate age, period and cohort effects in earnings equations." In *Cohort analysis in social research*. 137–150. Springer.

Hobcraft, John, Jane Menken, and Samuel Preston. 1985. "Age, period, and cohort effects in demography: a review." In *Cohort analysis in social research*. 89–135. Springer.

Karahan, Fatih, Sean Mihaljevich, and Laura Pilossoph. 2017. "Understanding Permanent and Temporary Income Shocks." *URL link retrieved on*, 3(02): 2018.

Machado, José AF, and JMC Santos Silva. 2019. "Quantiles via moments." *Journal of Econometrics*, 213(1): 145–173.

Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal sticky prices under rational inattention." *American Economic Review*, 99(3): 769–803.

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.

Powell, Jerome. 2019. "Hearing Before the Committee on Banking, Housing, and Urban Affairs." The Semiannual Monetary Policy Report to the Congress.

Reis, Ricardo. 2006. "Inattentive consumers." *Journal of Monetary Economics*, 53(8): 1761–1800.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50(3): 665–690.

Verbrugge, Randal J, and Carola Conces Binder. 2016. "Digging into the downward trend in consumer inflation expectations." *Economic Commentary*, , (2016-11).

Warren, John Robert, and Andrew Halpern-Manners. 2012. "Panel conditioning in longitudinal social science surveys." *Sociological Methods & Research*, 41(4): 491–534.

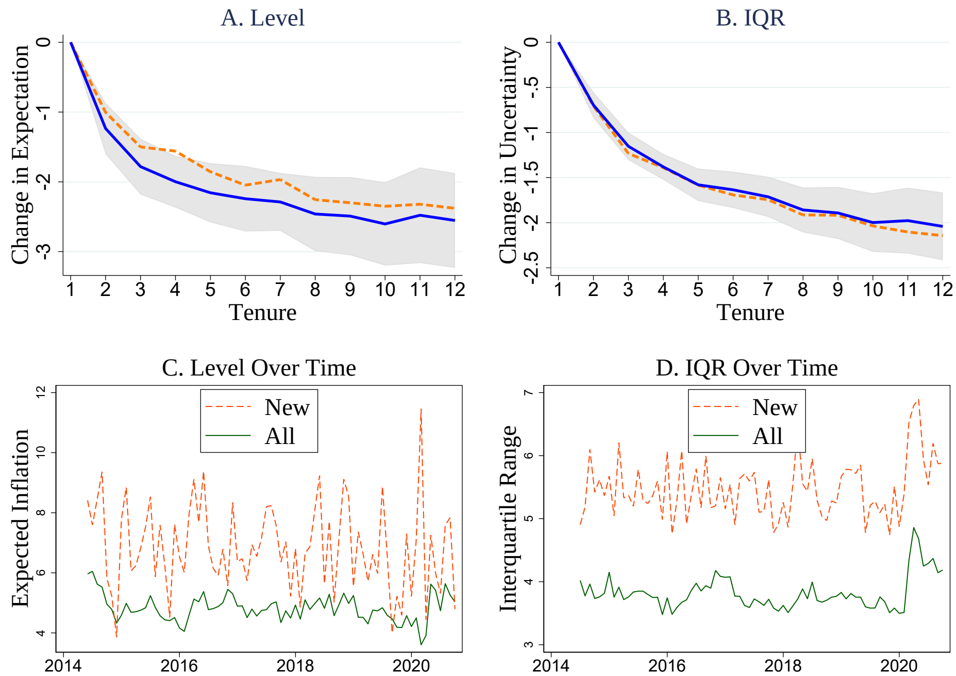

Figure 1. Average Learning-Through-Survey Effects on Inflation Expectations in the SCE

Note: Panels A and B show the change in responses of survey participants compared to their initial responses, in percentage points, estimated from regression (1). For Panel A, the dependent variable is the inflation point forecast, and for Panel B, the dependent variable is the interquartile range of the density forecast. The solid blue (dashed orange) lines correspond to one-year (three-year) ahead inflation forecasts. The gray area shows a 95% confidence interval for the solid blue line with Driscoll-Kraay standard errors of lag one. Survey tenure is shown on the x-axis. We restrict samples to respondents who eventually participate in the survey for twelve waves (non-attriters) and winsorize the top and bottom 5% of each dependent variable for each tenure group and period. A full regression table is in Appendix Table A1. Panels C and D show the mean inflation point forecast and mean inflation density interquartile range over time for new respondents and all respondents, in this case without the non-attrition restriction. Data is from the FRBNY Survey of Consumer Expectations, from June 2013 to October 2020 with monthly frequency.

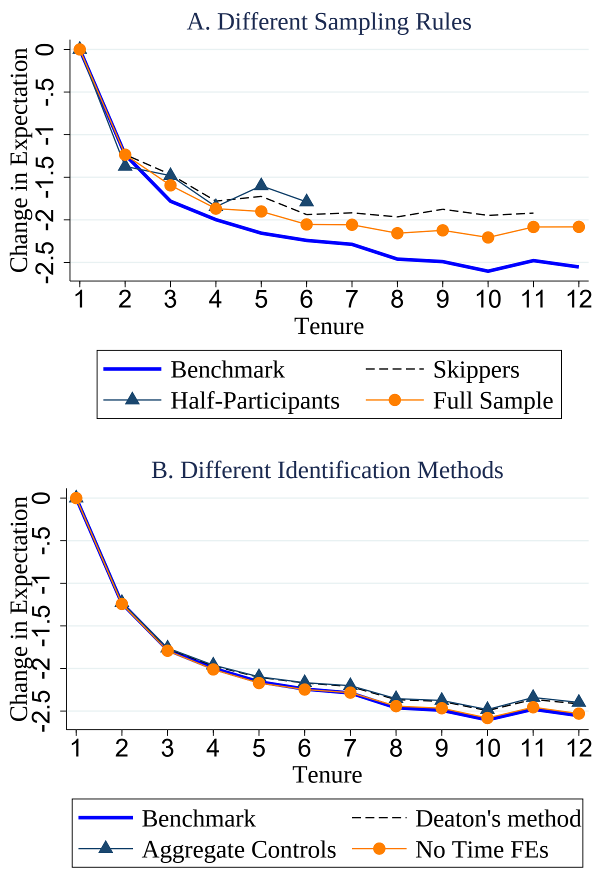

Figure 2. Baseline Results under Different Sampling Rules and Identification Methods

Note: Panel A reproduces the results from Panel A of Figure 1 under different sampling rules. "Benchmark" corresponds to the baseline results using non-attriters only. "Skippers" are respondents who skip a survey at least once. "Half-participants" participate in the SCE no more than six times. "Full sample" corresponds to the case when we do not make restrictions based on total survey participation.
Panel B reproduces the results of one-year-ahead inflation forecasts in Panel A of Figure 1 using different identification methods. "Benchmark" corresponds to the baseline results: linear panel fixed effects regression with quarterly time fixed effects and individual fixed effects. "Deaton's method" uses normalization of monthly time fixed effects following Deaton and Paxson (1994). "Aggregate Controls" replaces time-fixed effects with macroeconomic aggregate variables: monthly CPI inflation rates, the aggregate median of MSC one-year-ahead inflation forecasts, unemployment rate, monthly growth rate of the industrial production index, and log of average WTI oil prices. "No Time FEs" corresponds to the case when neither time-fixed effects nor aggregate control variables are used. We winsorize the top and bottom 5% of dependent variables for each tenure group and period.

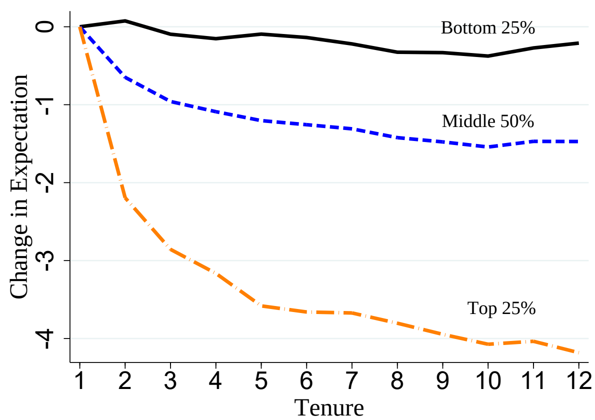

Figure 3. Learning Effects on Inflation Expectations by Initial Inflation Uncertainty

Note: The figure plots the learning-through-survey effects by initial inflation uncertainty ($IQR_i \in \{U_H, U_M, U_L\}$), which are estimated from the equation (2). For example, the top 25% line (long dashed orange line) corresponds to the case when a respondent had a high level of inflation uncertainty in the first interview, assuming $\alpha_i = 0, \gamma_t = 0$. The y-axis shows the change in one-year-ahead inflation expectation of respondents compared to their initial responses, in percentage points. Tenure is shown on the x-axis and corresponds to the total number of survey experiences of each respondent (including the current survey wave). Sample is restricted to non-attriters. We truncate the top and bottom 5% of the dependent variable for each tenure group and period.

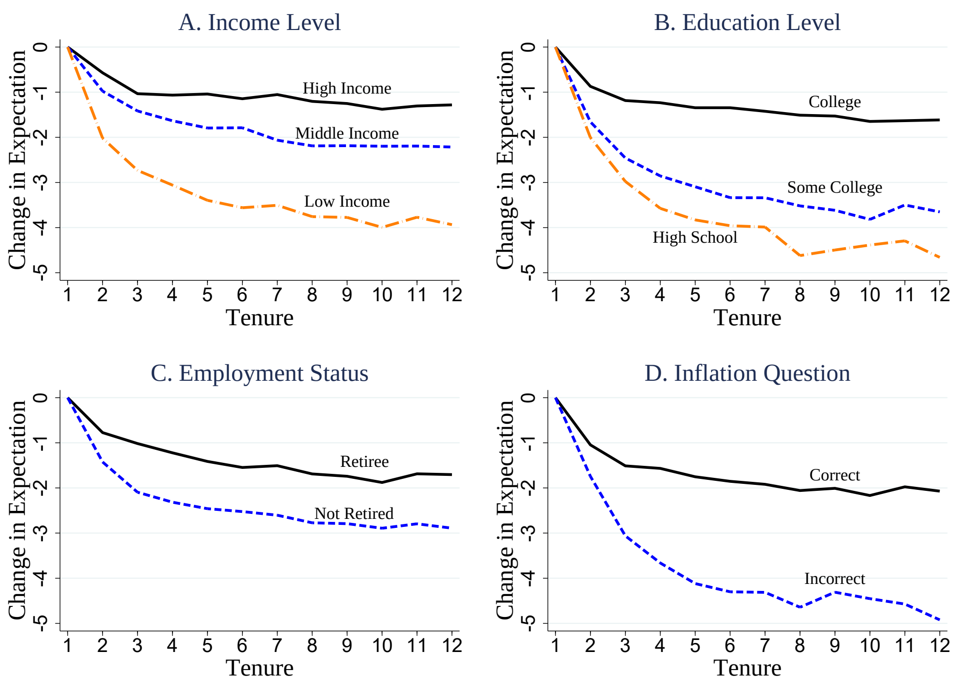

FIGURE 4. LEARNING EFFECTS BY DEMOGRAPHICS AND UNDERSTANDING OF INFLATION

Note: Each panel plots learning-through-survey effects on one-year-ahead inflation expectation by demographic variables and inflation understanding: income level, education level, retiree status, and whether respondents gave a correct answer to a question asking about inflation. Estimates are obtained from regression equation (3) using the indicated dummy variables as the interaction term. Sample is restricted to non-attriters. We winsorize the top and bottom 5% of the dependent variable for each tenure group and period.

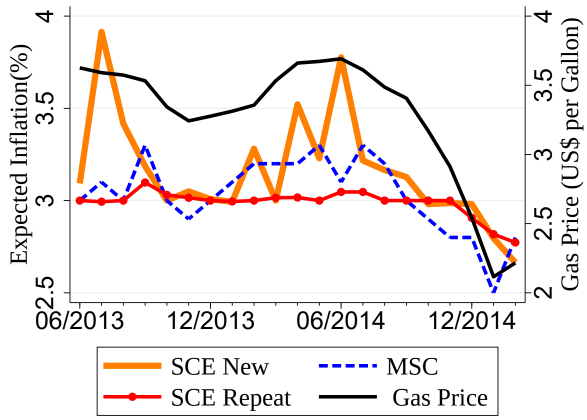

FIGURE 5. INFLATION EXPECTATION OF NEW AND REPEAT PARTICIPANTS

Note: The monthly average nominal WTI crude oil price per barrel in US$ is on the right y-axis (thin solid black line; "Oil Price"). For the left y-axis, one-year-ahead median density mean inflation expectations of new survey participants of the SCE (thick solid orange line; "SCE New"), repeat survey participants of the SCE (connected red line; "SCE Repeat"), and the median inflation expectations of Michigan Survey of Consumers (dashed blue line; "MSC") are presented in percentage points. Data is from the FRBNY Survey of Consumer Expectations and Federal Reserve Economic Data.

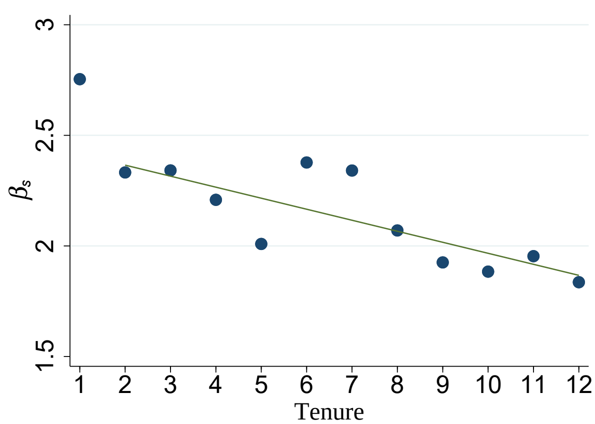

FIGURE 6. RESPONSES TO GAS PRICES BY SURVEY TENURE

Note: The regression coefficients $\{\beta_s\}_{s=1}^{12}$ obtained from our benchmark regression (4), which measures the response of density mean inflation expectation to the increase in log gas prices ($= \partial\pi^e/\partial log(Gas)$), are presented in the figure by each tenure group, s. Tenure is shown on the x-axis and corresponds to the total number of survey experiences of each respondent, including the current survey wave. A linear fitted line is presented for repeat participants (tenure>1). Data is from the FRBNY Survey of Consumer Expectations, July 2014 to February 2015 February. A full regression table is available in Appendix Table A4.

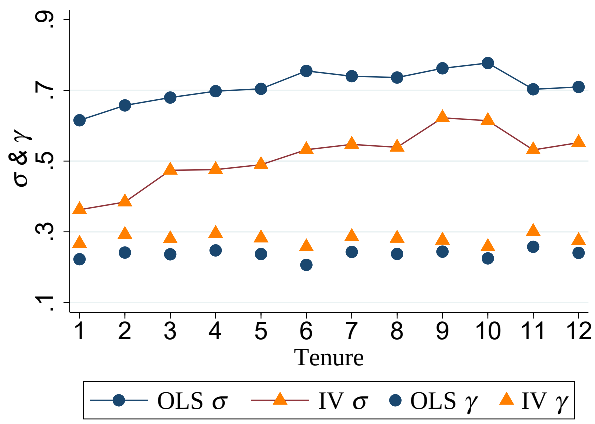

Figure 7. Estimates of EIS and Excess Sensitivity of Consumption by Survey Tenure

Note: We run a linear panel regression as in Crump et al. (2015), allowing regression coefficients to vary by survey experience : $ExpCG^i_{t,t+12} = -\sum_{s=1}^{12} \tau_s \sigma_s ExpInf^i_{t,t+12} + \sum_{s=1}^{12} \tau_s \gamma_s ExpIG^i_{t,t+12} + \alpha_i + \beta_t + \varepsilon_{i,t}$. The estimated regression coefficients $\{\hat{\sigma}_s\}_{s=1}^{12}$ (EIS) and $\{\hat{\gamma}_s\}_{s=1}^{12}$ (Excess Sensitivity) are presented in the figure. For the case of IV, the point inflation expectation is used as an instrument of density-implied mean inflation expectation. The sample is restricted to non-attriters. We truncate the top and bottom 5% of all point forecasts for each tenure group and period. Data is from the FRBNY Survey of Consumer Expectations, June 2013 to October 2020.

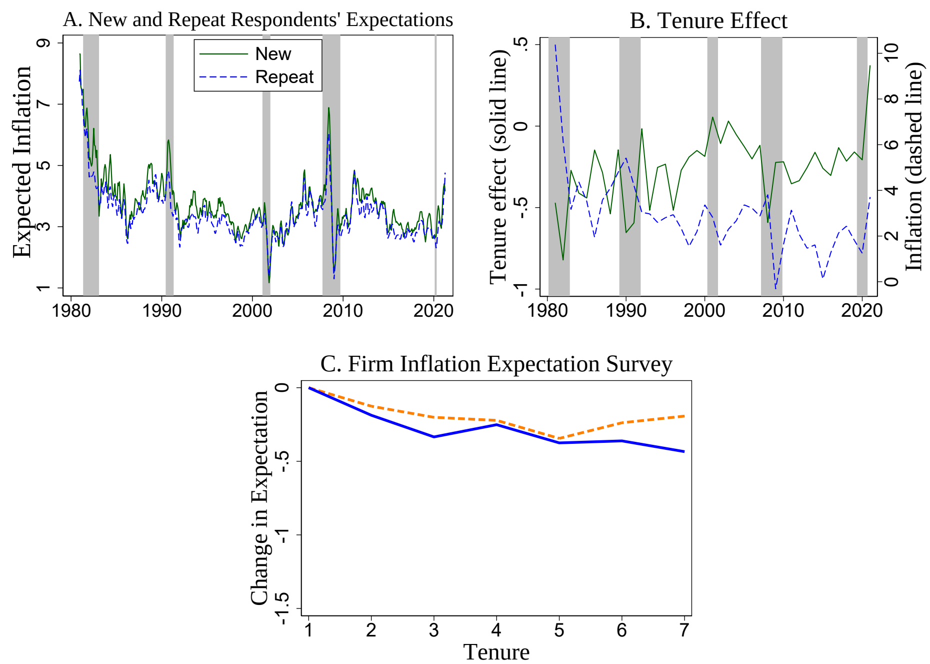

FIGURE 8. LEARNING-THROUGH-SURVEY EFFECTS ON THE OTHER SURVEYS

Note: Panel A shows the mean inflation expectations of new and repeat respondents from the Michigan Survey of Consumers (MSC), January 1981 to September 2021. Centered three-month moving average is shown for visual clarity. Sample is restricted to non-attriters and responses are winsorized as in previous section.

Panel B shows an annual time-series of learning-through-survey effects, which are estimated from a regression equation (6). The solid line is a plot of $\{\hat{\delta}_t\}_{t=1}^{t=T}$, which are regression coefficients attached to the tenure dummies. If $\hat{\delta}_t$ is negative, then the second-time interviewees have lower inflation expectations than the first-time interviewees in the period t. The dashed line is CPI inflation. Shaded bars indicate NBER recessions.

Panel C shows tenure effects from a U.S. firm survey. The percentage points change in inflation expectations of survey participants compared to their initial responses is presented on the y-axis. The solid blue line corresponds to the results from Deaton's method, which normalizes quarterly dummy variables following Deaton and Paxson (1994). The dashed orange line corresponds to the results when macroeconomic aggregate variables are used to control for time effects in a linear panel fixed effects regression, including monthly CPI inflation rates, the S&P 500 stock price return, unemployment rate, and the log of average WTI oil prices. We restrict samples to consist of firms who eventually participate in the survey more than three times and winsorize the top and bottom 5% of the data. The sample period is 2018Q2 to 2020Q2.