

Exposure to Grocery Prices and Inflation Expectations

Francesco D'Acunto

Boston College

Ulrike Malmendier

*University of California, Berkeley, Centre for Economic Policy Research,
and National Bureau of Economic Research*

Juan Ospina

Banco de la República de Colombia

Michael Weber

University of Chicago and National Bureau of Economic Research

Consumers rely on the price changes of goods in their grocery bundles when forming expectations about aggregate inflation. We use micro data that uniquely match individual expectations, detailed information about consumption bundles, and item-level prices. The weights consumers assign to price changes depend on the frequency of purchase, rather than expenditure share, and positive price changes loom larger than negative price changes. Prices of goods offered in the same store but not purchased do not affect inflation expectations, nor do other dimensions. Our results provide empirical guidance for models of expectations formation with heterogeneous consumers.

We thank Klaus Adam, Sumit Agarwal, George-Marios Angeletos, Rudi Bachmann, Olivier Coibion, Stefano Eusepi, Andreas Fuster, Nicola Gennaioli, Yuriy Gorodnichenko, Theresa Kuchler, Emanuel Moench, Ricardo Perez-Truglia, Chris Roth, Eric Sims, Johannes Stroebel, Giorgio Topa, Laura Veldkamp, Nate Vellekoop, Mirko Wiederholt, Johannes

Electronically published March 19, 2021

[*Journal of Political Economy*, 2021, vol. 129, no. 5]

© 2021 by The University of Chicago. All rights reserved. 0022-3808/2021/12905-0008\$10.00

I. Introduction

In his seminal islands model, Lucas (1972, 1973) posited that agents use the prices they directly observe in their daily lives to form expectations about aggregate inflation. As he discussed in Lucas (1975, 1122–23), “[T]he history of prices . . . observed by an individual is his source of information on the current state of the economy and of the market z in which he currently finds himself; equivalently, this history is his source of information on future prices.” Although Lucas did not aim to provide a literal description of reality, this assumption triggered a debate about its logical consistency and realism. To what extent are consumers relying on prices they personally observe to form expectations about aggregate inflation, rather than simply looking up money supply (or, nowadays, the inflation rate on the internet)? Despite the relevance of this assumption for modern models of belief formation, such as behavioral approaches as well as models of rational inattention, the evidence to assess its plausibility is scant. Assessing the empirical plausibility of this assumption is especially important in times of low interest rates and inflation (Summers 2018), in which the ability to manage households’ inflation expectations is key for the effectiveness monetary and fiscal policies (Feldstein 2002; Yellen 2016; Lagarde 2020).

In this paper, we bring the Lucas assumption to the data and investigate the extent to which consumers rely on the grocery-price changes they observe in their consumption bundles to form expectations about aggregate inflation. Our data uniquely link expectations, consumption bundles, and item-level prices at the consumer level. The richness of these data allows us to investigate the characteristics of price changes that matter the most in the expectations formation process.

Wohlfart, Basit Zafar, and conference and seminar participants at the 2017 American Economic Association annual meeting, the European Central Bank Conference on “Understanding Inflation: Lessons from the Past, Lessons for the Future,” the Ifo Conference on Macro and Survey Data, the 2019 Stanford Institute for Theoretical Economics workshop, the Cleveland Fed Conference on Inflation for valuable comments, and the University of Chicago. We also thank Shannon Hazlett and Victoria Stevens at Nielsen for their assistance with the collection of the PanelViews Survey. Yann Decressin and Krishna Kamepalli provided excellent research assistance. We gratefully acknowledge financial support from the University of Chicago Booth School of Business and the Fama-Miller Center to run our surveys. Weber also acknowledges the hospitality, during part of the research process for this paper, by the Vienna University of Economics and Business as Engelbert Dockner Fellow. Researchers’ own analyses were calculated (or derived) in part on the basis of data from The Nielsen Company (US) and marketing databases provided by the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Earlier versions of this paper have circulated with the titles “Salient Price Changes, Inflation Expectations, and Household Behavior” and “Exposure to Daily Price Changes and Inflation Expectations.” Data are provided as supplementary material online.

We find that the price changes of goods consumers purchase significantly influence their expectations about aggregate inflation. The weights consumers assign to different price changes in their grocery consumption bundle depend on the frequency of purchase, rather than the expenditure share, and positive price changes receive a larger weight than negative ones. The prices of goods in the same store that consumers do not purchase (any more) do not affect inflation expectations, nor do other dimensions of price changes, such as their volatility.

These results are a robust feature of the data and do not depend on details of the inflation calculation, such as considering gross prices rather than net prices (after discounts and coupons); using shopping trips or number of goods purchased to compute the frequency weights; varying the time horizon or the granularity of the definition of goods; moving from Laspeyres to alternative ways of defining the consumption bundle; excluding goods purchased at low frequencies; and using the maximum or median price changes or excluding temporary sales when calculating household-level inflation measures.

Our results are important in that they provide empirical guidance on which features of price signals are relevant, or irrelevant, to the formation of macro expectations. As such, they help advance models featuring heterogeneous beliefs.

To analyze the role of household-specific price changes on beliefs, we construct a novel data set. We combine detailed information about the quantity and prices of the nondurable consumption baskets of more than 90,000 households in the Kilts-Nielsen Consumer Panel (KNCP) with new survey data on expectations we elicited from all members of the Nielsen households in June 2015 and June 2016. These data allow us to construct household-level inflation measures and match them with the inflation expectations of each survey participant at the time they shopped for groceries. Because of this level of granularity, we can study in detail which price changes are most relevant to shape inflation expectations, while keeping constant a large range of individual-level characteristics as well as other personal and macroeconomic expectations.

We construct a variety of household-level inflation measures, which capture alternative features of personal grocery-price changes. Our first measure, the Household CPI, mirrors the consumer price index (CPI) but uses each household's nondurable consumption basket instead of a representative consumption basket. The Household CPI is a significant predictor of 12-months-ahead inflation expectations. For example, when we group households into eight equal-sized bins of Household CPI, the difference in expected inflation between households in the lowest and highest bin is 0.5 percentage points (pp). This difference is economically sizable, given a realized inflation rate of around 1% during the same period. The results hold, conditioning on a rich set of demographics

including age, income, gender, marital status, household size, education, employment status, and risk tolerance. Within-individual analyses across the two survey waves also confirm the results. Thus, time-invariant individual characteristics, such as cognitive abilities or financial sophistication, cannot explain our findings.

Building on the finding that personal price changes affect beliefs about aggregate inflation, we then ask whether consumers weigh price changes on the basis of expenditure shares, as the CPI assumes, or instead on the basis of their frequency of exposure. The latter would be consistent with consumers perceiving the price signals from frequently purchased goods as more precise (Angeletos and Lian 2016) or easier to recall (Georganas, Healy, and Li 2014). We construct a second measure, the Frequency CPI, which uses the frequency of purchases to weigh price changes. The positive association between the Frequency CPI and inflation expectations is 20%–40% larger than the association with the Household CPI. When we include both measures, the coefficient of the Household CPI shrinks to zero and loses statistical significance, whereas the statistical and economic significance of the Frequency CPI barely changes. The estimation results are also robust to computing alternative versions of the Frequency CPI based on the number of trips in which households purchase a good or considering only goods households purchase in high volumes.

We also consider a large array of specific features of price changes that prior research has suggested as potential determinants of consumers' belief formation, including their sign, volatility, horizon, and technical details of the inflation weighting. The one aspect that robustly matters is the sign of price changes: positive price changes influence expectations more than negative ones. This differential effect of positive over negative price changes is robust to using gross prices (instead of prices net of discounts) and to excluding temporary price cuts such as weekly sales in retail scanner data. In other words, the asymmetric overweighing of positive price changes does not appear to reflect differential persistence in the price changes. Instead, the result is consistent with Cavallo, Cruces, and Perez-Truglia (2017), who argue that households pay more attention to price increases than to price decreases.

Because we investigate many dimensions of price changes, one might be concerned about the role of multiple testing and searching across different measures for our results. We show that the frequency of purchase and the higher relevance of positive price changes compared to negative ones remain significant dimensions for the expectations formation process of individuals after adjustment of p -values for multiple testing.

We also assess the explanatory power of past observed price changes on individuals' inflation expectations in more detail. The R^2 estimated in the purely cross-sectional part of our baseline regressions amounts to less than 10%. Since the Nielsen panel captures about 20%–25% of

respondents' overall consumption and households naturally differ in the content, prices, and frequency of their remaining consumption, we might view an R^2 of 25% as a natural upper bound. This is, in fact, the degree of explanatory power we find when we exploit within-individual variation and thus keep constant the unobserved part of the consumption bundle. Hence, our findings leave room for other, complementary determinants of subjective expectations formation (Andre et al. 2019), such as house-price experiences (Kuchler and Zafar 2019), social interactions (Bailey et al. 2018), lifetime experiences (Malmendier and Nagel 2011), cultural norms (D'Acunto 2019, 2020), gender roles (D'Acunto, Malmendier, and Weber, forthcoming), socioeconomic status (Kuhnen and Miu 2017), and heterogeneous reactions to measures of economic policy (D'Acunto, Hoang, and Weber 2021; Hanspal, Weber, and Wohlfart 2021).

At the same time, it is likely that the R^2 from the baseline analysis underestimates the true explanatory power of personal exposure to price changes for expectations, since it is estimated on survey data. Survey data tend to suffer from noise and measurement error, also as a result of rounding and heaping.¹ In fact, simulations (in the appendix, available online) reveal that plausible amounts of noise in the micro data would generate the R^2 from our baseline specifications even if personal inflation exposure fully explained inflation expectations.

To assess the extent to which noise in survey expectations might partially obscure the true explanatory power of personal exposure for inflation expectations, we follow the approach of Card and Lemieux (2001) and reestimate our baseline model on increasingly coarser samples that result from averaging the micro data within economically meaningful partitions. This methodology aims to preserve economically relevant variation in inflation expectations and consumption baskets while reducing the role of noise, rounding, and heaping in lowering the R^2 .

The first dimension we consider is households' spatial distribution, because households in the same geographic location tend to face common variation in price changes and in their economic expectations (Stroebel and Vavra 2019). We find that the R^2 of our regressions increases monotonically with the size of the geographic areas, increasing up to 66% without any demographic controls when averaging over the largest feasible cells, which correspond to US census regions.

As a second dimension, we consider consumers' cohorts or, equivalently (given the cross-sectional nature of our data), consumer age. In

¹ Heitjan and Rubin (1990) are among the first to study the implications of noise, rounding, and heaping in survey data. Jappelli and Pistaferri (2010) discuss these issues when studying consumption and income inequality using survey-based, self-reported individual data from the Survey of Household Income and Wealth (SHIW). Binder (2017) and D'Acunto et al. (2019c) document the role of rounding in the elicitation of inflation expectations through surveys.

using this dimension, we follow Aguiar and Hurst (2005) and Malmendier and Nagel (2011), who show that cohort-level experiences and consumers' age are relevant in determining spending behavior and expectations. We further subsample by education because previous research documents its influence on consumption choices and inflation expectations (D'Acunto et al. 2019c). The R^2 of our resulting regressions increases monotonically with the size of the cohort-by-education groups, up to 25% for the largest partitions for which we still have enough observations to meaningfully estimate our regressions.

These results are consistent with substantial amounts of noise being present in the micro data and indicate that heterogeneity in price exposure goes a long way toward explaining heterogeneity in inflation expectations after accounting for survey-induced noise. At the same time, we acknowledge that it is not possible to conclusively distinguish between noise and other sources of unmeasured heterogeneity. We cannot precisely estimate the extent to which the low R^2 values are due to noise versus other individual-level unobserved determinants, which the Card-Lemieux approach might also average out. Further research in macroeconomics, microeconomics, marketing, and social and cognitive psychology is needed to investigate additional micro-level determinants of inflation expectations, especially in times when the effectiveness of monetary and fiscal policies hinges on their ability to shape households' inflation expectations (D'Acunto et al. 2019b).

Related literature.—Our analysis builds on prior work that demonstrates the large heterogeneity across households, both in terms of inflation in their consumption bundles (Kaplan and Schulhofer-Wohl 2017) and in terms of inflation expectations (Bachmann, Berg, and Sims 2015). Our household-level evidence suggests that consumers interpret price changes in their bundles as signals about aggregate price changes. We also build on Cavallo, Cruces, and Perez-Truglia (2017), who study the formation of inflation expectations in high- and low-inflation countries, on the basis of recording one grocery bundle for a cohort of grocery shoppers. Our data record household-level shopping bundles for several years and multiple shopping trips, which allows us to create several measures of realized inflation at the household level and to investigate which features of price changes and consumption goods do or do not matter in the formation of household-level expectations. We also observe both the realized and the expected inflation within consumers over time, which allows us to abstract from time-invariant individual characteristics. We further build on Kuchler and Zafar (2019), who show that individuals extrapolate from local house-price changes they observe in their counties to expectations about US-wide real estate inflation.

We also relate to recent work on the determinants of cross-sectional variation in inflation expectations: Malmendier and Nagel (2016) show

that cohorts form inflation expectations on the basis of their personal lifetime aggregate inflation experiences. Other work on heterogeneity in belief formation includes that of D'Acunto et al. (2019a, 2019b, 2019c), who show that cognitive abilities are strongly correlated with forecast accuracy, uncertainty about future inflation, and responses to measures of fiscal and monetary policy. Roth and Wohlfart (2020) show that macroeconomic expectations affect personal expectations and choice. Coibion, Gorodnichenko, and Weber (2019), Coibion et al. (2020), D'Acunto, Fuster, and Weber (2020), D'Acunto et al. (2020), and D'Acunto, Hoang, and Weber (2021) show that policy communication affects inflation expectations differently across demographic groups.

II. Data on Expectations and Consumption

Our data combine the Chicago Booth Expectations and Attitudes Survey (CBEAS), which we fielded in two waves in 2015 and 2016, and the KNCP. The KNCP is a panel of about 40,000–60,000 households from 2004–18. Households report demographic characteristics as well as the prices, quantities, and shopping outlets of their consumption bundles. To avoid measurement and reporting errors, panelists use a Nielsen-provided optical scanner similar to those grocery stores use to read barcodes. The sample spans through 52 major consumer markets and nine census divisions. It records purchases of 1.5 million unique products, which include groceries, drugs, small appliances, and electronics. Nielsen estimates that the KNCP covers about 25% of US households' consumption.²

The CBEAS is a customized survey with 44 questions, which we designed in March 2015 and fielded in June 2015 and June 2016. The final sample includes 92,511 households. In the first wave, 49,383 respondents from 39,809 unique households completed the survey (43% response rate). The second wave had 43,036 unique respondents from 36,758 unique households. Of those, 15,104 participated only in wave 1, 7,269 participated only in wave 2, and 18,373 participated in both waves.³ The survey builds on the Michigan Survey of Consumers (MSC) and the New York Fed Survey of Consumer Expectations (SCE), as well as the pioneering work of Bruine de Bruin, van der Klaauw, and Topa (2011), Armantier et al. (2013), and Cavallo, Cruces, and Perez-Truglia (2017).

We first elicit demographic information that the KNCP does not provide: college major, employment status, occupation, income expectations,

² For stores where Nielsen has point-of-sales (POS) information, Nielsen uses the average price for the UPC (universal product code) during the week of purchase to minimize the data-entry burden for panelists. For stores without POS information, households report item-level gross prices, and they always report whether they used coupons or bought the item on discount.

³ The average response time was 14 minutes and 49 seconds in the first wave and 18 minutes and 35 seconds in the second wave, which included a few more questions.

and rent, mortgage, and medical expenses. We also ask for the primary shopper of the household. We then elicit perceived inflation (over the previous 12 months) and expected inflation (over the next 12 months), in terms of both point estimates and the full probability distribution.⁴

Summary statistics.—The working sample consists of 59,126 individuals for whom we observe complete data from both the KNCP and survey responses. To limit the role of outliers, we winsorize all continuous variables at the 1%–99% level.

As shown in table 1, the average age is 61, and, as in Kaplan and Schulhofer-Wohl (2017), women outnumber men. Five percent of respondents are unemployed, and almost three-quarters own a house. The average household size is 2.2. Survey respondents are more educated and wealthier than the average US individual: almost half of the respondents hold a college degree. Survey participants expect, on average, stable income over the next 12 months, with a median income bracket of USD 45,000–60,000. In terms of racial and ethnic composition, 85% of the sample is white, 8.5% Black, and 3.1% Asian.

Participants expect, on average, one-year-ahead inflation of 4.67%. Figure 1A plots the distribution of 12-months-ahead expected inflation rates. Consistent with other surveys (e.g., Binder 2017), we see substantial mass between 0% and 5% and bunching at rounded multiples of 5%. The cross-sectional dispersion is substantial, ranging from –20% to +45%. Overall, our expectations data are similar to those in the MSC and SCE.

Table A.1 (tables A.1–A.7 are available online) reports summary statistics for these variables separately for respondents who participate only in the first wave, only in the second wave, or in both waves. No substantial differences in observables exist across these groups, which suggests that observable characteristics barely explain attrition.

III. Household CPI and Frequency CPI

In this section, we study the association between household-level inflation and inflation expectations.

A. Defining Household-Level Inflation

We define household-level inflation by mimicking the CPI:

$$\text{Household CPI}_{j,t} = \frac{\sum_{n=1}^N \Delta p_{n,j,t} \times \omega_{n,j}}{\sum_{n=1}^N \omega_{n,j}}, \quad (1)$$

⁴ We randomized between two sets of questions. The MSC-inspired questions ask about the prices of things on which respondents spend money. The New York Fed SCE's questions ask specifically about inflation.

TABLE 1
SUMMARY STATISTICS

	N	Mean	Standard Deviation	Minimum	25th Percentile	Median	75th Percentile	Maximum
Age	59,118	61.4	12.9	21	54	63	70	102
Male	59,126	.36	.48	0	0	0	1	1
Unemployed	59,126	.05	.22	0	0	0	0	1
Homeowner	59,126	.74	.44	0	0	1	1	1
Household size	56,227	2.19	1.11	1	1	2	3	9
College	59,126	.48	.50	0	0	0	1	1
Income outlook (1–3)	59,126	2.18	.90	1	1	3	3	3
Economic outlook (1–5)	59,126	2.69	1.04	1	2	3	4	5
Financial outlook (1–5)	59,126	3.00	.88	1	2	3	4	5
Expected inflation	59,126	4.67	8.20	–15	0	2	6	50
Perceived inflation	59,126	4.44	8.27	–20	0	2	5	45
Household CPI	59,126	.81	7.14	–17.5	–3.17	.23	4.02	27.16
Frequency CPI	59,126	1.61	5.85	–11.71	–1.91	.83	4.21	23.08

NOTE.—This table reports summary statistics of the main independent and dependent variables for our running sample. Demographic characteristics refer to respondents who took part in the CBEAS. Income outlook, economic outlook, and financial outlook are respondents' qualitative expectations on the soundness of income growth, personal financial conditions, and overall economic outlook of the country for the next 12 months and are bounded between 1 (very bad) and 5 (very good). Expected inflation and perceived inflation are survey-reported numerical expectations and perceptions of inflation rates for a 12-month period and are bounded between –100 and +100 pp. Household CPI and Frequency CPI are the measures of household-level grocery inflation based on scanner data from the KNCP. Both measures are computed over a horizon of 12 months before the respondent took part in the CBEAS.




FIG. 1.—Grocery shopping and inflation expectations: raw data. *A* plots the distribution of inflation expectations and *B* the averages of inflation expectations across households in eight equal-sized bins by realized inflation rates in households' consumption bundles. Inflation expectations are from the customized CBEAS, fielded in June 2015 and June 2016. We use the micro data from the KNCP to create different measures of realized inflation. We use the 12 months before June of the survey wave as the measurement period and the 12 months before that period as the base period. Household CPI uses the Nielsen expenditure shares in the base periods as weights, and Frequency CPI uses the frequencies of purchase in Nielsen in the base period as weights. A color version of this figure is available online.




FIG. 2.—Time line of inflation measurement and surveys. A color version of this figure is available online.

where $\Delta p_{nj,t}$ is the log price change of good n bought by household j at time t , and $\omega_{nj} = p_{nj,0} \times q_{nj,0}$ is the weight of good n in the inflation rate for household j , with $q_{nj,0}$ being the amount of good n household j purchased in the base period. We use June 2013–May 2014 as the base period for the first survey wave and calculate price changes until the month before we fielded the first survey, that is, June 2014–May 2015. The timing varies accordingly for the second wave, fielded in June 2016 (see fig. 2).

Defining expenditure shares and price changes at the household level poses a set of conceptual and empirical challenges that do not necessarily arise in a representative-bundle setting. One such issue is seasonality in spending. We follow Kaplan and Schulhofer-Wohl (2017) and calculate volume-weighted average prices during both the base year, $p_{nj,0}$, and the year over which we measure inflation, $p_{nj,1}$. Another issue is that households might stop purchasing specific products over time. In this case, we impute entries on the basis of the price of the good at the finest geographic partition available (county, state of residence, country).⁵ All results are virtually identical if we do not impute any prices.

B. Household CPI and Inflation Expectations

Our baseline analysis estimates the following model by OLS (ordinary least squares):

$$\begin{aligned} \mathbb{E}\pi_{i,t \rightarrow t+1} = & \alpha + \beta \times \pi_{j,t-1 \rightarrow t} + X_i' \gamma + \mathbb{E}_i' \gamma + \eta_w + \eta_q \\ & + \eta_k + \eta_i + \eta_l + \epsilon_i, \end{aligned} \quad (2)$$

where $\mathbb{E}\pi_{i,t \rightarrow t+1}$ is the inflation rate individual i expects for the next 12 months, measured in percentage points; $\pi_{j,t-1 \rightarrow t}$ is the Household CPI; X_i is a vector of individual characteristics (age, age squared, sex, employment status, homeownership status, marital status, household size, college dummy, race dummies, risk tolerance), and $\mathbb{E}_i$ is a vector of expectations about household income, the aggregate economic outlook,

⁵ If we still cannot find the price, we assume no price change. The last two steps almost never arise.

and the personal financial outlook for the next 12 months. The survey-wave fixed effects η_w allow for systematic differences in (expected and realized) inflation between June 2015 and June 2016. The inflation-question fixed effects η_q allow for systematic differences in expected inflation when asked about inflation versus changes in prices. County fixed effects η_k absorb unobserved time-invariant differences across counties. Individual fixed effects η_i are included in the most restrictive specifications and absorb unobserved time-invariant differences across individuals. The income fixed effects η_I consist of the 16 income dummies from Nielsen. We cluster standard errors at the household level to allow for arbitrary correlation in residuals across respondents within household, all of whom experience the same household-level inflation.

Columns 1–3 of table 2 report the estimation results. We find a significantly positive relation between expected inflation and the Household CPI. A 1-standard-deviation increase in Household CPI is associated with a 0.17-pp increase in expected inflation, about 4% of the average expected inflation in the sample. The size of the association barely changes when we partial out a rich set of demographics, other individual expectations, and county fixed effects. The within-individual association in column 3 is slightly higher, which suggests that unobserved differences across consumers are unlikely to explain our findings. These results support the assumption in Lucas (1975) which, to the best of our knowledge, had not been formally tested with individual data.

C. *The Role of Purchase Frequency: Frequency CPI*

The Household CPI assumes that consumers weigh price changes by expenditure shares. Recent research in macroeconomics, though, proposes that price changes agents observe more often might be perceived as more precise signals (e.g., Angeletos and Lian 2016) and/or might be easier to recall. We thus test whether frequently purchased goods have a larger impact on expectations. We define a Frequency CPI using the frequency of purchase in the base period as the weight in the household's consumption basket, $\omega_{nj} = f_{nj,0 \rightarrow 1}$, where $f_{nj,0 \rightarrow 1}$ is the total quantity household j purchases of good n throughout the 12-month base period.

The distributional properties of the Frequency CPI and the Household CPI differ. Figure 1*B* sorts survey respondents into eight bins, separately for each measure, and reports average expected inflation for each bin. The resulting range in expected inflation is 0.5 pp for the Household CPI but 40% larger, 0.7 pp, for the Frequency CPI. This value is sizable, as it corresponds to about 47% of realized inflation in the United States during the period we consider.

Columns 4–6 of table 2 confirm the association from the raw data. Replicating specifications of columns 1–3 using the Frequency CPI instead of

TABLE 2
GROCERY SHOPPING AND INFLATION EXPECTATIONS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Household CPI	.171*** (4.50)	.174*** (4.59)	.192*** (2.82)	.199*** (5.19)	.221*** (5.83)	.304*** (3.40)	.046 (.77)	.014 (.24)	.070 (.78)
Frequency CPI							.164*** (2.73)	.211*** (3.56)	.243*** (2.04)
Observations	59, 126	56, 220	56, 220	59, 126	56, 220	56, 220	59, 126	56, 220	56, 220
Adjusted R^2	.028	.090	.245	.028	.091	.245	.028	.091	.245
Demographic controls		X	X		X	X		X	X
Expectation controls		X	X		X	X		X	X
County fixed effects		X	X		X	X		X	X
Individual fixed effects			X			X			X

NOTE.—This table reports OLS estimates of regressing individuals' inflation expectations on the inflation rates in their household consumption bundles. Inflation expectations are from the customized CBEAS, fielded in June 2015 and June 2016. The inflation question is randomized to ask about changes in prices (as in the MSC) or about inflation (as in the SCE). Measures of household-level inflation are constructed from the KNCP. We use the 12 months before the June of each survey wave to measure price changes and the 12 months before that period as the base period. The Household CPI uses the Nielsen expenditure shares in the base periods as weights; the Frequency CPI uses the frequencies of purchase (overall quantity) in the base period as weights; both CPIs use volume-weighted net prices (gross prices net of discounts). Demographic controls include age, square of age, sex, employment status, 16 income dummies, homeownership, marital status, household size, college dummy, four race dummies, and reported risk tolerance. Expectation controls include household income expectations, aggregate economic outlook, and personal financial outlook. All columns include survey-wave, inflation-question, and county fixed effects. Standard errors are clustered at the household level; t -statistics are in parentheses.

** $p < .05$.

*** $p < .01$.

the Household CPI, we estimate the association with inflation expectations to be 20%–50% larger. When we include both measures, in columns 7–9, the coefficient on the Household CPI shrinks toward 0 and is no longer significant. The point estimate on the Frequency CPI, instead, barely changes relative to that in columns 4–6 and remains statistically significant in all cases.

D. Robustness

These results are a robust feature of the data.⁶ They are very similar when we use changes in gross rather than net prices (minus discounts) to compute inflation (table A.2) or when we use the share of shopping trips in which an item is purchased and overweighing goods sold at higher volumes (table A.3). Neither of the alternative frequency definitions explains the cross-section of inflation expectations beyond the Frequency CPI (cols. 3 and 6 in table A.3).

We also explore the role of price changes over shorter horizons. In table A.4, columns 1–3, we include alternative CPIs that calculate household-level inflation over the prior 1, 6, and 12 months. These specifications also address concerns about reverse causality from consumers' perceptions and expectations to what to buy—consumers expecting worse times (and low inflation) buying goods with smaller price increases. Under such a mechanism, we would expect the price changes of the recently purchased goods to drive our results. Empirically, however, these price changes do not explain the cross-sectional variation of expectations conditional on the Frequency CPI.

Another aspect of the Frequency CPI that we explore is the use of average prices in the base and measurement periods to construct price changes. Although the average summarizes information about all price changes consumers observe, values such as the maximum or median might be more memorable and hence matter more in the expectations formation process. Columns 4 and 5 of table A.4 show that the changes in neither the maximum nor the median price explain expectations beyond the Frequency CPI.

A third aspect we consider is the level of granularity. The Frequency CPI defines price changes at the UPC level—the finest possible category of goods consumers observe. What if consumers think about price changes in broader categories, such as group, department, or module? Table A.5 shows that these broader categories, or using the prices at the stores instead of the ones scanned by households, do not add explanatory power.

Finally, we consider alternative weighting schemes. Columns 2–5 of table A.6 show that indices using Fisher, Paasche, or other weights do not

⁶ We thank the editor Greg Kaplan and four anonymous referees for suggesting several of the variations we study below.

add explanatory power to the baseline Frequency CPI, which follows the Laspeyres index construction.

IV. Which Price Changes Matter Most?

Our results so far reveal that the price changes consumers are exposed to most frequently help explain their inflation expectations. We now ask whether there are particular types of goods or types of price changes that matter most. We test several hypotheses that emerge from prior work, such as Cavallo, Cruces, and Perez-Truglia (2017), in particular on the sign of the price changes and the set of consumption items consumers focus on. We also show that our results remain statistically significant after we account for multiple testing.

A. *Positive Price Changes*

Cavallo, Cruces, and Perez-Truglia (2017) argue that consumers pay more attention to price increases than to price decreases. In table 3, column 1, we replace the Frequency CPI from the baseline specification with two CPIs that use only positive or only negative price changes, Positive-Price-Changes F-CPI and Negative-Price-Changes F-CPI, respectively. We find that positive observed price changes significantly influence expectations, whereas negative past observed price changes do not matter.

A similar insight emerges from the specification in column 2, in which we modify the Frequency CPI to overweigh positive price changes by a factor of 2 and a factor of 4 (Positive $\times$ 2 and Positive $\times$ 4 F-CPI). The CPI that overweighs positive changes by a larger factor drives the explanatory power of past observed inflation. We also distinguish the higher explanatory power of positive price changes from a possible role of “frequent price changes.” In column 3, we compute the Frequency CPI separately for goods whose prices displayed price volatility in households’ baskets above or below the median (High-Volatility and Low-Volatility F-CPI). Neither has explanatory power.

Finally, we take some steps to ensure that our results are not confounded by a differential persistence of positive versus negative price changes. Price increases tend to be more permanent, whereas price cuts often reflect temporary sales that revert within days or weeks (Eichenbaum, Jaimovich, and Rebelo 2011). The construction of our measures makes this explanation unlikely, since we do not use trip-to-trip price changes. Rather, we calculate the log price change between the volume-weighted price in the base period and another volume-weighted average in the observation period for each individual good.

Two additional results help to differentiate sign from persistence directly. First, we can observe whether individuals purchased goods on discounts

TABLE 3
WHICH PRICE CHANGES MATTER: SIGN AND VOLATILITY

	POSITIVE PRICE CHANGES AND VOLATILITY		
	Sign of Price Change (1)	Overweight Positive Price Changes (2)	Volatility of Price Changes (3)
Positive-Price-Changes			
F-CPI	.211*** (4.63)		
Negative-Price- Changes F-CPI	-.040 (-.84)		
Positive×4 F-CPI		.315** (2.04)	
Positive×2 F-CPI		-.078 (-.25)	
High-Volatility F-CPI			.025 (.87)
Low-Volatility F-CPI			-.039 (-.51)
Observations	56,212	56,220	49,568
Adjusted <i>R</i> ²	.042	.0042	.042
Demographic controls	X	X	X
Expectation controls	X	X	X
County fixed effects	X	X	X

NOTE.—This table reports OLS estimates of regressing individuals' inflation expectations on the inflation rates in their household consumption bundles. Inflation expectations are from the customized CBEAS, fielded in June 2015 and June 2016. The inflation question is randomized to ask about changes in prices (as in the MSC) or about inflation (as in the SCE). Measures of household-level inflation are constructed from the KNCP. We use the 12 months before the June of each survey wave to measure price changes and the 12 months before that period as the base period. The Frequency CPI employs the frequencies of purchase (overall quantity) in the base period as weights and uses volume-weighted net prices (gross prices net of discounts). The main independent variables are, in col. 1, separate indices for positive and negative price changes; in col. 2, two measures that weigh positive price changes by a factor of 4 and 2, respectively; and in col. 3, two separate Frequency CPIs based on the volatility of price changes in the Kilts-Nielsen Retail Panel. Demographic controls include age, square of age, sex, employment status, 16 income dummies, homeownership, marital status, household size, college dummy, four race dummies, and reported risk tolerance. Expectation controls include household income expectations, aggregate economic outlook, and personal financial outlook. All columns include survey-wave, inflation-question, and county fixed effects. Standard errors clustered at the household level are reported in parentheses.

** $p < .05$.
*** $p < .01$.

using coupons. As we show in table A.2, results are virtually identical to those in table 3 when we use gross rather than net prices (after discounts). Second, we follow Gorodnichenko and Weber (2016) and apply a V-shaped sales filter to the Nielsen weekly retail scanner data. That is, we compute alternative household-level CPIs when filtering temporary price changes. We exclude temporary sales if the price returns to the presale price within 1 week, 2 weeks, or 3 weeks. In these cases, we use the regular

prices to calculate realized inflation at the household level. The results, reported in columns 6–8 of table A.4, show that these alternative measures do not add any information about inflation expectations beyond the Frequency CPI.

Overall, the sign of price changes emerges as a significant factor: consumers appear to put more weight on positive than on negative price changes they observe, a feature that should be incorporated into models of expectations formation. Volatility and persistence, instead, do not appear to play a significant role in our setting.

B. Price Changes of Goods Not Purchased

Our data also allow us to consider price changes of goods that a consumer does not purchase but that are offered in the same store at the same time. Testing for the influence of such goods, though, requires a consideration set that avoids a mechanical nonresult: if we used all goods in the shopping outlet, a nonresult would be unsurprising, as consumers would not even have noticed many of them. To avoid this confound, we consider only goods that households have bought in the past. Shoppers are likely aware of their prices and, in fact, might not have purchased them because of a large, salient price increase. In column 1 of table 4, we augment our baseline model by adding an alternative definition of the Frequency CPI, the Imputation-in-Measurement-Period CPI. This measure uses the price changes of all goods the household purchased in the base period, even though they stopped purchasing such goods in the measurement period. We find that this measure does not add any additional information about inflation expectations beyond the Frequency CPI.

We also consider restricting, rather than expanding, the set of goods a household may take into account when forming beliefs about inflation. In column 2, we include a measure that restricts the Frequency CPI calculation to goods bought at least twice in the base period (Recurring-Purchases-Base CPI) and in column 3 one that restricts it to goods bought at least once in the measurement period (Purchase-in-Measurement-Period CPI). Neither alternative CPI measure has explanatory power relative to the default Frequency CPI.

V. Multiple Testing and Explanatory Power

One important concern that is typically underappreciated in economics research is the issue of multiple testing. By constructing several measures of realized inflation at the household level, we might find some being significant predictors of inflation expectations by pure chance. One common way to address the issue of specification searches or multiple testing in general is through adjustments to p -values such as the Bonferroni, Holm, and Benjamini, Hochberg, and Yekutieli adjustments.

TABLE 4
WHICH PRICE CHANGES MATTER? GOODS NOT PURCHASED IN BOTH PERIODS

	VARIATION IN SAMPLE		
	Purchased in Base Period Only (1)	Purchased at Least Twice in Base Period (2)	Purchased at Least Once in Measurement Period (3)
Frequency CPI	.212*** (5.47)	.218*** (4.51)	.229*** (5.59)
Imputation-in- Measurement- Period CPI	-.046 (-1.25)		
Recurring-Purchases- Base CPI		.024 (.52)	
Purchase-in- Measurement- Period CPI			-.017 (-.40)
Observations	51,957	56,191	56,195
Adjusted R^2	.092	.091	.091
Demographic controls	X	X	X
Expectation controls	X	X	X
County fixed effects	X	X	X

NOTE.—This table reports OLS estimates of regressing individuals’ inflation expectations on the inflation rates in their household consumption bundles. Inflation expectations are from the customized CBEAS, fielded in June 2015 and June 2016. The inflation question is randomized to ask about changes in prices (as in the MSC) or about inflation (as in the SCE). Measures of household-level inflation are constructed from the KNCP. We use the 12 months before the June of each survey wave to measure price changes and the 12 months before that period as the base period. The Frequency CPI employs the frequencies of purchase (overall quantity) in the base period as weights and uses volume-weighted net prices (gross prices net of discounts). In each specification, we propose a horse race between the Frequency CPI and a version of the Frequency CPI measured using an alternative definition. The Imputation-in-Measurement-Period CPI uses goods the consumer did not buy in the measurement period (but bought in the base period). The Recurring-Purchases-Base CPI includes only goods the consumer purchased at least twice in the base period, and the Purchase-in-Measurement-Period CPI includes only goods the consumer purchased at least once in the measurement period. Demographic controls include age, square of age, sex, employment status, 16 income dummies, homeownership, marital status, household size, college dummy, four race dummies, and reported risk tolerance. Expectation controls include household income expectations, aggregate economic outlook, and personal financial outlook. All columns include survey-wave, inflation-question, and county fixed effects. Standard errors clustered at the household level are reported in parentheses.

*** $p < .01$.

An important caveat to keep in mind in these adjustments is that none of the measures we tested were purely arbitrary but instead all our measures were motivated by theoretical reasons and findings in earlier literature, which, as Harvey, Liu, and Zhu (2016) argue, reduces the concern that our results might be driven by chance.

To directly rule out this concern, we consider the Bonferroni adjustment, which is the most conservative of the standard p -value adjustments for multiple testing. It implies rejecting the null hypothesis of no association only if the p -value of a t -test for significance of an estimated coefficient is smaller than 0.05 divided by the number of measures tested throughout the analysis. In total, we tested 21 different measures of realized inflation at the household level. Hence, any estimate with a t -statistic larger than about 3.01 would be significant at the 5% level after adjustment for multiple testing according to the Bonferroni adjustment.

The coefficients attached to our baseline measures—Frequency CPI and Household CPI—are highly statistically significant in across-individual specifications and significant at the 10% level in within-individual specifications, even with the most stringent adjustment for multiple testing and ignoring the theoretical justification for testing these very measures. Crucially, the estimate on the positive price change CPI in table 3, which is the most relevant dimension we uncover as related to inflation expectations, has a t -statistic of almost 5 and hence is highly statistically significant even after the Bonferroni adjustment for multiple testing is applied.

As a final step, we assess the explanatory power of households' personal exposure to grocery-price inflation for the observed heterogeneity in inflation expectations. In the purely cross-sectional part of our baseline regressions, the estimated R^2 amounts to less than 10%. Since the Nielsen panel captures about 20%–25% of the overall consumption bundle for the average household and since households naturally differ in their remaining consumption bundle, the prices they pay, and the frequency of purchase, we might view an R^2 of 25% as a natural upper bound. This is, in fact, the degree of explanatory power we find when we exploit within-individual variation and thus keep constant the unobserved part of the consumption bundle. Hence, our findings on the role of exposure to grocery-price changes leave room for other, complementary determinants of expectations formation, such as house-price experiences (Kuchler and Zafar 2019), social interactions (Bailey et al. 2018), or lifetime experiences (Malmendier and Nagel 2011).

At the same time, it is likely that our baseline R^2 significantly underestimates the true explanatory power of personal exposure to price changes, since it is estimated on survey data. Estimations using survey data tend to have a low R^2 even if the estimated model was correct because of noise in individually reported values and the tendency of respondents to round to integers or multiples of 5.⁷

⁷ Heitjan and Rubin (1990) are among the first to study the implications of noise, rounding, and heaping in survey data. Jappelli and Pistaferri (2010) discuss these issues when studying consumption and income inequality using survey-based, self-reported individual data from the SHIW.

In fact, we show with a simple simulation exercise (see sec. A.1 and table A.7 in the appendix) that, even if personal inflation exposure fully explained inflation expectations, implying an R^2 of 1, empirical estimations would generate an R^2 similar to that in our baseline specifications for plausible amounts of noise and rounding in the micro data. These simulation results do not mean that the lower R^2 in our baseline specifications is necessarily fully driven by noise and rounding in survey data, but they suggest that noise and rounding might indeed play a relevant role in the goodness of fit of our regressions.

To further assess the role of noise in our empirical data, we follow the approach of Card and Lemieux (2001). Their methodology relies on averaging the micro data within economically meaningful dimensions. The goal is to preserve economically relevant variation (here, in inflation expectations, consumption baskets, and good-level prices) while reducing the impact of rounding and heaping on R^2 by canceling out noisy values of opposite signs. The R^2 estimated on the coarser data would then provide for a more informative benchmark to assess the amount of cross-sectional variation in inflation expectations that is explained by household-level grocery-price changes. Of course, in these specifications we do not add additional controls to ensure that the regressions' R^2 do not increase as a result of variation of the outcome variable explained by the controls, rather than by household-level inflation.⁸

The first dimension we consider is households' geographic location. This analysis builds on work by Stroebel and Vavra (2019), who find that households in the same geographic location tend to face commonality in price changes and display comoving economic expectations. Moreover, geographic splits provide aggregation of partitions with different levels of granularity that are fully contained within each other, which allows us to average out more and more noise as we move to coarser partitions, but still maintaining the same meaningful geographic-level variation within partitions.

In table 5, we collapse the individual-level data within geographic cells whose size increases moving to the right: zip code, county, three-digit FIPS (Federal Information Processing Standard) code, state, and US census region. The three-digit FIPS code is assigned to counties within each state, and the same codes are used across all 50 states. Thus, this partition creates groups of counties that belong to different states.⁹ We find that, when moving from the finest to the broadest geographic partition, the R^2 increases monotonically, consistent with substantial amounts of noise

⁸ All the estimated coefficients are quantitatively and qualitatively similar if we add averages of the demographic controls from table 2 at the level of each partition.

⁹ This specific collapse of the data allows us to verify that the averaging of noise, rather than common geographic shocks, explains the increase in R^2 .

TABLE 5
GOODNESS OF FIT: GEOGRAPHIC AND COHORT-BY-EDUCATION PARTITIONS

	GEOGRAPHIC PARTITIONS				COHORT-BY-EDUCATION PARTITIONS			
	Zip Code (1)	County (2)	3-Digit FIPS code (3)	State (4)	Census Region (5)	Birth Month and College (6)	Birth Quarter and College (7)	Birth Year and College (8)
Frequency CPI	.209*** (4.15)	.322*** (2.85)	.493*** (2.75)	.162* (1.66)	.296** (1.98)	.194* (1.74)	.236* (1.85)	.409* (1.93)
Observations	21,177	4,452	472	98	18	1,753	617	160
Adjusted R^2	.028	.021	.058	.331	.656	.029	.032	.247

NOTE.—This table reports the estimates of regressing inflation expectations averaged at the geographic and cohort-by-education levels reported on top of each column and across survey waves on the average inflation rates faced by households who are part of such partitions. Inflation expectations are from the customized CBE/AS, fielded in June 2015 and June 2016. The inflation question is randomized to ask about changes in prices (as in the MSC) or about inflation (as in the SCE). Measures of household-level inflation are constructed from the KNCP. We use the 12 months before the June of each survey wave to measure price changes and the 12 months before that period as the base period. The Frequency CPI employs the frequencies of purchase (overall quantity) in the base period as weights, and uses volume-weighted net prices (gross prices net of discounts). Zip code, county, FIPS code, state, and census region represent the geographic partitions within which we average variables. The three-digit FIPS codes are assigned in alphabetical order on the basis of county name within each state. This partition thus creates groups of counties that belong to different states. For the cohort-by-education partitions, birth month, birth quarter, and birth year represent the cohort definitions by which we average variables. Within each cohort partition, we further average variables separately for respondents with and without college education, which produces two observations for each cohort. All columns include averages of the survey-wave and inflation-question dummies at the partition level on the right-hand side. Huber-White standard errors are reported in parentheses.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

in the micro data. With the maximum amount of noise averaged out (collapsing at the level of US census regions), we obtain an R^2 of up to 65.6%.

As a second dimension, we consider consumers' cohorts or, equivalently (given the cross-sectional nature of our data), consumers' age. In using this dimension, we follow Aguiar and Hurst (2005) and Malmendier and Nagel (2011), who show that cohort-level experiences and consumers' age are relevant in determining spending behavior and expectations. We also note that, within a cohort/age group, observable dimensions such as education and cognitive abilities generate systematic differences in the composition of consumption bundles and in the formation of economic expectations (D'Acunto et al. 2019c). We therefore include aggregations of the data at the cohort-by-education level—within each cohort group, we aggregate the data separately for cohort members who hold a college degree and those without a college degree.

Columns 6–8 of table 5 reveal that the R^2 of our regressions increases monotonically with the size of the cohort-by-education groups. It amounts to 24.7% for the largest partitions for which we still have enough observations to meaningfully estimate the empirical model. Note that this partition (col. 8) is based on 160 cohort-by-education observations, which is a number of observations similar to that for the state-level partition in column 4, and the size of the R^2 in these two partitions is similar.

Overall, as we aggregate across larger partitions, the R^2 s of our regression models increase, which, on the basis of the approach in Card and Lemieux (2001), indicates that the low R^2 in regressions on the individual-level micro data might be driven by a substantial amount of noise, which then gets averaged out at the partition level. At the same time, it remains possible that unexplained individual heterogeneity that is orthogonal to both geography and age and education will also be averaged out. Although the Card and Lemieux (2001) strategy cannot distinguish between noise and unexplained heterogeneity, the robustness of our findings across partitions points to the well-known role of survey noise as the key factor.

VI. Conclusions

We document that household-level grocery-price changes significantly shape inflation expectations. We use unique, representative US data that link individual expectations to items purchased, frequency and outlet of purchase, and prices paid. These rich data also reveal which features of observed price changes matter in the formation of inflation expectations—the frequency of purchase and the positive sign of price changes—which informs advances in heterogeneous-beliefs models.

We focus on inflation expectations not only to test the Lucas (1975) assumption but also because understanding how households form inflation expectations is especially important in times of low interest rates

and inflation (Summers 2018). Under these conditions, traditional monetary-policy tools are unviable, and managing households' inflation expectations directly is a key form of unconventional monetary and fiscal policy to stimulate aggregate demand (Feldstein 2002; Yellen 2016; D'Acunto, Hoang, and Weber 2018; Lagarde 2020). Our results motivate additional work to further understand how consumers form aggregate expectations about inflation and other macroeconomic variables, as well as how these expectations feed into economic and financial decision-making.

Future work should also aim to understand how price changes in the nongrocery part of households' bundles interfere with grocery-price changes. Another fruitful avenue for research is understanding how the inflationary environment in which consumers form expectations interacts with the role of personally observed price changes. For instance, is it optimal for consumers to focus on personal shopping exposure when forming expectations in a stable inflation environment but to shift the focus on aggregate inflation in volatile times, as Frache and Lluberá (2018) suggest using firms' inflation expectations? The extent to which the increasing substitution of in-store shopping with online shopping affects the role of personal inflation on inflation expectations is also an interesting direction for future research.

References

- Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *J.P.E.* 113 (5): 919–48.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2019. "Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample." Working paper. <http://dx.doi.org/10.2139/ssrn.3355356>.
- Angelesos, G.-M., and C. Lian. 2016. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." In *Handbook of Macroeconomics*, vol. 2A, edited by John B. Taylor and Harald Uhlig, 1065–240. Amsterdam: North-Holland.
- Armantier, Olivier, Wändi Bruine de Bruin, Simon Potter, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2013. "Measuring Inflation Expectations." *Ann. Rev. Econ.* 5:273–301.
- Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims. 2015. "Inflation Expectations and Readiness to Spend." *American Econ. J.: Econ. Policy* 7 (1): 1–35.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018. "The Economic Effects of Social Networks: Evidence from the Housing Market." *J.P.E.* 126 (6): 2224–76.
- Binder, Carola C. 2017. "Measuring Uncertainty Based on Rounding: New Method and Application to Inflation Expectations." *J. Monetary Econ.* 90:1–12.
- Bruine de Bruin, Wändi, Wilbert van der Klaauw, and Giorgio Topa. 2011. "Expectations of Inflation: The Biasing Effect of Thoughts about Specific Prices." *J. Econ. Psychology* 32 (5): 834–45.

- Card, David, and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." *Q.J.E.* 116 (2): 705–46.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices." *American Econ. J.: Macroeconomics* 9 (3): 1–35.
- Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber. 2020. "Forward Guidance and Household Expectations." Working Paper no. 26778 (February), NBER, Cambridge, MA.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2019. "Monetary Policy Communications and Their Effects on Household Inflation Expectations." Working Paper no. 25482 (January), NBER, Cambridge, MA.
- D'Acunto, Francesco. 2019. "Identity and Choice under Risk." Working paper.
- . 2020. "Tear Down This Wall Street: Anti-market Rhetoric, Motivated Beliefs, and Investment." Working paper.
- D'Acunto, Francesco, Andreas Fuster, and Michael Weber. 2020. "Diverse Policy Committees Are More Effective." Res. Paper 2020–38, Booth School Bus., Chicago.
- D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2019a. "Cognitive Abilities and Inflation Expectations." *AEA Papers and Proc.* 109:562–66.
- . 2019b. "Human Frictions to the Transmission of Economic Policy." Working paper.
- . 2019c. "IQ, Expectations, and Choice." Working Paper no. 25496 (September), NBER, Cambridge, MA.
- . 2020. "Effective Policy Communication: Targets versus Instruments." Res. Paper 2020–148, Booth School Bus., Chicago.
- D'Acunto, Francesco, Daniel Hoang, and Michael Weber. 2018. "Unconventional Fiscal Policy." *AEA Papers and Proc.* 108:519–23.
- . 2021. "Managing Households' Expectations with Unconventional Policies." *Rev. Financial Studies*, forthcoming.
- D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber. Forthcoming. "Gender Roles and the Gender Expectations Gap." *Proc. Nat. Acad. Sci. USA*.
- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference Prices, Costs, and Nominal Rigidities." *A.E.R.* 101 (1): 234–62.
- Feldstein, Martin. 2002. "The Role for Discretionary Fiscal Policy in a Low Interest Rate Environment." Working Paper no. 9203 (September), NBER, Cambridge, MA.
- Frache, Serafin, and Rodrigo Lluberas. 2018. "New Information and Inflation Expectations." Working Paper no. 781, Bank Internat. Settlements, Basel.
- Georganas, Sotiris, Paul J. Healy, and Nan Li. 2014. "Frequency Bias in Consumers' Perceptions of Inflation: An Experimental Study." *European Econ. Rev.* 67:144–58.
- Gorodnichenko, Yuriy, and Michael Weber. 2016. "Are Sticky Prices Costly? Evidence from the Stock Market." *A.E.R.* 106 (1): 165–99.
- Hanspal, Tobin, Annika Weber, and Johannes Wohlfart. 2021. "Exposure to the COVID-19 Stock Market Crash and Its Effect on Household Expectations." *Rev. Econ. and Statis.*, forthcoming.
- Harvey, Campbell R., Yan Liu, and Heqing Zhu. 2016. "... and the Cross-Section of Expected Returns." *Rev. Financial Studies* 29 (1): 5–68.

- Heitjan, Daniel F., and Donald B. Rubin. 1990. "Inference from Coarse Data via Multiple Imputation with Application to Age Heaping." *J. American Statis. Assoc.* 85 (410): 304–14.
- Jappelli, Tullio, and Luigi Pistaferri. 2010. "Does Consumption Inequality Track Income Inequality in Italy?" *Rev. Econ. Dynamics* 13 (1): 133–53.
- Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. "Inflation at the Household Level." *J. Monetary Econ.* 91:19–38.
- Kuchler, Theresa, and Basit Zafar. 2019. "Personal Experiences and Expectations about Aggregate Outcomes." *J. Finance* 74 (5): 2491–542.
- Kuhnen, Camelia M., and Andrei C. Miu. 2017. "Socioeconomic Status and Learning from Financial Information." *J. Financial Econ.* 124 (2): 349–72.
- Lagarde, Christine. 2020. "ECB President Press Conference." June 4. European Central Bank. <https://www.ecb.europa.eu/press/pressconf/2020/html/ecb.is200604~b479b8cfff.en.html>.
- Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *J. Econ. Theory* 4 (2): 103–24.
- . 1973. "Some International Evidence on Output-Inflation Tradeoffs." *A.E.R.* 63 (3): 326–34.
- . 1975. "An Equilibrium Model of the Business Cycle." *J.P.E.* 83 (6): 1113–44.
- Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?" *Q.J.E.* 126 (1): 373–416.
- . 2016. "Learning from Inflation Experiences." *Q.J.E.* 131 (1): 53–87.
- Roth, Christopher, and Johannes Wohlfart. 2020. "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?" *Rev. Econ. and Statis.* 102 (4): 731–48.
- Stroebel, Johannes, and Joseph Vavra. 2019. "House Prices, Local Demand, and Retail Prices." *J.P.E.* 127 (3): 1391–436.
- Summers, Lawrence H. 2018. "The Threat of Secular Stagnation Has Not Gone Away." <http://larrysummers.com/2018/05/06/the-threat-of-secular-stagnation-has-not-gone-away/>.
- Yellen, Janet L. 2016. "Macroeconomic Research after the Crisis." Speech no. 915, October 14, Board of Governors, Fed. Reserve System.