Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment*

Luis Armona, Andreas Fuster, and Basit Zafar†

February 27, 2018

Abstract

Home price expectations are believed to play an important role in housing dynamics, yet we have limited understanding of how they are formed and how they affect behavior. Using a unique "information experiment" embedded in an online survey, this paper investigates how consumers' home price expectations respond to past home price growth, and how they impact investment decisions. After eliciting respondents' priors about past and future local home price changes, we present a random subset of them with factual information about past (one- or five-year) changes, and then re-elicit expectations. This unique "panel" data allows us to identify causal effects of the information, and provides insights on the expectation formation process. We find that, on average, year-ahead home price expectations are revised in a way consistent with short-term momentum in home price growth, though respondents tend to underpredict the strength of momentum. Revisions of longer-term expectations show that respondents do not expect the empirically-occurring mean reversion in home price growth. These patterns are in line with recent behavioral models of housing cycles. Finally, we show that home price expectations causally affect investment decisions in a portfolio choice experiment embedded in the survey.

Keywords: housing, expectation formation, information, updating

*The paper has greatly benefited from the comments of the editor (Nicola Gennaioli) and four referees. We also thank John Beshears, Richard Crump, Adam Guren, Ida Johnsson, Theresa Kuchler, Charles Nathanson, Robert Shiller, Johannes Stroebel, and participants at numerous conferences and seminars for helpful discussions and comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

†Armona: Department of Economics, Stanford University; E-mail: luisarmona@gmail.com. Fuster: Research and Statistics, Federal Reserve Bank of New York; E-mail: andreas.fuster@ny.frb.org. Zafar: Department of Economics, W.P. Carey School of Business, Arizona State University; E-mail: basitak@gmail.com

1 Introduction

Home price expectations play a prominent role in many accounts of the housing boom that occurred during the early- to mid-2000s, both in the US and globally (e.g. Shiller 2005, Foote et al. 2012, Glaeser et al. 2013, Kaplan et al. 2017). Beyond this particular episode, home prices display patterns such as strong momentum at a relatively short horizon (e.g. Case and Shiller 1989, Guren 2016) and mean reversion at a longer horizon (e.g. Cutler et al. 1991, Glaeser et al. 2014) that researchers in this area have found challenging to explain within a fully rational framework. As a consequence, in recent years there has been increasing interest in exploring theories of home price expectations that, to varying degrees, depart from full rationality and instead feature some form of extrapolation from recent growth.[1] However, so far there exists very little direct empirical evidence on home price expectations that such theories could be validated against.

In this paper, we present new evidence on how home price expectations are formed, and how they affect behavior. Specifically, we rely on a novel "information experiment" within an online household survey to test how respondents update their expectations about future home price growth in their local area when they are provided with objective information about recent home price growth. We furthermore embed an incentivized portfolio choice experiment in the survey that enables us to study the causal effect of expectations on a housing-related investment decision.

The survey has three main stages. In the first stage, respondents are asked about their perceptions of home price changes in their local area over the past one and five years, and about their expectations of future local home price changes over the next one and five years. Individuals also make a hypothetical decision on an investment with payoffs linked to future local home price changes— specifically, respondents are asked how they would allocate $1,000 between a housing market fund with returns tied to local year-ahead house price growth and a risk-free savings account. In the intermediate stage, respondents are randomly exposed to either objective information about actual local home price changes over the past one year, or over the past five years, or no information (control group). In the final stage, future home price expectations are re-elicited, and respondents are again presented with the investment decision, which is now incentivized.

This empirical design allows us to study two main questions. First, we test *whether* and *how* respondents revise their expectations after being provided with information that may differ from their

[1] Work in this vein includes Adam et al. (2012), Burnside et al. (2016), DeFusco et al. (2017), Gao et al. (2015), Gelain and Lansing (2014), Glaeser and Nathanson (2017), Granziera and Kozicki (2015), Guren (2016), and Piazzesi and Schneider (2009); see Glaeser and Nathanson (2014) for a review.

priors about recent home price growth in their local area. For instance, if a respondent thought that prices had increased by 3% over the past year, and expects them to increase by 2% over the coming year before we provide her with the information, how does she react after learning that according to a house price index (HPI), prices had in fact increased by 6%? If she believes in momentum in house prices, we would expect her to adjust her expectations of future growth upward, while a belief in mean reversion would lead her to revise her expectations downward. On the other hand, if she believes that home prices follow a random walk, there should be little systematic revisions in response to the information. Second, the investment decision allows us to investigate whether home price expectations are linked with (hypothetical and actual) behavior; the panel aspect of our design allows us to study this link both in the cross-section as well as within individual.

Our design has several advantages over alternative approaches. In general, investigating expectation formation faces the issue that neither the true nor the perceived data generating processes, nor the individuals' information set, are known. Furthermore, expectations may be measured with error, and the way in which they are elicited may matter. Our study advances the literature by moving beyond documenting correlations between past home price changes and subjective home price expectations (e.g. Case et al. 2012, Kuchler and Zafar 2015), which could be partly driven by omitted variables. By directly manipulating individuals' information set, we can provide a causal interpretation to the relationship between past changes and expectations, and estimate the strength of this relationship more cleanly than from cross-sectional data alone. Our design does not rely on any assumptions on either the respondents' perceived data generating process or prior information set; instead, we generate rich quantitative data on priors and expectation updating over different horizons that models can be quantitatively compared against. Relatedly, by studying how updating varies with respondents' personal characteristics or across locations, we shed further light on different theories of belief formation. While some measurement error is surely present in our data, we have several ways of ensuring that neither measurement error nor a particular way of eliciting priors and expectations drive our qualitative conclusions.

We find that, when provided with information about past year local home price growth, respondents on average update their year-ahead local home price expectations in an extrapolative manner: for each percentage point underestimation (overestimation) of past growth relative to the HPI, respondents adjust their expectations upward (downward) by 0.20 percentage points. In contrast, information about price growth over the previous five years has no significant effect on revision of year-ahead expectations (although directionally respondents also extrapolate).

A natural question to ask is how these findings compare with actual serial correlation in home price growth. Home price growth exhibits strong positive autocorrelation at the one-year horizon (Case and Shiller, 1989). The coefficient of a regression of local one-year home price growth on lagged one-year growth, averaged across the zip codes in our sample, is a precisely estimated 0.53; the coefficient in the case of one-year growth regressed on lagged five-year growth is 0.14 (and imprecisely estimated). Thus, over the short horizon, the average respondent directionally updates in a "rational" manner, that is, one that is consistent with data. However, the average respondent tends to *underreact* to past growth, given the strong short-term momentum in actual home prices.

A different picture emerges in the case of medium-term expectations (which we define as expectations for the two-to-five year horizon). Home prices tend to exhibit mean reversion over longer horizons. However, in our experiment, respondents tend to update their medium-term expectations in an extrapolative manner (though with smaller estimated effect sizes than at the one-year horizon). In addition, we do not find evidence of revisions being systematically related with respondents' baseline confidence or uncertainty in their priors, as Bayesian updating would predict. Thus, our findings appear most consistent with "behavioral" models of housing market dynamics; for instance, the estimated effect sizes are close to the calibration in Glaeser and Nathanson (2017). From a broader perspective, these patterns support the view of extrapolation or an underappreciation of mean reversion as a potentially important driver of fluctuations in financial markets (e.g. DeLong et al. 1990, Barberis et al. 1998, Barsky and DeLong 1993, Fuster et al. 2012, Barberis et al. 2015, Bordalo et al. 2017a,b).

We also study heterogeneity in updating behavior. Treatment respondents (that is, those who receive the information) are more likely to update their expectations than a control group. Conditional on updating, treatment respondents are much more likely to be "extrapolators" (revising their expectations in the direction of the gap between revealed past HPI growth and their prior about it) than to be "mean reverters" (doing the opposite) for expectations at both horizons. We find mixed evidence for models of age-dependent updating (Malmendier and Nagel, 2016): younger respondents and those with shorter tenures in their locality are not more likely to update than their counterparts. However, conditional on updating, they are much more likely to be extrapolators. Perhaps our most intriguing result is that individuals residing in areas with inelastic housing supply (or with stronger long-term mean reversion in home prices) exhibit a higher propensity to extrapolate from past growth at both horizons we study. This is arguably rational behavior at the shorter horizon (since inelastic areas tend to have stronger momentum), but not for the longer horizon.

Turning to our second question of how expectations affect behavior, we find that expectations have an economically and statistically significant effect on respondents' investment allocation, both across respondents and within-respondent (meaning the change in the housing fund share between the hypothetical and incentivized rounds is related to the change in expectations following the information provision). Outside the stylized investment experiment, we also document significant correlations between respondents' baseline expectations and various intended housing-related behaviors. These findings suggest that survey measures of house price expectations contain meaningful information to understand behavior, and are therefore important variables to track for policy makers and housing market analysts.

While the survey design is discussed in more detail later in the paper, we point out a few noteworthy features here. First, we randomize our respondents into different question "frames" when eliciting their perceptions and expectations to ensure that our results are not exclusive to a given frame. Specifically, half the respondents are asked for their perceptions and forecasts in terms of house price levels (from which we then calculate percent changes) while the others are directly asked about percent changes. Our main results hold within both frames. Second, the information provision (and re-elicitation of expectations) does not happen immediately after the respondents' priors are elicited, but only after they have gone through various other (unrelated) survey questions. This makes it unlikely that the effects of the information are driven by "demand effects" or a desire to give the "correct" answer. Our design also features a control group that is not provided with information, so that we can account for the effects on expectations that merely completing the survey may have. Third, we test whether the information provision has persistent effects on our respondents' beliefs by re-eliciting them in a separate follow-up survey two months after the initial one. We find that indeed, the average effect of the information on short-term expectations remains almost the same as within the main survey.

The empirical design in this paper is closest to that used in a recent literature that employs information experiments in surveys to understand expectation formation (Armantier et al. 2016a, Cavallo et al. 2014, and Coibion et al. 2015). The actual dependence in home prices (and the regional variation in it) provides us with a natural benchmark against which we can evaluate the updating patterns of our respondents. The information experiment in our survey is also related to other experimental work in lab settings (e.g. Schmalensee 1976, Haruvy et al. 2007, Rötheli 2010, and Beshears et al. 2013).

Our work further relates directly to other survey-based studies on expectation formation. In the housing market, Case and Shiller (2003) and Case et al. (2012) measure expected future home price growth in a sample of recent homebuyers across four metropolitan areas, finding evidence consistent with extrapolation at one-year and ten-year forecast horizons. Niu and van Soest (2014) and Ma (2016) study home price expectations in the American Life Panel and the Michigan Survey of Consumers, respectively, while Bover (2015) conducts a similar exercise in Spanish data. Kuchler and Zafar (2015) study how experienced local home price growth (as measured by a HPI) affects expectations about future national home price growth.[2] Our approach is unique in that we directly measure respondents' perceptions of recent local home price growth and test whether changing this perception through information provision affects future expectations. Other work has used surveys to study the properties of stock market expectations (e.g. Vissing-Jorgensen 2004, Amromin and Sharpe 2014, Greenwood and Shleifer 2014) and inflation expectations (e.g. Malmendier and Nagel 2016, Madeira and Zafar 2015). Gennaioli et al. (2016) present evidence that corporate CFOs' expectations of future earnings growth are extrapolative, and affect firm behavior.[3]

The remainder of the paper is organized as follows: the next section describes the design of the survey, how it was administered, and details about the respondent sample. In order to provide a benchmark for our experimental setting, Section 3 analyzes the dependence in actual home prices over different horizons. Section 4 characterizes respondents' perceptions and expectations at the baseline (prior to the information provision). Section 5 presents the experimental results of the effects of information on expected future home price growth. Section 6 studies the effect of expectations on behavior, and Section 7 concludes.

2 Survey Design and Administration

Our data come from two original online surveys, both fielded as part of the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE). The SCE is an internet-based survey of a rotating panel of approximately 1,200 household heads from across the US, with the goal of eliciting expectations about a variety of economic variables, such as inflation and labor market conditions. Respondents participate in the panel for up to twelve months, with a roughly equal number rotating

[2] Bailey et al. (2017) study how a qualitative survey measure of the attractiveness of housing as an investment is affected by the home price experiences of (out-of-town) friends, and also how these experiences affect housing-related behavior; we return to this study in the conclusion.

[3] A large literature in macroeconomics studies the role of information frictions in expectation formation (see e.g. the sticky/noisy information models of Mankiw and Reis 2002 or Sims 2003). In our setting we exogenously provide (differential) information to our respondents, thereby alleviating these frictions for them. Coibion and Gorodnichenko (2012, 2015) present a unifying framework to empirically test and distinguish between several sticky/noisy information models based on *aggregate* forecast data from surveys; these tests would not be applicable to our individual-level data.

in and out of the panel each month. Respondents are invited to participate in at least one survey each month.[4]

The first and main survey is a special module on housing, fielded in February 2015. Active panel members who had participated in a SCE monthly survey in the prior eleven months were invited to participate in the housing module. Out of a total sample of 1,383 household heads on the panel that were invited, 1,205 participated, implying a response rate of 87%.

The housing module contains multiple blocks of questions, some differing between owners and renters. The respondents are asked, among other things, about their perceptions of past local home price changes and expectations for future local home price changes, (current and future) financing conditions, past housing-related behavior (such as buying a home, and housing debt), and the future likelihood of buying a home. Respondents also provide information about their zip code location, their household income, and many other demographic variables. The median survey time was 34 minutes, with owners having a median completion time 7 minutes higher than renters, since they answered many more questions. When appropriate, questions had built-in logical checks (for instance, percent chances of an exhaustive set of events had to sum to 100). Item non-response is extremely rare, and almost never exceeds one percent for any question.

The second survey is the regular monthly SCE survey, and was fielded during April 2015. Respondents of the housing module who still remained in the SCE rotating panel were invited to participate in a short follow-up module. Of the 978 household heads still in the panel, 856 did so, for a repeat response rate of 87.5%.

2.1 Survey Design

We next describe the relevant sections of the two surveys.

The experimental setup in the first survey consisted of three stages:

1. **Baseline Stage**: The first stage elicited respondents' perceptions about home price changes in their zip code over the past 12 months and the past 5 years. We also elicited respondents' expectations regarding home price changes in their zip code over the next 12 months, and the next 5 years (the precise questions will be discussed below).[5] In addition, respondents

[4]The survey is conducted over the internet by the Demand Institute, a non-profit organization jointly operated by The Conference Board and Nielsen. The sampling frame for the SCE is based on that used for The Conference Board's Consumer Confidence Survey (CCS). Respondents to the CCS, itself based on a representative national sample drawn from mailing addresses, are invited to join the SCE internet panel. The response rate for first-time invitees hovers around 55%. Respondents receive $15 for completing each survey. See Armantier et al. (2016b) for additional information.

[5]Furthermore, respondents were asked to rate the attractiveness of housing in their zip code as a financial investment

were presented with a hypothetical investment scenario where they were asked to allocate \$1,000 between a fund indexed to year-ahead home price growth in their local area, and a 2% risk-free interest savings account.[6]

2. **Treatment Stage**: A block of other housing-related questions taking roughly 15 minutes separated the baseline and treatment stages. In the treatment stage, respondents were randomly assigned to one of three groups:

    - *1-year Treatment ("T1")*: Respondents were informed about the percentage change in home prices in their zip code over the 2014 calendar year. This information was based on the Zillow Home Value Index (ZHVI), which is freely available online.[7]
    - *5-year Treatment ("T5")*: Respondents were informed about the total percentage change in home prices in their zip code over the past 5 years, from the beginning of 2010 to the end of 2014.
    - *Control group*: Respondents in this group got no information on past home price changes.

3. **Final Stage**: This stage followed right after the treatment stage. All survey respondents were re-asked their expectations of zip code level home price changes at the one- and five-year horizons—the same forecast horizons for which expectations were initially elicited at the baseline stage. The investment scenario that respondents had seen in the first stage was also presented again. It was identical to the initial scenario, except that the decision was now incentivized—respondents were informed that two people taking the survey would be paid in a year's time depending on the return of their investments.[8]

on a 1-5 scale. We analyze this question in Appendix A.6.

[6]The exact question was: "*Consider a situation where you have to decide how to invest \$1,000 for one year. You can choose between two possible investments. The first is a fund that invests in your local housing market, and pays an annual return equal to the growth in home prices in your area. The second is a savings account that pays 2% of interest per year. What proportion of the \$1,000 would you invest in (1) the housing market fund, (2) the savings account?*"

[7]Respondents were shown the following: *"Zillow is one of the best-known sources of information about home prices. According to Zillow.com, home prices in your zip code during 2014 [increased/decreased] by [X]%."* (X, the respondent-specific local home price change, was shown with one decimal place.)

For more information on the construction of the ZHVI, see http://www.zillow.com/research/zhvi-methodology-6032/ (accessed on September 6, 2017). We use the ZHVI as of January 2015, the month prior to the survey.

The coverage of ZHVI is incomplete at the zip code level, so if we do not have zip code level information, we use the state-level ZHVI change (respondents were told *"In cases where zip code level information is not available, we use the state-level change in home prices (or, in very few cases where no state-level information is available, the national change)."* 70.3% of our respondents' reported zip codes were covered by the ZHVI. In the very rare cases where we do not have state-level data (Maine and Kansas), we report national changes; 14 of our 1,205 (1.16%) respondents were in this category.

[8]Paying only a randomly chosen subset of respondents is commonly done in large-scale economic experiments (e.g. Dohmen et al., 2011) and has been found to generate very similar behavior to paying all subjects (March et al., 2016). Respondents were told: *"Note that you have a chance of earning extra money by answering this question. At the end of the month,*

The follow-up questions were fielded to respondents in the April 2015 SCE monthly survey. Respondents were asked their expectations of zip code level home price changes at the one- and five-year horizons.

Some features of the study design merit further discussion. We include treatments that provide information on short- and longer-term home price changes since home price changes tend to exhibit momentum in the short term and mean reversion over a longer horizon, as will be discussed in the next section. The reason for including a control group was that the simple act of taking a survey about housing may make respondents think more carefully about their responses, and may lead them to revise their home price expectations even if they are not provided with any new information (see Zwane et al. 2011 for a discussion of how surveying people may change their subsequent behavior). Since we are interested in revisions in expectations that are directly attributable to the information, we identify them from differences between the treatment and control groups' changes in expectations. The investment task allows us to investigate, in a direct fashion, whether home price change expectations impact both hypothetical and incentivized behavior, in the cross-section as well as at the individual level. Finally, the follow-up survey allows us to test whether the effect of the treatment, if any, persists beyond the initial survey horizon.

Home price perceptions (for the past one and five years) and expectations (for one and five years ahead) were elicited in two different formats. All respondents were first asked for the dollar value of a typical home in their zip code today. Each respondent was then randomly assigned to one of two "frames" which determined how the questions about the past and future were asked:[9]

- **(L)evel-frame**: The perception and expectation questions were asked in terms of house price *levels*. For example, past one year home price change perceptions were elicited as follows: "*You indicated that you estimate the current value of a typical home in your zip code to be [X] dollars. Now, think about how the value of such a home has changed over time. (By value, we mean how much that typical home would approximately sell for.). What do you think the value of such a home was one year ago (in February 2014)?*" We refer to this frame as the L-frame.

we will randomly pick 2 survey participants. These 2 participants will be paid in Spring 2016 according to the investment choice they made (that is, the $1,000 and the return on their choices). If you are chosen, your payment will depend on how you had invested the money, so answer this question carefully.

To determine the return on the housing market fund, we will use the Zillow home price index for your current zip code. In cases where zip code level information is not available, we use the state-level index (or, if that is not available, the national index)."

Respondents were not informed of how many total respondents there were. Two of them were drawn at random in early March 2015 (after all respondents had taken the first survey) and notified that they had been picked to be paid; the payments were made in April 2016.

[9]This randomization was orthogonal to the randomization into treatment (T1, T5, or Control). Each respondent remained in the same frame throughout both surveys.

- **(C)hange-frame**: The perception and expectation questions were asked in terms of *percent changes*. For example, when eliciting past one year home price change perceptions, respondents were first asked if they thought home prices had increased or decreased over the past one year, and next asked for the percentage change: "*By about what percent do you think the value of such a home has [increased/decreased] over the past 12 months? Please give your best guess.*" We refer to this as the C-frame in the analysis.

These two approaches for eliciting perceptions and expectations were motivated by the finding of Glaser et al. (2007) that survey respondents' predictions of stock performance are influenced by whether they are asked to forecast future returns or future price levels. In the former case, expectations appear to be extrapolative, whereas when asked for levels, respondents appear to believe in mean reversion. We therefore want to study whether our findings are robust to the elicitation mode. In our analysis, we control for the frame assignment whenever the analysis is done on the full sample. For our main results on expectations, we also discuss how findings differ across frames.

Respondents, at the baseline stage, were also asked about their subjective distribution for both one- and five-year ahead home price growth. In the case of one-year ahead expectations, for example, respondents were asked to assign probabilities to four intervals that future year-ahead home price changes may lie in (less than -5%; between -5% and 0%; between 0% and 10%; more than 10%). We use this to measure the respondent's belief of downside risk in home prices, and to construct a measure of the respondent's prior uncertainty.

In order to reduce the importance of outliers and to screen out individuals who arguably did not take the survey seriously, the analysis in the paper removes respondents with extreme observations for our key variables: baseline perceptions of price changes over the last 12 months and past five years, and baseline as well as final stage home price expectations over the two horizons. Specifically, respondents who report answers in the top and bottom 2% of the response distribution for those variables are dropped. In addition, we drop 12 respondents who provide a response of less than $10,000 for the value (today; in the past, or; in the future) of a typical home in their zip code. This leaves us with 1,020 individuals (from a total of 1,205 respondents who took the survey). Our results are qualitatively similar if we trim observations at 1% or 5%, or if we instead winsorize extreme responses.

2.2 Sample Characteristics

The first column of Table 1 displays the demographic characteristics of our sample. The sample aligns well with average demographic characteristics of the United States along most dimensions. For instance, the average age of our respondents is 50.4 years, and 52.9% of them report annual household income of less than $60,000, while the corresponding numbers among US household heads are 53.7 years and 54.5%.[10] 74.3% of respondents are homeowners, compared to a national homeownership rate in 2015:Q1 of 63.7% according to the Census. One notable divergence between our sample and the US population is in education. Our sample is significantly more educated than the overall population: 55% of our respondents have at least a Bachelors' degree, while only a third of the US household heads fall in this category. This may partly be a result of differential internet access and computer literacy across education groups in the US population.

The table also shows some other demographic variables, such as labor force status, tenure in the respondent's town or city, and numeracy.[11] Columns (2)-(4) of the table show that the demographic characteristics are not statistically different across the three treatment groups (the only exception being the proportion of males). This should not be surprising, since random assignment should have largely preserved balance between the three groups.

The last column of Table 1 shows the characteristics of the follow-up sample (excluding respondents who are removed based on being outliers in the initial survey). We also conduct (unreported) pairwise tests for the equality of the means of characteristics for the follow-up sample (column 6) and the initial sample (column 1). There are no significant differences in observables across the two surveys, meaning there is no evidence for selection on observables into the follow-up survey.

3 Dependence in Actual Home Price Changes

Before turning to the empirical analysis, it is useful to investigate the actual dependence in home prices over different horizons. These patterns provide us with a benchmark of how individuals in the treatment groups *should* respond to objective information about home price changes in the last one or five years (at least if one is willing to assume that these past patterns will continue to hold going forward).

[10] The statistics on the United States population come from the 2014 ACS 1-year sample of household heads.

[11] We ask respondents when they enter our survey panel to answer 5 questions that evaluate their numeracy. The questions are taken from Lipkus et al. (2001) and Lusardi (2009). Those who answer at least 4 of the 5 numeracy questions correctly are classified as having high numeracy.

For this purpose, we estimate time series regressions of home price changes on lagged home price changes, over different time horizons. In particular, we test how strongly past one-year and five-year growth (the information we provide in the treatments) relate to future growth over the next one year or the next 2-5 years. These two horizons are chosen because they are the "short" and "medium" horizon that we will use in our analysis of respondents' expectations (revisions), as explained in Section 4.2.

Using CoreLogic Home Price Index data that covers the years 1976-2015, we estimate autoregressive coefficients at the zip code level using the following specification:

$$\Delta_h \log(HPI_{g,t+h})/h = \alpha_g + \phi_g \Delta_l \log(HPI_{g,t})/l + \varepsilon_{g,t},$$

where $HPI_{g,t}$ is CoreLogic's Home Price Index in year t in zip code g, h is the horizon over which the change in the dependent variable is computed (i.e., one or 2-5 years), and l is the horizon over which the change in the independent variable is computed (one or five years). Dividing by h and l means that we annualize all home price changes. The parameter ϕ_g indicates persistence in home price growth for a given zip code g. We estimate the model using ordinary least squares (OLS) with Newey-West standard errors in order to account for the serial correlation in error terms due to overlapping observations.

Table 2 reports various statistics (mean; standard deviation; median) of the estimates across the zip codes, as well as proportion of the zip-code-level estimates that are statistically significantly positive or negative at $p < 0.05$. For example, for the regression of one-year home price changes on lagged one-year home price changes, the average estimate of ϕ_g across the zip codes of respondents in our sample is 0.53 (the median is 0.55, and the standard deviation across the zip code level estimates is 0.14). This means that on average, a one percentage point higher growth rate in year t is followed by about a 0.5 percentage point higher growth rate in year $t+1$. The AR(1) coefficient is estimated to be significantly positive (at $p < 0.05$) for 91.2% of the zip codes in the sample. This indicates strong momentum in home price changes over short horizons, a pattern that has been well documented in the literature (e.g. Case and Shiller, 1989; Guren, 2016).[12] On the other hand, the average estimate of a regression of one-year home price changes on lagged five-year changes is 0.14, but indistinguishable from zero for the vast majority—more than 80%— of the zip codes in the

[12] Anenberg and Laufer (2017) show that some of the very high-frequency positive autocorrelation in price changes (at a horizon of a few weeks) can be explained by the time lag between buyer-seller agreement and contract closing. However, they find that this issue appears to have little effect on measured autocorrelations over horizons of two months or more, and thus should play little role in the patterns we emphasize.

sample, and significantly positive for 15% of the the zip codes.

Turning to the regressions of medium-term home price growth (that is, over 2-5 years) on lagged changes, we first note that the average coefficient on lagged one-year changes is very close to zero. The estimate is significantly negative (positive) for only 8% (1.7%) of zip codes. Thus, the most recent growth alone has little predictive power for the longer horizon. In contrast, we see stronger evidence of mean reversion in the case of a regression of 2-5 year growth on lagged 5-year growth, where the average estimate is -0.38, and the estimate is statistically significantly negative for more than half of the zip codes. This longer-horizon mean reversion is again in line with patterns detected in earlier work (e.g. Cutler et al., 1991; Glaeser et al., 2014).

In sum, there is strong momentum in home price changes over short horizons, and mean reversion over longer horizons. Appendix Table A-1 shows that our qualitative conclusions are similar if we use county- instead of zip-code-level indices, or if we instead use the Zillow Home Value Index (which covers more zip codes than the CoreLogic index but starts only in 1996), though mean reversion over the five-year horizon is notably stronger in that case.

4 Analysis of Baseline Perceptions and Expectations

In this section, we analyze the perceptions and expectations in the first (baseline) stage. These provide the "input" for our subsequent experimental analysis, but are also of interest by themselves.

4.1 Perceptions and Perception Gaps

Respondents were first asked for their perceptions of past home price changes in their zip code over the past twelve months and over the past five years. C-frame respondents directly report their beliefs in percentage point terms, but for L-frame respondents who report beliefs in levels, we compute percentage point changes. Summary statistics of respondents' perceptions of past home price changes are reported in Panel A of Table 3. Respondents, on average, perceive that home prices in their zip code increased by 3.8% over the past 12 months. The perceived average change over the past five years, annualized, is 1.5%. The large standard deviations, and the fact that average absolute perceptions are meaningfully larger than the average perceptions, indicate that there is substantial heterogeneity in perceived home price changes. Respondents report a mean confidence of 3.2 (on a 1-5 scale) in their past recall, suggesting sizeable uncertainty in perceptions.[13] The aver-

[13]After reporting their past perceptions, respondents were asked: *"How confident are you in your answers?"* on a five-point scale, where 1 meant "Not at all confident" and 5 meant "Very confident".

age perceptions are similar across the three groups (as indicated by the p-value in the fifth column of the table), which should not be surprising since assignment to groups is random. Columns (6) and (7) of the table show that the two question frames yield different responses, with the average perceived growth being significantly higher in the L-frame.[14]

A key ingredient in our analysis is a measure of respondents' ex-ante informedness about the treatment information. The measure we use to capture this is the difference between what the realized percentage point home price change over the past t years actually was in i's zip code according to the information source that we used (which we denote as $\pi_{i,t}$), and what respondent i believes the percentage point change in home prices was in her zip code (which we denote as $\hat{\pi}_{i,t}$). Note that the objective information (from Zillow) presented to the respondent is individual-specific and depends on her zip code. We refer to this difference as the "perception gap", $\alpha_{i,t} = \pi_{i,t} - \hat{\pi}_{i,t}$, with a positive (negative) gap reflecting an underestimation (overestimation) of past home price changes relative to the Zillow measure. For the five-year horizon, the perception gap is annualized.[15]

Panel B of Table 3 shows that the mean perception gap in our sample is 1.4 for the one-year horizon, and -0.5 for the (annualized) five-year horizon. That is, on average, respondents' perceptions of past home price growth are reasonably close to the Zillow HPI, with an underestimation at the one-year horizon and a slight overestimation at the five-year horizon. However, the corresponding standard deviations of 7.0 and 4.1, respectively, imply substantial heterogeneity in the perception "accuracy" among respondents; similarly, the average absolute perception gaps are quite large. This implies that on average, the information shown to respondents in treatments T1 and T5 is appreciably different from their priors.

We next investigate the correlates of these absolute perception gaps. Table 4 shows estimated coefficients from OLS regressions of the absolute perception gaps at the one- and five- year horizon on a rich set of demographic controls. We see that college-educated, higher-income, and high-numeracy respondents on average have smaller absolute gaps at both horizons (the estimates are however only significant at the one-year horizon). Respondents who report being more confident in their past perceptions (that is, those reporting 4 or more on the 5-point scale), and those who have checked housing websites over the past 12 months also tend to have smaller absolute gaps, as

[14] The difference between the frames is larger (and more significant) at longer horizons for both perceptions (Panel A) and expectations (Panel C). This may be partly due to respondents failing to appreciate compounding; specifically, if a respondent thinks that house prices increased annually by x% on average over the past five years, they may report $5x$, rather than $100(1+\frac{x}{100})^5 - 100 > 5x$.

[15] We annualize the five-year perception gap as follows: $[1+(\pi_{i,5} - \hat{\pi}_{i,5})]^{1/5} - 1$. We continue to use the notation $\alpha_{i,5}$ to refer to the annualized five-year perception gap. The perception gap is annualized so that the analysis is comparable across the two horizons.

one might expect; however, the estimates are not significant at conventional levels. It is notable that tenure in one's town, being a homeowner, or planning to buy or sell a home soon are not associated with smaller gaps; the latter finding suggests that perceptions differing from objective measures are unlikely to be a result of rational inattention. Unsurprisingly, respondents residing in volatile housing markets (defined as areas with above-median volatility in home prices over the past five years) have significantly larger absolute perception gaps on average. Notably, the R-squared of these two regressions indicate that less than 7% of the variation in perceptions can be explained by these controls.[16] Thus, the extent to which respondents are "surprised" by the provided information is largely orthogonal to demographics.

4.2 Expectations of Future Home Price Growth

As mentioned above, we elicit respondents' home price expectations (at the zip code level) for the next one year and five years. We would expect a significant correlation between the five-year and year-ahead expectations simply because the five-year expectation is a combination of a respondent's expectations of year-ahead home price changes and 2-5 years ahead home price changes. We, therefore, separately analyze respondents' 2-5 year-ahead expectations. This is simply $y_{i,2-5} = \left[1 + \frac{(y_{i,5} - y_{i,1})}{(1+y_{i,1})}\right]^{1/4} - 1$, where $y_{i,h}$ is i's expectations about home price changes (in percent terms—with, for example, a percentage point change denoted as 0.01) at horizon h. We refer to these as "medium-term" expectations.

Panel C of Table 3 displays summary statistics of home price expectations at the baseline. We see that respondents, on average, expect a 3.5% increase in house prices in their zip code over the next 12 months, 11.0% over the next five years, and an annualized change of 1.7% at the 2-5 year horizon. The sizable standard deviations highlight the substantial heterogeneity in beliefs in the sample. As was the case for perceptions, the L-frame elicitation method yields higher means, particularly for the longer horizon. Finally, note that average expectations are similar to average past perceptions (reported in Panel A of the table), potentially the result of extrapolation from the (perceived) past to the future. We turn to this topic next.

As noted in Section 2.1, respondents also reported their subjective probability distribution across four intervals of future home price growth at both horizons. Following the approach developed by

[16] When looking at individual demographic characteristics in a univariate framework, Appendix Table A-2 shows that males, higher-income respondents, college-educated individuals, high-numeracy respondents, married individuals, those who frequently check housing websites and other sources, and those confident in their recall have significantly smaller average absolute perception gaps at the one-year horizon. For the five-year horizon, homeowners, higher-income individuals, and C-frame respondents have smaller absolute gaps, on average.

Engelberg et al. (2009), we fit a generalized beta distribution to each respondent's stated probabilistic beliefs (or a uniform distribution if the respondent assigns all her mass to a single bin). We then generate the standard deviation of the respondent's fitted distribution—this is the respondent's *Prior Uncertainty*. The last two rows of Panel C show that respondents tend to perceive substantial uncertainty in their forecast for future home price growth. For example, the average (individual-specific) standard deviation at the 1-year horizon is 5.6%.

4.3 Home Price Expectations and Past Perceptions

Table 5, using the cross-sectional variation in the sample, regresses home price expectations onto past perceptions, and documents a significant correlation between the two. Column (1), for example, shows that a one percentage point higher perceived past one-year local home price change is associated with a 0.26 percentage point higher year-ahead local home price expectation. Thus, respondents who report higher past home price growth also tend to report higher expected future growth, consistent with extrapolation. Interestingly, our estimate is very similar to that by Case et al. (2012), who find a coefficient of 0.23 in a regression of expected year-ahead MSA-level home price changes on lagged actual 12-month changes (for a sample of recent homebuyers in four MSAs over 2003-2012). Controlling for expectations about various fundamentals in column (2) reduces the coefficient only slightly, even though the R-squared increases substantially.

Columns (3) and (4) show similar extrapolation from perceived longer-term past changes to year-ahead home price change expectations. The last four columns show that even medium-term expectations are positively related to past perceptions, though the estimates are substantially smaller than those in the case of near-term expectations. This latter finding is somewhat different from Case et al. (2012), whose respondents report more extreme long-term (10-year) forecasts.

Note that these estimates cannot be given a causal interpretation due to various individual-specific as well as geographic confounds and potentially other omitted variables. For example, a respondent who is optimistic may report both higher past home price changes as well as future expectations. Furthermore, Appendix Table A-3 shows that the C-frame elicitation method yields a stronger correlation between year-ahead expectations and past one-year perceptions, as well as between medium-term expectations and past five-year perceptions. Thus, it appears difficult to reach convincing conclusions about the link between past (perceived) home price growth and expected future growth based on an analysis of cross-sectional variation alone. Our experimental framework, discussed next, allows us to get around these issues.

5 Experimental Analysis

5.1 Hypotheses on Updating Behavior

In general, we expect our information intervention to cause respondents to revise their home price expectations under two conditions. First, their expectations need to be influenced by their beliefs about the measures we use in our information treatments, i.e., past short- and long-term home price changes. This would not be the case, for instance, if respondents believed that home prices follow a random walk. Second, respondents are not already fully informed about the true values of these past changes (as we confirmed in Section 4.1).

If respondents' expectations evolved in a "data-consistent" way (that is, in line with actual movements in home prices, analyzed in Section 3), we would expect to see updating that is consistent with momentum in the T1 group for short-term expectations. That is, we would see an under- (over-) estimation of past one-year home price changes leading to an upward (downward) revision in year-ahead home price expectations. Recall that underestimations correspond to positive perception gaps. Therefore, in this case, year-ahead home price expectation revisions would be expected to be positively related to the one-year perception gap for T1 respondents. The relationship between medium-term expectation revisions and one-year perception gaps should be weaker. Turning to the T5 treatment, data-consistent updating would predict little systematic relationship between annualized five-year perception gaps and year-ahead expectation revisions (though directionally, the relationship in actual home price changes is positive). In contrast, respondents should realize that there tends to be a negative relationship between past five-year growth and future 2-5 year growth—so that, if they learn that home prices grew faster over the past five years than they had thought, they should revise their 2-5 year expectations downward.

Behavioral theories of expectation formation would typically predict extrapolation at both horizons, meaning that respondents would fail to perceive longer-term mean reversion. Models embedding such expectation formation in an equilibrium model of the housing market, such as Glaeser and Nathanson (2017), may also predict that individuals *underreact* to recent home price changes when forming their short-term expectations—that is, they extrapolate, but not enough. The reason for this in their model is that naive buyers fail to realize that future buyers will also extrapolate.

We will initially use the data from our information experiment to distinguish between these hypotheses based on average updating behavior. We also study to what extent revisions appear consistent with Bayesian updating. Then, we explore heterogeneity (across individuals as well as geographic areas) in updating patterns in order to shed additional light on different theories of expectation formation. For example, we investigate differences in updating by respondents' age and tenure in their location to evaluate predictions of theories that emphasize such heterogeneity, such as Malmendier and Nagel (2016).

5.2 Non-Parametric Analysis

We first proceed with a non-parametric analysis of updating behavior. Panel D of Table 3 shows the revisions in home price expectations between the baseline and the final stage. The average revision in the sample is an increase of 0.3 percentage points at the one-year horizon, and a decrease of 0.1 percentage points for the 5-year forecast. While average revisions are similar across the three groups (the Control and two treatment groups), absolute revisions tend to be larger in the treatment groups. The final two rows show the fractions of respondents that change their expectations in the final stage (relative to the baseline stage). While even in the control group a majority of respondents update their expectations, this fraction is significantly higher in the treatment groups, suggesting that the information provision does affect respondent expectations.[17]

Next, we provide graphical evidence on the relationship between perception gaps and home price expectation revisions. The first row of Figure 1 shows the mean year-ahead expectation revisions for each of the three groups, conditional on one-year perception gap decile bins. While the one-year perception gap can be constructed for each respondent (since past perceptions are elicited from all respondents), the one-year past home price change according to Zillow is only revealed to the T1 group. Hence, we expect to observe a systematic relationship between revisions and the perception gap for the T1 group but not the other groups. That is exactly what we see in the first row of Figure 1. There is a nearly monotonic relationship between year-ahead revisions and one-year perception gaps for the T1 group, with greater underestimation of past home price changes leading to a larger upward revision of year-ahead expectations. This pattern is consistent with respondents perceiving momentum in the short-term, as observed in actual home price changes.

The second row of Figure 1 shows the average medium-term (that is, 2-5 years) home price expectation revisions, conditional on one-year perception gap decile bins. Here, for none of the

[17] There are different potential interpretations of the forecast updates for control group respondents—they may simply reflect "noise" (people expressing their "true" forecast with some random error in the initial and/or the final stage), or they could reflect genuine updating that occurs due to the respondents taking the survey that makes them think about various issues related to the housing market. Appendix Table A-4 presents regressions of the absolute revisions of control group respondents on various observable characteristics. The explanatory power of these characteristics is very limited; a strong relation to proxies for sophistication might have pointed toward the noise interpretation.

three groups do we see a strong relationship between expectation revisions and perception gaps.

Figure 2 displays the relationship between expectation revisions and (annualized) past five year perception gaps. The top row shows a weak monotonic relationship between perception gap bins and average year-ahead revisions, for T1 and T5.[18] The bottom row of the figure shows little systematic relationship in the case of medium-term expectation revisions. We certainly do not see a negative relationship for T5, as would have been the case if updating were consistent with mean reversion as observed in actual longer-term home price movements (see Section 3).

We next proceed with a more precise, regression-based evaluation.

5.3 Regression Analysis

Our main regression model for home price expectation updating is as follows:

$$\Delta y_{i,h} = \beta_0 + \beta_1 T_{1,i} + \beta_2 T_{5,i} + \beta_3 \alpha_{i,1} + \beta_4 \alpha_{i,5} + \beta_5 (T_{1,i} * \alpha_{i,1}) + \beta_6 (T_{5,i} * \alpha_{i,5}) + \beta_7 1_{\text{C-frame},i} + \varepsilon_{i,h}, \quad (5.1)$$

where $\Delta y_{i,h}$ is the revision in home price expectations, for horizon h. The model is estimated separately for the one-year horizon ($h = 1$) and the 2-5 year horizon ($h =$2-5). $T_{1,i}$ ($T_{5,i}$) is an indicator that equals 1 if respondent i is assigned to treatment T1 (T5); $\alpha_{i,H}$ is i's perception gap for the past H years, where $H = \{1, 5\}$; and $1_{\text{C-frame},i}$ is an indicator that equals 1 if i's expectations were elicited using the C-frame. The βs are the parameters of interest.

The constant term, β_0, captures the average revision for those Control group respondents who have a perception gap of zero and whose expectations are elicited using the L-frame. $\beta_0 + \beta_1$, for example, reflects the average revision for respondents in the T1 group (in the L-frame elicitation group) with a perception gap of zero. β_3 and β_4 capture revisions related to the one-year and five-year perception gaps, respectively, for respondents that are not shown the relevant information. β_7 allows for the possibility that revisions may depend on the elicitation method.

The main coefficients of interest are β_5 and β_6. β_5, for example, measures the sensitivity of home price expectations with respect to the one-year perception gap for the T1 group—it provides an estimate of the causal effect of the one-year past information on home price expectation revisions. β_5 and β_6 will be different from zero if revisions are systematically driven by the difference between the revealed Zillow information and a respondent's prior. As discussed earlier, data-consistent

[18] That we observe a relationship for T1 respondents may be somewhat surprising since the five-year perception gap is never revealed to them. However, this is likely because of the high level of correlation between one- and five-year perception gaps within respondents (Spearman rank correlation of 0.377, significant at $p < 0.001$).

updating would imply that β_5 would be positive when the dependent variable is $\Delta y_{i,1}$, and that β_6 would be negative when the dependent variable is $\Delta y_{i,2-5}$.

Equation (5.1) is estimated using ordinary least squares, with robust standard errors.[19] Columns (1) and (2) of Table 6 show the estimates for the short-term and medium-term expectation revision, respectively. In column (1), we see that the estimate of β_5 is positive and significant: the coefficient of 0.20 implies that, for each percentage point underestimation (overestimation) of past one-year home price changes, T_1 respondents revise up (down) their year-ahead expectations by 0.20 percentage points. For comparison, the average AR(1) coefficient of home price growth in our respondents' zip codes is 0.53 (see Table 2). This implies that the average respondent, when forming her expectations, may undercorrect for momentum present in her local housing market. It is also notable that the estimate of β_5 is quite similar to the cross-sectional estimate of 0.26 in Table 5. These comparisons are done more formally in Appendix Table A-5, which reports the results from a bootstrap analysis to compare the coefficients from the information experiment in this section with those from the cross-sectional analysis in Section 4.3 as well as the actual serial dependence in house prices (Section 3). In the case of β_5 (1-year past to 1-year expectations), both the experimental and the cross-sectional coefficients, while statistically indistinguishable from one another, are significantly smaller than the average coefficient from actual home prices.

The estimate of β_6 is 0.07 but not statistically different from zero. It is also not significantly different from the average coefficient across zip codes when regressing actual one-year growth in zip code home prices on lagged five year growth (0.14), or from the corresponding coefficient of 0.21 in the cross-sectional analysis (see Table A-5).

Looking at the other estimates in column (1), we see that β_1 and β_2 are indistinguishable from zero. This implies that there is no effect of the treatments on home price expectation revisions (relative to control group revisions), other than what is explained by the size of respondents' perception gaps. Likewise, both β_3 and β_4 are small in magnitude and not significantly different from zero, meaning that perception gaps are not significantly related to revisions for those respondents that are not shown the information. The estimate of β_0 indicates that control respondents in the L-frame, on average, revise their expectations up by 0.11 percentage points, an (economically and statistically) insignificant revision. β_7 is also not statistically significant; that is, mean revisions for respondents in the C-frame are not different from those in the L-frame.[20]

[19] Demographics are not included in the regression because random assignment to treatment groups should ensure demographics are irrelevant to treatment effects. Indeed, when we control for demographics (not shown), there is no notable difference in estimates; also, we find that the demographics are not significantly related to revisions.

[20] In Appendix Table A-6, we estimate a version of equation (5.1) in which we add interactions of all variables with a

Turning to column (2) in Table 6, we see a positive relationship between medium-term expectation revisions and perception gaps. Estimates of β_5 and β_6 imply that individuals revise up (down) their medium-term expectations by 0.04-0.05 percentage points per percentage point under- (over-) estimation of past home price changes. This apparent extrapolation is inconsistent with actual home price patterns. In particular, there is no evidence that, on average, our respondents view higher growth over the past five years as predictive of lower future medium-term growth, even though this is at least directionally the case in actual home prices. Table A-5 shows that the coefficient on past five-year changes when predicting future 2-5 year changes using actual home prices, which averages -0.38 across zip codes, is significantly different from the experimental estimate; the table also shows that the estimates from our experiment are again not significantly different from those from the earlier cross-sectional analysis.

5.3.1 Bayesian Updating. Could the revisions we observe be consistent with Bayesian updating? In a Bayesian framework, conditional on the perception gap, respondents who are less confident in their past perceptions should be more responsive to the treatment information. Likewise, respondents who are more uncertain about future home price changes should be more responsive to the treatment information. Since we collect data on respondents' confidence in their past perceptions (as mentioned in Section 4.1) and on their subjective distribution of future home price expectations at both horizons (as mentioned in Section 2.1), we investigate this next.

Recall that confidence is elicited on a 1-5 scale, where 5 is very confident. For ease of interpretation, we center the responses on a -2 to 2 scale, so that 0 denotes "somewhat confident". We augment the baseline specification in equation (5.1) by adding the *Confidence* variable and its interaction with all the other terms in the specification. The parameters of interest are the interactions of *Confidence* with the $T_{1,i} * \alpha_{i,1}$ and $T_{5,i} * \alpha_{i,5}$ terms. Bayesian-consistent updating would predict these triple interaction terms to be negative, i.e., more confident respondents should be less responsive to the information. Columns (3) and (4) of Table 6 show the relevant parameters from this specification, for the two horizons. For year-ahead expectations, there is no evidence that revisions are systematically related to confidence in past perceptions. In the case of medium-term expectation revisions (column 4), there is some evidence of higher confidence leading to lower revisions, though only for the T1 respondents.

C-frame dummy. This allows us to test if the impact of information differs systematically by whether expectations are elicited in levels or changes. The main interaction terms are not statistically significant (that is, the variables $T_{1,i} * \alpha_{i,1}$ and $T_{5,i} * \alpha_{i,5}$ interacted with the C-frame dummy), suggesting that the results are not being driven by a particular frame. This is in contrast to the findings for the cross-sectional relationship between perceptions and expectations (Table A-3, where the elicitation frame mattered.

We next turn to the relationship between prior uncertainty and revisions. For this purpose, we use the *Prior Uncertainty* measure that we described earlier in Section 4.2. We standardize the measure to have a mean of zero and standard deviation of 1. As in the case of confidence, the baseline specification is augmented by *Prior Uncertainty* and its interactions with all the other variables. As before, the parameters of interest are the interaction terms with $T_{1,i} * \alpha_{i,1}$ and $T_{5,i} * \alpha_{i,5}$. Under Bayesian updating, these should be positive. The last two columns of Table 6 show the relevant estimates from this model. We find no systematic relationship between prior uncertainty and revisions. It is also notable that controlling for baseline confidence or uncertainty has no impact on the uninteracted terms—those estimates continue to be similar to the baseline estimates. In sum, our analysis in this subsection finds little evidence that respondents update in a Bayesian manner.

5.3.2 Measurement Error. Like any survey data, our data on subjective expectations and perceptions may contain measurement error. A particular concern might be that measurement error in the independent variable (in this case, the perception gap) is attenuating the estimates of β_5 and β_6 towards zero.[21] This is particularly relevant for the interpretation of the estimate of β_5 for T1 respondents in the case of year-ahead expectation revisions. The low estimate, relative to actual dependence in home prices, could either be a result of measurement error in the independent variable, or could reflect underprediction of short-term momentum by respondents. We use two complementary approaches to investigate this.

The first approach instruments the perception gap with the information that is presented in the treatment stage. For example, we instrument the 1-year perception gap of T1 (T5) respondents with the local past one- (five-) year information that is presented to them in the experiment. Estimates of this model are presented in columns (3) and (4) of Table 7. The tests reported in the bottom of the table show that the instruments yield a strong first stage. In the case of year-ahead expectation revisions in column (3), instrumenting for the perception gap increases the estimate of β_5 to 0.31, up from 0.20 in the baseline model (estimates are statistically different at $p = 0.04$). At the same time, the estimate of 0.31 remains significantly lower than the average AR(1) coefficient in actual data of 0.53 ($p = 0.002$). The estimate of β_6 now indicates substantial extrapolation from past five year home price changes to year ahead expectations.[22] Turning to medium-term expectation revisions in column (4), our qualitative conclusions remain unchanged—in particular, β_6 remains positive,

[21]Measurement error in the dependent variable— expectation revisions—should not bias our coefficients of interest.

[22]This large estimated extrapolation partially derives from the high degree of predictive power that 5-year information has on 5-year perception gaps for T5 respondents in the first stage; combined with the lower amount of variation in the information relative to the perception gaps, small changes in shown information will mechanically result in large estimated swings in forecast updates under this IV specification.

indicating extrapolation even at the horizon where empirically mean reversion tends to occur.

Note that this instrument is valid if the presented information impacts respondents' home price expectations through the perception gap only. This exclusion restriction would be violated if, for example, respondents naively anchored their expectations to the presented information (something we discuss further below); it is hard to rule out this possibility completely. As an alternative, in columns (5) and (6) of Table 7 we instead instrument the perception gap with county-level home price changes; this information was not shown to treatment respondents (as explained earlier, the presented information was at the zip code level in most cases). Given the high correlation between zip and county level price changes, the first stage remains strong, although standard errors are larger than in columns (3) and (4). Our conclusions remain the same: attenuation bias does not seem to drive the result that on average, respondents underestimate the extent of short-term momentum.

For our second approach, we use the simulation-extrapolation (SIMEX) procedure (Cook and Stefanski, 1994). This approach simulates a model's estimation under additional measurement error, and then uses these simulations to extrapolate backwards what the true coefficient would be if the observed variables were free of measurement error.[23] This method is truly complementary to the first one, since it does not depend on the existence of a "good" instrument but instead on how valid the parametric assumptions are that we make about the type of measurement error in the data. Details of the procedure are presented in Appendix A.1. Using this approach, we find that the coefficient predicted for 1-year horizon forecast updates for the T1 group (β_5) under the extreme assumption that perception gaps only reflect measurement error is 0.26, which is lower than that found using the instrumental variables approach, and still very close to the originally estimated coefficient of 0.20 in our baseline regression. Other coefficients also remain close to those in Table 6.

Overall, this suggests that our finding of respondents on average underestimating the strength of short-term momentum in house price growth is not due to measurement error.

5.3.3 Size of the Perception Gap. A potential concern about external validity of our estimates stems from the fact that identification in the regressions above may come from those respondents that had particularly large perception gaps in the baseline stage, and those respondents may be non-representative in some ways. For instance, they may simply pay less attention to the housing market. To assess this concern, first recall from Table 4 that perception gaps were not strongly related to respondent characteristics, including proxies for whether one paid more or less attention

[23]Examples of economics papers that have used this approach to correct for measurement error include Benjamin et al. (2012) and Weizsäcker and Zankiewicz (2017).

to the housing market. Second, in Appendix Table A-7 we show results from a version of the above regressions where we add an additional interaction of our main terms of interest with the absolute size of a respondent's perception gap. These interaction terms are far from statistically significant, implying that identification is not driven by respondents with particularly large gaps. Third, the instrumental variable analysis in Section 5.3.2 speaks directly to this point as well; by instrumenting for perception gaps with the local house price growth, these regressions exploit only variation in perception gaps that is exogenous to respondents' priors. The fact that our conclusions remain qualitatively unchanged, in particular for the one-year horizon, suggests that our results are not driven by differential attention or other types of selection across our survey respondents.

5.3.4 Other Robustness Checks. Appendix A.2 discusses a series of additional specifications. For example, we test for asymmetric effects of positive and negative perception gaps, finding little evidence for differential effects. We also exploit respondents' subjective assessment of how the provided information compared to their prior belief about past changes in local home prices, and show that those who state that the shown house price growth was higher (lower) than what they had thought revise their expectations upward (downward) in a significant way. In other robustness checks, we restrict the sample to only treatment respondents to rule out the possibility that idiosyncratic revisions in the Control group are driving our results. Similarly, to test the sensitivity of our results, we bring back "outlier" respondents who have been trimmed in the main analysis, and instead winsorize these observations at the 2% level. We also restrict the sample to those respondents who are able to recall their baseline perceptions accurately (as measured by the qualitative question about subjective informedness). Finally, we report results from a falsification exercise. The results are robust across these checks, and corroborate our main findings reported in Table 6.

5.4 Heterogeneity in Updating Behavior

Our within-subject design allows us to investigate heterogeneity in updating. There are at least two reasons for doing so. First, the previous analysis may mask substantial differences across individuals in how they update. Second, heterogeneity in updating by either respondents' observable characteristics or geographic factors can inform us about the theoretical models of belief formation that best fit the data. For example, models of experience-based learning (Malmendier and Nagel, 2016) would predict that younger respondents should be more responsive to our treatment.[24] We

[24]The broader prediction from such models is that individuals' expectations are shaped by their lifetime experiences. This is difficult to directly test in our setting, since it would be reflected in respondents' priors (in the baseline stage),

also test whether, in line with data-consistent updating, individuals residing in areas with inelastic housing supply (which tend to exhibit stronger momentum in the short term and stronger mean-reversion in the long term) are more likely to update in a way consistent with momentum (mean reversion) when updating their short- (medium-) term expectations.

We denote the individual's updating type by $v_{i,h}$, where h denotes the horizon over which the respondent is forecasting (one year ahead, or five years ahead). That is, we allow the respondent to exhibit different behavior at different horizons (for example, extrapolation for short-term expectations, and mean reversion for long-term expectations). The three update types are:

- **Non-Updater (NU)**: This type does not update following treatment:

$$v_{i,h} = \text{NU if } \Delta y_{i,h} = 0.$$

- **Extrapolator (E)**: This type updates in a way consistent with momentum in home prices. If the perception gap, α_i, is positive (negative) — that is, the respondent under-estimated (over-estimated) past home price changes relative to the Zillow index — she revises up (down) her home price expectations. Formally, the definition is:

$$v_{i,h} = \text{ E if } (\alpha_i > 0, \Delta y_{i,h} > 0) \text{ or } (\alpha_i < 0, \Delta y_{i,h} < 0).$$

- **Mean Reverter (MR)**: This type updates in a way consistent with mean reversion in home prices. For example, if she learns that prices in the past actually increased by more than previously thought (that is, $\alpha_i > 0$), she revises her future forecast downward. Formally:

$$v_{i,h} = \text{ MR if } (\alpha_i > 0, \Delta y_{i,h} < 0) \text{ or } (\alpha_i < 0, \Delta y_{i,h} > 0).$$

For this analysis, we focus on two quantities. One, the proportion of non-updaters in the two treatment groups pooled together, relative to the the proportion of non-updaters in the Control group. Second, conditional on updating, the odds of being an extrapolator versus a mean reverter; this quantity is defined for those in the treatment groups only (since it is based on the respondent's perception gaps as implied by the revealed information). The first metric can be viewed as a measure of updating on the extensive margin, while the latter is a measure on the intensive margin.[25] We focus our analysis in this section on binary cuts of the data; however, in Appendix A.3 we

which could however also reflect other local aspects (e.g., economic conditions). Thus, we focus on the more easily testable prediction that an additional piece of information should receive more weight for somebody with less experience.

[25] Dominitz and Manski (2011) use a similar approach of classifying heterogeneity in subjective expectations about stock market returns. Their approach relies on the cross-sectional variation in expectations.

conduct a multivariate analysis of updating that also allows for differences in updating by other characteristics.

The first row of Table 8 shows the updating types for the treatment sample, separately for the year-ahead and five-year ahead expectations. At both horizons, treatment respondents are about 0.8 times as likely as control respondents to not revise their expectations. Conditional on revising, respondents are much more likely to update in an extrapolative manner at both horizons.

The remaining rows of the table show the updating patterns for various cuts of the sample. We also report the p-values to test the equality of proportions of non-updaters and updating types of the paired groups.

We do not see much difference in the propensity to update by treatment type for year-ahead expectations. However, the odds of extrapolation are significantly higher for T1 respondents. In the case of medium-term expectations, however, T5 information leads to a significantly larger impact on the extensive margin of updating, but no differential impact on the intensive margin.

Panel A of Table 8 shows cuts of the sample based on individual characteristics that allow us to test for experience-based updating.[26] We see little evidence of a differential propensity to update by either age or tenure in one's locality. While these extensive margin patterns seem at odds with predictions of models of age-dependent updating, the intensive margin results for the one-year horizon conform with these models: conditional on updating, younger respondents and those with a shorter history in a location exhibit higher odds of extrapolating at both horizons.

Turning to individuals who have had negative experiences in the housing market (they are currently "underwater" on their mortgage, or went through a foreclosure or a short sale in the past), we see they are relatively more likely to update but have lower odds of extrapolating their short-term expectations; the differences are however not statistically significant, in part due to the small subsample of households with negative experiences. Over the medium-term horizon, perhaps wary of their past adverse experiences, such individuals are significantly less likely to be extrapolators in both relative and absolute terms—in fact, conditional on updating, they are more likely to update their medium-term expectations in a manner that is consistent with mean reversion (in stark contrast to the rest of the sample).

Is it the case that active housing market participants, relative to their less active counterparts,

[26]To further investigate demographic heterogeneity in the experimental treatment effects, we estimate a version of equation (5.1) where the treatment terms are interacted with selective demographics (homeownership status, age, gender, education, and income). Results are presented in Table A-8. We find little evidence of significant heterogeneity, with the exception that more educated respondents extrapolate more at the medium term (2-5 year) horizon (something we also find in the multivariate analysis of updating in Appendix A.3).

are more responsive to the information treatment, as would be predicted under models of rational inattention? We investigate this in the last two rows in Panel A. We see that active housing market participants—proxied by respondents who report a probability of 50% or more of either buying a home over the next three years or (for owners) selling their home over the next year—are not more likely to update their expectations at either horizon. Conditional on updating, there is some evidence of such respondents being more likely to extrapolate at the one-year horizon.

Panel B presents sample cuts based on housing-related factors in the respondent's location. The first cut we look at is by housing supply elasticity of the respondent's location. We use the Saiz (2010) MSA-level elasticity measure based on land topology factors. We study this cut because momentum in home prices in the short term and mean reversion in the long term is stronger in areas with relatively inelastic supply (Glaeser et al., 2008).[27] Treated respondents residing in below-median supply elasticity areas (that is, the more inelastic areas) are more likely to update their expectations at both horizons (though the difference is not statistically significant). For both horizons, conditional on updating, these respondents are also substantially more likely to extrapolate. While this is "data-consistent" at the short horizon, the opposite is true at the longer horizon: those respondents that should be more likely to perceive mean reversion are in fact less likely to do so.

The next two cuts of the table exploit variation in the serial dependence in respondents' local home prices in both the short and long term.[28] These cuts yield results that are similar to those based on supply elasticity, except that they utilize a larger sample since the Saiz elasticity measure is only available for a subset of locations. Individuals residing in areas with high short-term momentum "correctly" exhibit higher odds of extrapolating for one-year expectations, although the difference is not statistically significant. The cut by the strength of longer-term mean reversion shows that, conditional on updating, respondents residing in areas in above-median long-term mean reversion are twice as likely to be extrapolators than mean reverters, a pattern that is opposite to what "data-consistent" updating would predict.

Finally, the last two cuts of the table investigate whether households extrapolate more in locations that have experienced stronger recent growth in prices. For the first cut, we split the sample by house price growth over January 2012 to January 2015 (given that January 2012 was a "turning point" for the US housing market overall, starting the recovery). Conditional on updating, respon-

[27] The average AR(1) estimate from a regression of one-year home price changes on lagged one-year changes for respondents in below-median elasticity zip codes is 0.57, versus 0.49 for above-median elasticity zip codes (difference significant at $p < 0.01$). The average estimate from a regression of 2-5 year home price changes on lagged five year changes is -0.54 for below-median elastic zip codes, versus -0.21 for the above-median group (difference significant at $p < 0.001$).

[28] For this purpose, in order to maximize coverage of our respondents' locations, we use county-level house price patterns; these are very similar to the zip-code-level ones shown in Table 2.

dents in areas that have experienced stronger recent growth are weakly more likely to extrapolate at both horizons (the difference across locations is not statistically significant). For the second cut, we split the sample by the severity of the house price drop around the Great Recession, to test whether recent experience of a large decline may have led to an enhanced recognition of medium-term mean reversion. Instead, we find that respondents in the hardest-hit areas are more likely to update their medium-term forecast, and conditional on updating are more likely to extrapolate (although the difference is not statistically significant). This provides an interesting contrast with the finding above, which suggested that *individual* negative experiences may be associated with an enhanced recognition of mean reversion.

It is worth discussing the fact that a non-trivial proportion of treated respondents—65 of the 345 respondents in the T1 group and 52 of the 339 respondents in the T5 group—do not revise their forecasts at either horizon. There could be several potential reasons for this. In standard models of expectation formation, an individual should respond to the information if it is ex ante unknown (and relevant). Thus, one possible reason for why some respondents do not update could be that they have small perception gaps (that is, the treatment information is close to their perceptions). Appendix Table A-12, however, shows that the perception gap size is only weakly significant as a predictor of updating. In addition, average perception gaps are almost identical for those who update and those who do not. Another possibility is that the non-updating behavior is driven by respondents who do not pay much attention to the survey. However, as shown in the last two rows of Panel A in Table 8, the updating behavior of respondents who have a reason to be more attentive—namely, those more likely to be active in the housing market soon—does not differ from their counterparts, suggesting that this is not a major factor. Our preferred explanation instead is that many non-updaters simply believe that home prices follow a random walk, meaning they do not realize there is serial dependence in home price changes. Another plausible explanation, and one that we cannot rule out, is that some respondents may have "sticky" expectations and do not revise them even when facing large perception gaps. Out setup is not designed to test for why expectations may be sticky, but this could be due to costs of processing information (Sims, 2003). Regardless, our main findings that individuals tend to underpredict momentum in the short term and do not perceive mean reversion over the medium term continue to hold even if we exclude non-updating respondents (as shown in columns (3) and (4) of Appendix Table A-10).

5.5 Discussion: Implications for Modeling Home Price Expectations

Our results suggest that on average, respondents perceive momentum in home price changes over short horizons—they respond positively to the gap in their one-year past perceptions when revising their year-head expectations. While average revisions at the one-year horizon are directionally consistent with observed momentum in actual home prices, the average respondent seems to *undercorrect* for momentum: our baseline estimate of perceived momentum is less than half of the coefficient of dependence in actual home price movements. On the other hand, we do not find evidence of the average respondent believing in mean reversion in medium-term home price changes relative to 5-year lagged growth rates (as is the case for actual dependence in home prices). If anything, the average respondent also appears to extrapolate when updating their medium-term expectations.

Thus, updating behavior, while directionally correct for short-term expectations, does not otherwise appear to be consistent with actual home price data. In contrast, it is at least qualitatively in line with behavioral models that assume extrapolation. In fact, a calibration in Glaeser and Nathanson (2017) implies a one-year updating coefficient of 0.2, exactly what we estimate. It may seem surprising that their non-rational agents extrapolate too *little* at the short horizon; however, this is essential for the model, since if buyers fully expected the extent of price increases, these expectations would be reflected in prices immediately and would prevent momentum from being as strong as in the data. Other behavioral theories of expectation formation such as Fuster et al. (2012) would imply that while agents fail to appreciate longer-term mean reversion (in line with our results), they correctly anticipate short-term momentum.[29]

The geographic heterogeneity patterns strengthen our findings of directionally data-consistent updating at the short horizon, but incorrect extrapolation at the longer horizon. We also do not find much evidence consistent with Bayesian updating, reinforcing the case for behavioral models as being most consistent with our evidence. Our heterogeneity analysis uncovers some patterns consistent with models of expectation formation based on personal experiences, although primarily on the intensive margin (the direction of updating conditional on doing so). At a more basic level, however, the fact that not all respondents update when receiving new information, and that some believe in momentum while others believe in mean reversion, suggests that allowing for belief or "type" heterogeneity in housing market models is likely important (e.g. Piazzesi and Schneider,

[29]A number of these theories (e.g. Barberis et al., 1998; Barsky and DeLong, 1993; Fuster et al., 2012) consider a situation where agents extrapolate fundamentals, and this is what in equilibrium generates predictable medium-term mean reversion in returns (which surprises the agents). Since we do not ask (or inform) respondents about fundamentals (which in the case of the housing market would be rents), but instead about "returns" (or rather price growth) directly, it is difficult to directly map such models to our setting, even though they capture the qualitative nature of our findings.

2009; Burnside et al., 2016; Guren, 2016). Behavioral models that assume all agents to be biased in the same way, while being more tractable, fail to capture this heterogeneity, and future work should aim to understand to what extent this may lead to counterfactual implications of such models.

For the interpretation of our results, it bears noting that when forming their expectations, our respondents may of course take into account information other than past home price growth, and we do not know their "mental model" nor their information set. For example, they could use a model such as $E(\Delta HP_{t+1}) = \alpha + \beta \Delta HP_t + \Gamma \mathbf{X_t}$, where the vector $\mathbf{X_t}$ includes variables capturing local macroeconomic conditions (and could also involve forecasts of those macro conditions). Our discussion for the most part abstracts from these $\mathbf{X_t}$ and, roughly speaking, focuses only on β. But of course a person's β could be (rationally) different if they also include $\mathbf{X_t}$ in their model, and one could worry that there is something akin to omitted variable bias plaguing our analysis.

However, this issue does not alter our conclusions, for several reasons. Most importantly, to the extent that a respondent's model is "correct," the resulting expectations should still capture the serial dependence pattern in house price growth correctly. One channel through which this could work is that when we reveal ΔHP_t to them, they then (rationally) update their beliefs about $\mathbf{X_t}$, which should in turn feed through to their expectation about ΔHP_{t+1}.[30] Furthermore, empirically, it is not the case that adding local economic control variables leads to a much lower estimated β at the one-year horizon. For instance, adding concurrent and lagged county-level growth in population and income-per-capita as right-hand-side variables to an analysis as in Table A-1 (the county-level version of Table 2) only has relatively small effects on the estimated βs; the average estimated one-year β for our respondents' counties decreases slightly from 0.60 to 0.52, but remains substantially higher than the average β of 0.2 implied by their updating behavior at the one-year horizon.

5.6 Persistence and Anchoring

A natural question to ask is whether the effects of the information provision persist beyond the relatively short time frame of the main survey. To investigate this, as described in Section 2.1, we re-elicited respondents' expectations in April 2015, about two months after the original survey.

We refer to the difference between the expectations elicited in the follow-up survey and the

[30] One reason why they might not get it right is if they fail to think through what the past HP growth being different from what they thought could mean for local macro conditions. This failure to update in a Bayesian way (which is somewhat reminiscent of Glaeser and Nathanson 2017) could provide a plausible "model" for the failure to update beliefs correctly in response to an information shock as provided in our experiment. However, outside of the updating context of our experiment, this model would not predict that there is too little extrapolation (at the one-year horizon) in people's beliefs, which is what we find in the cross-sectional analysis in Section 4.3.

baseline stage in the main survey as "follow-up" revisions, opposed to the (within-survey) "initial" revisions in the main survey. The first two columns of Table 9 focus on revisions of year-ahead home price expectations. The dependent variable in the first column is the initial revision. That is, the column reports estimates of equation (5.1), restricting the sample to those respondents who also take the follow-up survey. The β_5 estimate for this subsample is 0.18, about the same magnitude and precision as the full sample. We next estimate the same specification as in equation (5.1), except that the dependent variable now is the follow-up revision. If the impact of the intervention is long-lasting, we expect the estimate of β_5 to be qualitatively similar. Estimates for this specification are presented in column (2) of Table 9. β_5 is significant at the 5% level and positive. The point estimate declines in magnitude, but is not statistically different from the corresponding estimate in column (1). This indicates that the information intervention has persistent effects on our respondents' year-ahead expectations; the slight attenuation of β_5 may be the natural consequence of respondents' receiving additional information between the first survey and the follow-up survey (which could also be due to information frictions as e.g. in Coibion and Gorodnichenko, 2012, 2015). The last two columns of the table show that the effect of information on medium-term expectations, which was much smaller in the main survey, is no longer statistically significant in the follow-up.[31]

This analysis also directly addresses a potential concern with our design, namely that our information intervention may cause respondents to simply anchor their revised forecasts to the statistic presented to them in the treatment (Tversky and Kahneman, 1974), thereby explaining the correlation we find (at least for one-year information and expectations). Given that the information effect persists in the follow-up survey, two months after the time when the information was shown, it is unlikely that one can attribute the effect of information entirely to anchoring. Also, as noted earlier, the effects of information hold both in the C-frame and the L-frame, whereas one might expect anchoring to operate primarily in the C-frame (since the information is presented in terms of changes as well). Finally, anchoring alone should be equally strong for T1 and T5 and both expectation horizons, not consistent with the differences in treatment effects that we find.[32]

In principle, it is possible that when making their forecast in the follow-up survey, respondents

[31] A non-parametric analysis yields similar results. Follow-up revisions for one-year expectations are systematically correlated with the initial revisions: the correlation between the two for one-year expectations is 0.51, 0.39, and 0.28 for the T1, T5, and Control groups, respectively. That is, the persistence is the strongest for the T1 group, as one would expect if the impact of information was long-lasting. In the case of the medium-term (2-5 years) home price expectations, there is little difference in the correlation across the three groups.

[32] In an analysis not reported here (but available upon request), we investigate the correlates of the tendency to give a forecast in the final stage (post-information) that is close to the number presented in the treatment—that is, the tendency to anchor to the presented statistic. We find no evidence of this tendency being less pronounced for individuals with higher numeracy or education, as one might have expected if respondents were naively anchoring their responses.

anchor to their forecast in the final stage of the main survey, which would generate the persistence we observe (and be indistinguishable from genuine belief updating). This appears very unlikely, given that it would require a strong focus of respondents on their revised forecast while ignoring many other anchors they may encounter in the two months between the surveys.

5.7 Ex-post Accuracy

Another interesting question is whether our intervention impacts the ex-post accuracy of respondents' expectations. For this purpose, we compute the absolute difference between the respondents' year-ahead home price expectations and the realized local home price change between February 2015 and February 2016 (according to the Zillow zip code level data). We refer to this absolute difference as the "ex-post forecast gap." This analysis is restricted to respondents for whom Zillow zip code level data are available (81.4% of the sample).[33] Caution is warranted in using an ex-post realized outcome as a benchmark for accuracy of ex-ante expectations, since (1) home price changes are uncertain, and (2) respondents' point forecasts may refer to various statistics (i.e. mean, median, mode, etc.) of their subjective probability distributions (Engelberg et al. 2009). Nevertheless, we find such a comparison useful as suggestive evidence for whether information helps respondents form more accurate expectations.

Figure 3 shows the cumulative density plot of the ex-post forecast gap for the Control and Treatment respondents (combining the T1 and T5 groups, for which the distributions look similar). The distribution for the control group is shifted to the right, indicative of our treatment moving respondents closer to the ex-post realized outcome; however, we cannot reject the equality of the two distributions at conventional levels of significance ($p = 0.155$, Kolmogorov-Smirnov test). We find that 33% of treatment respondents are within 2 percentage points of the ultimately realized home price change, versus 26% for the control group. The average ex-post forecast gap is also smaller for the treatment group (4.36%, versus 4.88% for the control group; the p-value of a t-test for equality of means is 0.08). Thus, we find that the treatment seems to cause respondents' one-year forecasts to become marginally more accurate, based on the criterion above.

6 Expectations and Behavior

Our interest in home price expectations stems from the belief that they influence individuals' current and planned economic activity and economic outcomes. In this section, we investigate the link

[33] The Zillow HPI coverage of our respondents' zip codes slightly increased between 2015 and 2016.

between home price expectations and actual as well as intended choices. While expectations play a key role in economic models of decision-making under uncertainty, there is surprisingly little *direct* empirical evidence on how subjective expectations impact financial decisions. This is largely a result of data limitations— establishing a direct link between expectations and behavior requires data on both from the same individuals, something generally not available.[34]

6.1 Investment in Housing Fund

As explained in Section 2.1, respondents were asked to allocate $1,000 for a year between a risk-free savings account (with a 2% annual return) and a housing fund that pays an annual return equal to the one-year growth in home prices in the respondent's area. This allocation was first elicited hypothetically (in the baseline stage) and then incentivized (in the final stage). While the scenario is clearly stylized, it offers a clean setting to test the link between expectations and behavior.

The dependent variable of interest is the share (on a 0-100 scale) that is allocated to the housing fund. Respondents, on average, allocate slightly more than half of their $1,000 endowment to the housing fund (54% in the baseline stage, 59% in the final stage). The standard deviation of the housing share is roughly 34% (in both stages), meaning there is substantial heterogeneity. Column (1) of Table 10 reports estimates from an OLS regression of the housing share onto year-ahead home price expectations reported in the baseline as well as an extensive set of controls.[35] The housing share is significantly and positively related to home price expectations: a percentage point increase in year-ahead growth expectations is associated with a 0.82 percentage point higher investment in the housing fund. To put this estimate into context, a one standard deviation increase in baseline home price expectations is associated with a 3.1 percentage point increase in the housing share, while a one standard deviation increase in log household income is associated with a 4.7 percentage point increase in the housing share. Thus, the estimated effect size is economically meaningful. Coefficient estimates on several of the other controls are also sensible: for instance, individuals who report

[34] Existing literature on financial decision-making that analyzes the role of expectations in choices usually has data on subjective expectations or actual choices, but not both. Thus, the literature either uses proxies for expectations or intended choices. See, for example, Malmendier and Nagel (2016), Bachmann et al. (2015), Crump et al. (2015), D'Acunto et al. (2015), and Bailey et al. (2017). On the other hand, Bover (2015) investigates the link between home price expectations and past purchases of durables. Also, these studies mostly exploit cross-sectional variation in expectations and cannot entirely rule out the role of confounds or the possibility of reverse causality. An exception is Armantier et al. (2015), who establish a link between inflation expectations and a financially incentivized investment decision where future inflation affects payoffs. In other contexts, such as education, the link between subjective expectations and actual choices is fairly well-established (Jensen 2010; Wiswall and Zafar 2015).

[35] Because the dependent variable is not continuous but a fraction, we check the robustness of our results by estimating a fractional probit regression, following the methodology of Papke and Wooldridge (1996). The fractional probit specification yields estimates that are almost identical to the OLS model (results available from the authors upon request).

being confident in their perception of past home prices and those with self-reported risk aversion below the median invest a higher share in the housing fund. The same is true for individuals who report having checked websites or other sources that provide information on property prices; this could proxy for an individual's enthusiasm about housing as an investment.

Column (2) reports estimates from the same specification as in column (1), except that we now also include respondents' perceived downside risk in year-ahead home price changes as an additional covariate.[36] Baseline home price expectations continue to be a significant correlate of the housing share (though the estimate declines to 0.52), and a higher perceived downside risk in home prices is associated with a lower share allocated to the housing fund. Column (3) of Table 10 reports estimates from the same specification as in column (1), except that we now use the revised housing share (from the final stage) as the dependent variable and revised year-ahead home price expectations as the explanatory variable of interest. The estimate is more than twice the corresponding estimate in column (1) of the table; the two estimates are statistically different (p-value = 0.011). The stronger link between expectations and behavior could be due to the fact that respondents answered several housing-related questions between the baseline and the post-information stage, prompting them to think harder about housing investments, or due to the post-treatment choice being incentivized. Column (4) supplements this specification by adding in controls for both the baseline housing share and an indicator for whether this share was a corner solution (that is, zero or 100). The coefficient on expected home price growth declines in magnitude, but remains highly statistically significant.

Columns (1)-(4) investigate the link between expectations and the housing share in the cross-section, controlling for an extensive set of demographic variables. However, one might be concerned about unobservable differences across individuals confounding the analysis. The within-survey panel on investment choices and home price expectations, generated as a result of our information experiment, allows us to investigate whether the relationship between home price expectations and behavior holds within-individual, and to give a causal interpretation to the relationship. Columns (5) and (6) of Table 10 report estimates of a regression of the within-individual change in the housing share onto changes in year-ahead home price expectations. Remarkably, the estimate is nearly identical to that in column (1): it implies that a percentage point increase in home price expectations leads to a 0.83 percentage point increase in the housing share. In column (7), we further interact expectation revisions with a dummy for being in the control group. The idea is to see

[36] Here, we sum the probabilities that the respondent assigns to year-ahead home price changes being less than -5%, and being between -5% and 0%.

whether the expectation revisions of those respondents that were not shown information are also linked to behavior. The strongly negative interaction coefficient implies that this is not the case, suggesting that expectation revisions in the control group may largely reflect noise.

The previous specification uses all variation in measured expectation revisions, but we can also isolate the effect of expectation revisions that happen due to our exogenous information treatment. The last two columns of Table 10 thus present estimates of an instrumental variable (IV) regression, where we instrument for the home price expectation revisions by the perception gap interacted with the corresponding treatment indicators. Column (7) shows that a one-percentage point predicted increase in home price expectations (due to the information intervention) leads to a 3.7 percentage point increase in the housing share. Note that while the first-stage relationship has a slightly low F-statistic of 9.7, the Kleibergen and Paap (2006) test comfortably rejects the hypothesis that the model is underidentified (p-value < 0.001). [37, 38]

6.2 Other Housing-related Behaviors

While the investment choice provides us with a clean setting to investigate the role of home price expectations, its stylized setup arguably makes the role of home prices overly salient, relative to real-world choices. In Appendix A.4, we investigate how home price expectations are related to stated "real-world" behaviors, such as buying a non-primary home, the likelihood of buying (rather than renting) their next primary residence, or (for current owners) making investments in the home. In each case, controlling for an extensive set of observable respondent characteristics, we find a (statistically and economically) significant correlation between expectations and intended behavior.

7 Conclusion

Households' expectations are potentially a key driver of fluctuations in the housing market, one of the most important asset markets from a macro and household portfolio perspective. In this paper, we have made progress toward a better understanding of how households form their expectations,

[37] While the IV estimate is almost five times larger than the OLS one, the implied impact on the housing share is qualitatively similar: a one standard deviation increase in the baseline one-year perception gap for T1 respondents leads to an increase of 1.43 percentage points in year-ahead home price expectations, which in turn results in a 5.2 percentage point increase in the housing share. This compares with an OLS-implied (column 5) impact of a 3.3 percentage point increase in the housing share for a standard deviation increase in home price expectation revisions. The larger IV impact may be due to the OLS estimate being attenuated due to "noise" in the measured expectation revisions.

[38] In Appendix Table A-11 we show the direct impact of the information treatment on housing share revisions (the "reduced form" of the IV regression above). The qualitative impacts are similar: for example, for T1 respondents, a 1 point underestimation of past one year home price change leads to a 0.73 percentage point increase in the share assigned to the housing fund.

and how expectations affect behavior. Using a novel information experiment embedded in a survey, we have found that expectations about future home price growth react to information about past local home price growth, in a way not fully consistent with actual patterns in home prices. Specifically, on average respondents extrapolate from past information, but too little at a short (one-year) horizon, where actual momentum is strong, and too much at a longer (five-year) horizon, where home price growth tends to mean revert. We have also established a meaningful link between expectations and behavior, implying that our elicited expectations have information content and are not just "noise."[39] One implication is that survey expectations are important for policy makers and housing market analysts to track.

The housing market is unique in that short-term extrapolation (for instance, from past-year growth to year-ahead growth) is actually rational; in other financial markets, momentum is much weaker or inexistent. It is an interesting question whether the behavior we observe is a manifestation of a more general "extrapolation bias," and whether households are aware of differences in serial correlation patterns across markets. While we believe that a lot more research on this question is needed, Appendix A.5 presents some evidence from a separate qualitative survey question (asked to a different group of respondents) that suggests that on average, households tend to also believe in momentum in the aggregate stock market, but that they are aware that it is weaker than in the housing market.

Of course, expectations are not just affected by (perceived or actual) past local home price movements. They may be affected by personal experiences in the housing market, and also by social interactions. This latter channel is emphasized in recent work by Bailey et al. (2017) that is complementary to ours. In Appendix A.6, we offer a rough comparison of our estimates with theirs, and find similar effect sizes for the "individual" and "social" channels in affecting expectations. Measuring the importance of different drivers of expectations (including own experiences and social channels) and investigating individuals' "mental models" based on which they form expectations remains a priority for future research. We believe that survey-based information experiments, as implemented here, provide a powerful tool to do so.

[39] Of course, our main evidence here comes from a fairly stylized portfolio allocation decision; it would be desirable in future research to further investigate the link between direct measures of expectations and actual real-world behavior.

References

ADAM, K., P. KUANG, AND A. MARCET (2012): "House Price Booms and the Current Account," *NBER Macroeconomics Annual 2011*, 26, 77–122.

AMROMIN, G. AND S. A. SHARPE (2014): "From the Horse's Mouth: Economic Conditions and Investor Expectations of Risk and Return," *Management Science*, 60, 845–866.

ANENBERG, E. AND S. LAUFER (2017): "A More Timely House Price Index," *Review of Economics and Statistics*, forthcoming.

ARMANTIER, O., W. BRUINE DE BRUIN, G. TOPA, W. KLAAUW, AND B. ZAFAR (2015): "Inflation expectations and behavior: Do survey respondents act on their beliefs?" *International Economic Review*, 56, 505–536.

ARMANTIER, O., S. NELSON, G. TOPA, W. VAN DER KLAAUW, AND B. ZAFAR (2016a): "The Price is Right: Updating Inflation Expectations in a Randomized Price Information Experiment," *Review of Economics and Statistics*, 98, 503–523.

ARMANTIER, O., G. TOPA, W. VAN DER KLAAUW, AND B. ZAFAR (2016b): "An Overview of the Survey of Consumer Expectations," Staff Report No. 800, Federal Reserve Bank of New York.

BACHMANN, R., T. O. BERG, AND E. R. SIMS (2015): "Inflation expectations and readiness to spend: cross-sectional evidence," *American Economic Journal: Economic Policy*, 7, 1–35.

BAILEY, M., R. CAO, T. KUCHLER, AND J. STROEBEL (2017): "The Economic Effects of Social Networks: Evidence from the Housing Market," *Journal of Political Economy*, forthcoming.

BARBERIS, N., R. GREENWOOD, L. JIN, AND A. SHLEIFER (2015): "X-CAPM: An extrapolative capital asset pricing model," *Journal of Financial Economics*, 115, 1–24.

BARBERIS, N. C., A. SHLEIFER, AND R. W. VISHNY (1998): "A Model of Investor Sentiment," *Journal of Financial Economics*, 49, 307–343.

BARSKY, R. B. AND J. B. DELONG (1993): "Why Does the Stock Market Fluctuate?" *Quarterly Journal of Economics*, 108, 291–311.

BENJAMIN, D. J., O. HEFFETZ, M. S. KIMBALL, AND A. REES-JONES (2012): "What Do You Think Would Make You Happier? What Do You Think You Would Choose?" *American Economic Review*, 102, 2083–2110.

BESHEARS, J., J. J. CHOI, A. FUSTER, D. LAIBSON, AND B. C. MADRIAN (2013): "What Goes Up Must Come Down? Experimental Evidence on Intuitive Forecasting," *American Economic Review*, 103, 570–74.

BORDALO, P., N. GENNAIOLI, R. LA PORTA, AND A. SHLEIFER (2017a): "Diagnostic Expectations and Stock Returns," Working Paper.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2017b): "Diagnostic Expectations and Credit Cycles," *Journal of Finance*, forthcoming.

BOVER, O. (2015): "Measuring expectations from household surveys: new results on subjective probabilities of future house prices," *SERIEs*, 6, 361–405.

BURNSIDE, C., M. EICHENBAUM, AND S. REBELO (2016): "Understanding Booms and Busts in Housing Markets," *Journal of Political Economy*, 124, 1088–1147.

CASE, K. E. AND R. J. SHILLER (1989): "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 125–137.

——— (2003): "Is There a Bubble in the Housing Market?" *Brookings Papers on Economic Activity*, 2, 299–342.

CASE, K. E., R. J. SHILLER, AND A. K. THOMPSON (2012): "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets," *Brookings Papers on Economic Activity*, Fall, 265–298.

CAVALLO, A., G. CRUCES, AND R. PEREZ-TRUGLIA (2014): "Inflation Expectations, Learning and Supermarket Prices: Evidence from Survey Experiments," *American Economic Journal: Macroeconomics*, 9, 1–35.

COIBION, O. AND Y. GORODNICHENKO (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120, 116–159.

——— (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–2678.

COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2015): "How Do Firms Form Their Expectations? New Survey Evidence," Working Paper 21092, National Bureau of Economic Research.

COOK, J. R. AND L. A. STEFANSKI (1994): "Simulation-Extrapolation Estimation in Parametric Measurement Error Models," *Journal of the American Statistical Association*, 89, 1314–1328.

CRUMP, R. K., S. EUSEPI, A. TAMBALOTTI, AND G. TOPA (2015): "Subjective intertemporal substitution," Staff Report 734, Federal Reserve Bank of New York.

CUTLER, D. M., J. M. POTERBA, AND L. H. SUMMERS (1991): "Speculative Dynamics," *Review of Economic Studies*, 58, 529–546.

D'ACUNTO, F., D. HOANG, AND M. WEBER (2015): "Inflation Expectations and Consumption Expenditure," Working Paper.

DEFUSCO, A. A., C. G. NATHANSON, AND E. ZWICK (2017): "Speculative Dynamics of Prices and Volume," Working Paper 23449, National Bureau of Economic Research.

DELONG, J. B., A. SHLEIFER, L. H. SUMMERS, AND R. WALDMANN (1990): "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance*, 45, 379–395.

DOHMEN, T., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP, AND G. G. WAGNER (2011): "Individual risk attitudes: Measurement, determinants, and behavioral consequences," *Journal of the European Economic Association*, 9, 522–550.

DOMINITZ, J. AND C. F. MANSKI (2011): "Measuring and interpreting expectations of equity returns," *Journal of Applied Econometrics*, 26, 352–370.

ENGELBERG, J., C. F. MANSKI, AND J. WILLIAMS (2009): "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters," *Journal of Business & Economic Statistics*, 27, 30–41.

FOOTE, C. L., K. S. GERARDI, AND P. S. WILLEN (2012): "Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis," Public Policy Discussion Paper 12-2, Federal Reserve Bank of Boston.

FUSTER, A., B. HEBERT, AND D. LAIBSON (2012): "Natural Expectations, Macroeconomic Dynamics, and Asset Pricing," *NBER Macroeconomics Annual 2011*, 26, 1–48.

GAO, Z., M. SOCKIN, AND W. XIONG (2015): "Learning about the Neighborhood: Supply Elasticity and Housing Cycles," Working paper, Princeton University.

GELAIN, P. AND K. J. LANSING (2014): "House prices, expectations, and time-varying fundamentals," *Journal of Empirical Finance*, 29, 3–25.

GENNAIOLI, N., Y. MA, AND A. SHLEIFER (2016): "Expectations and Investment," *NBER Macroeconomics Annual 2015*, 30, 379–431.

GLAESER, E. L., J. D. GOTTLIEB, AND J. GYOURKO (2013): "Can Cheap Credit Explain the Housing Boom?" in *Housing and the Financial Crisis*, ed. by E. L. Glaeser and T. Sinai, University of Chicago Press, 301–359.

GLAESER, E. L., J. GYOURKO, E. MORALES, AND C. G. NATHANSON (2014): "Housing dynamics: An urban approach," *Journal of Urban Economics*, 81, 45–56.

GLAESER, E. L., J. GYOURKO, AND A. SAIZ (2008): "Housing supply and housing bubbles," *Journal of Urban Economics*, 64, 198 – 217.

GLAESER, E. L. AND C. G. NATHANSON (2014): "Housing Bubbles," Working Paper 20426, National Bureau of Economic Research.

——— (2017): "An Extrapolative Model of House Price Dynamics," *Journal of Financial Economics*, 126, 147–170.

GLASER, M., T. LANGER, J. REYNDERS, AND M. WEBER (2007): "Framing effects in stock market forecasts: The difference between asking for prices and asking for returns," *Review of Finance*, 11, 325–357.

GRANZIERA, E. AND S. KOZICKI (2015): "House price dynamics: Fundamentals and expectations," *Journal of Economic Dynamics and Control*, 60, 152–165.

GREENWOOD, R. AND A. SHLEIFER (2014): "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 27, 714–746.

GUREN, A. (2016): "House Price Momentum and Strategic Complementarity," *Journal of Political Economy*, forthcoming.

HARUVY, E., Y. LAHAV, AND C. N. NOUSSAIR (2007): "Traders' Expectations in Asset Markets: Experimental Evidence," *American Economic Review*, 97, 1901–1920.

JENSEN, R. (2010): "The (Perceived) Returns to Education and the Demand for Schooling," *Quarterly Journal of Economics*, 125, 515–548.

KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2017): "The Housing Boom and Bust: Model Meets Evidence," Working Paper 23694, National Bureau of Economic Research.

KLEIBERGEN, F. AND R. PAAP (2006): "Generalized reduced rank tests using the singular value decomposition," *Journal of Econometrics*, 133, 97 – 126.

KUCHLER, T. AND B. ZAFAR (2015): "Personal Experiences and Expectations about Aggregate Outcomes," Staff Report 748, Federal Reserve Bank of New York.

LIPKUS, I. M., G. SAMSA, AND B. K. RIMER (2001): "General Performance on a Numeracy Scale among Highly Educated Samples," *Medical Decision Making*, 21, 37–44.

LUSARDI, A. (2009): "Household Saving Behavior: The Role of Literacy, Information and Financial Education Programs," in *Policymaking Insights from Behavioral Economics*, ed. by C. Foote, L. Goette, and S. Meier, Federal Reserve Bank of Boston, 109–149.

MA, C. (2016): "Momentum and Reversion to Fundamentals: Are They Captured by Households' Subjective Expectations of House Prices?" Working paper, Xiamen University.

MADEIRA, C. AND B. ZAFAR (2015): "Heterogeneous Inflation Expectations, Learning, and Market Outcomes," *Journal of Money, Credit, and Banking*, 47, 867–896.

MALMENDIER, U. AND S. NAGEL (2016): "Learning from Inflation Experiences," *Quarterly Journal of Economics*, 131, 53–87.

MANKIW, N. G. AND R. REIS (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117, 1295–1328.

MARCH, C., A. ZIEGELMEYER, B. GREINER, AND R. CYRANEK (2016): "Pay Few Subjects But Pay Them Well: Cost-Effectiveness of Random Incentive Systems," Working Paper Series No. 5988, CESifo.

NIU, G. AND A. VAN SOEST (2014): "House Price Expectations," Discussion paper no. 8536, IZA.

PAPKE, L. AND J. WOOLDRIDGE (1996): "Econometric Methods for Fractional Response Variables with an Application to 401 (K) Plan Participation Rates," *Journal of Applied Econometrics*, 11, 619–32.

PIAZZESI, M. AND M. SCHNEIDER (2009): "Momentum Traders in the Housing Market: Survey Evidence and a Search Model," *American Economic Review*, 99, 406–411.

RÖTHELI, T. F. (2010): "Pattern-Based Expectations: International Experimental Evidence and Applications in Financial Economics," *Review of Economics and Statistics*, 93, 1319–1330.

SAIZ, A. (2010): "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125, 1253–1296.

SCHMALENSEE, R. (1976): "An Experimental Study of Expectation Formation," *Econometrica*, 44, 17–41.

SHILLER, R. J. (2005): *Irrational Exuberance, Second Edition*, Princeton: Princeton University Press.

SIMS, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

TVERSKY, A. AND D. KAHNEMAN (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131.

VISSING-JORGENSEN, A. (2004): "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear With Wealth? Evidence from Expectations and Actions," *NBER Macroeconomics Annual 2003*, 18, 139–194.

WEIZSÄCKER, G. AND C. ZANKIEWICZ (2017): "Measuring Applicant Quality to Detect Discrimination In Peer-to-Peer Lending," Working Paper, Humboldt University Berlin.

WISWALL, M. AND B. ZAFAR (2015): "Determinants of College Major Choice: Identification using an Information Experiment," *Review of Economic Studies*, 82, 791–824.

ZWANE, A. P., J. ZINMAN, E. VAN DUSEN, W. PARIENTE, C. NULL, E. MIGUEL, M. KREMER, D. S. KARLAN, R. HORNBECK, X. GINÉ, ET AL. (2011): "Being surveyed can change later behavior and related parameter estimates," *Proceedings of the National Academy of Sciences*, 108, 1821–1826.

Figure 1. Average Revision in Home Price Expectations, conditional on 1-year Perception Gaps.

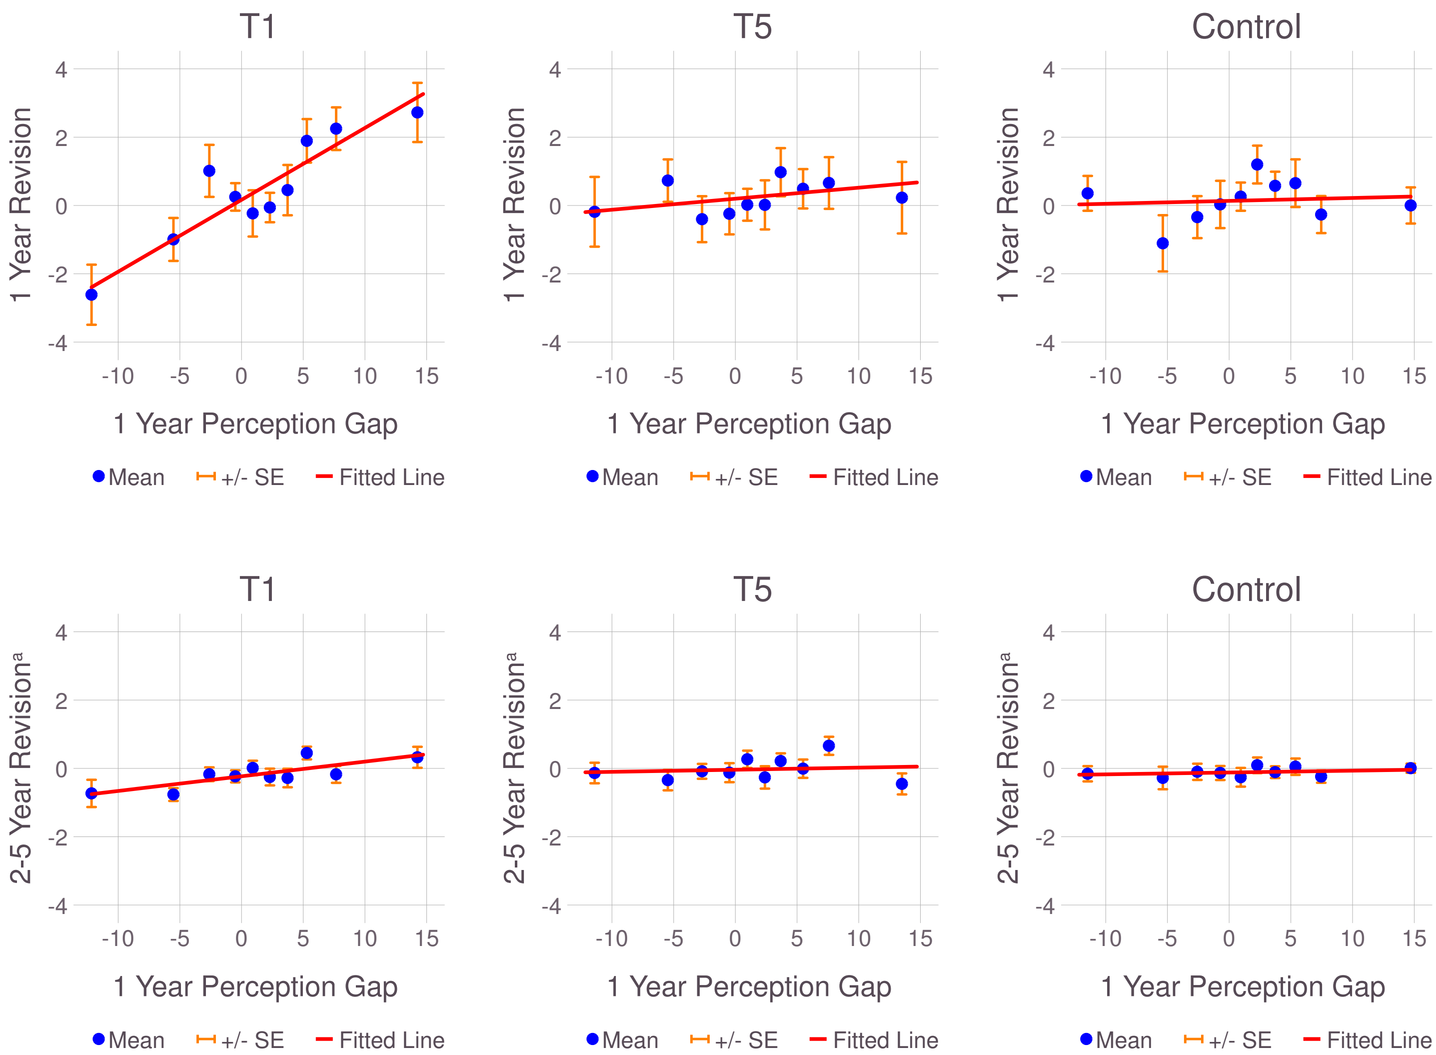

[a] Annualized

Figure 2. Average Revision in Home Price Expectations, conditional on 5-year Perception Gaps.

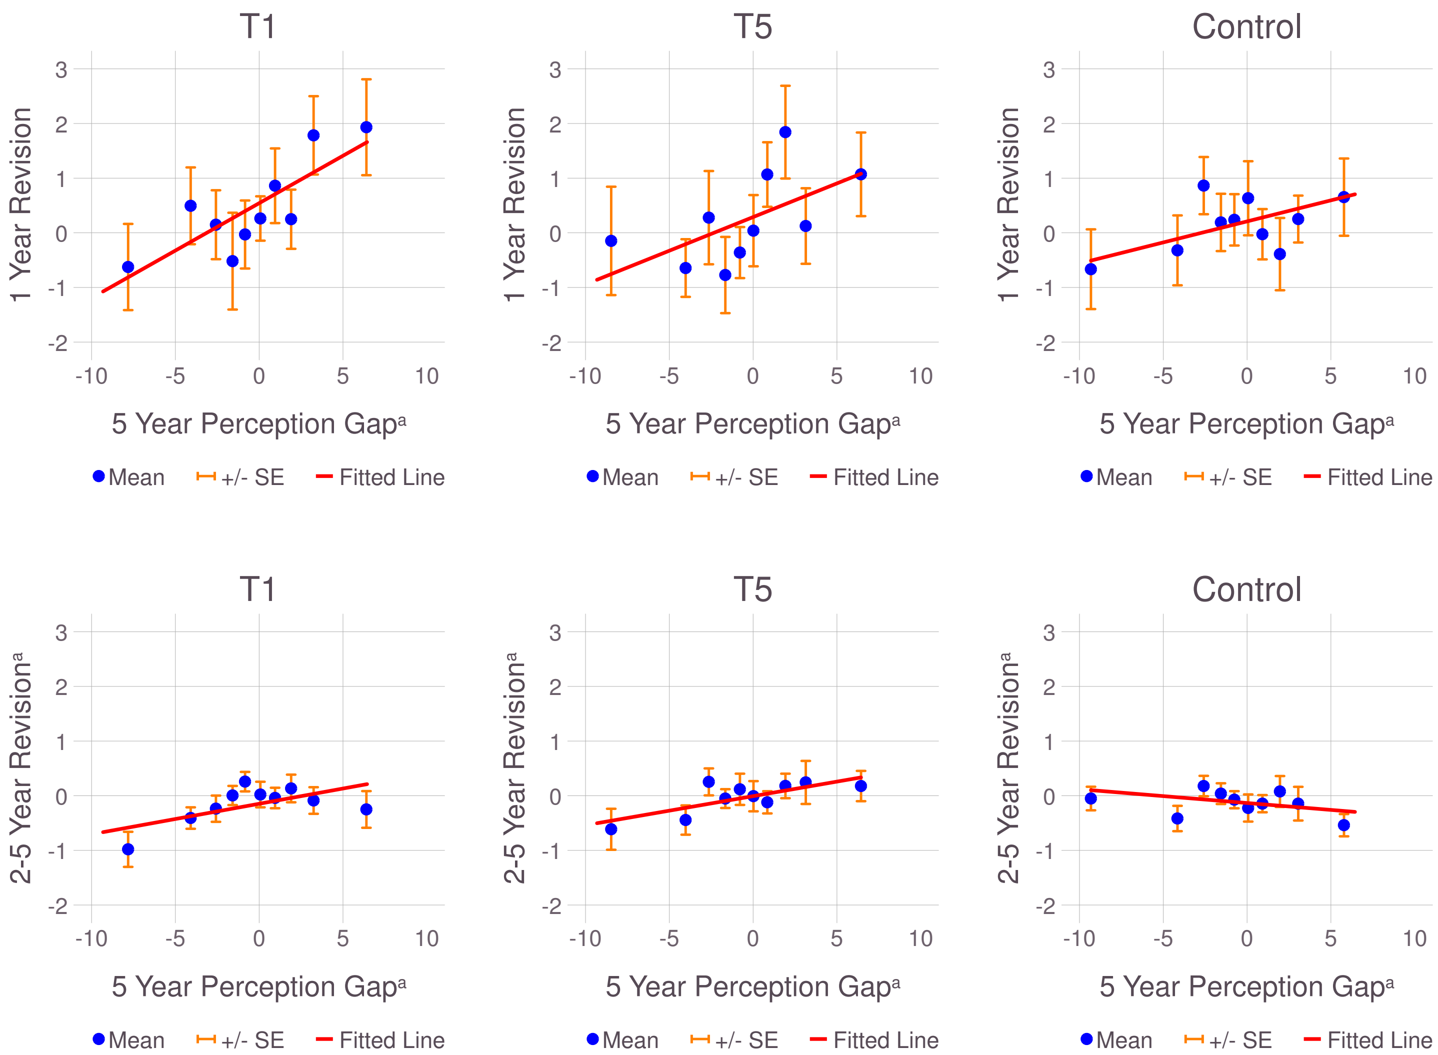

[a] Annualized

Figure 3. Cumulative Distribution Function of Realized Forecast Gap

Forecast gap = absolute value of (Final-stage 1-year expectation − Realization)

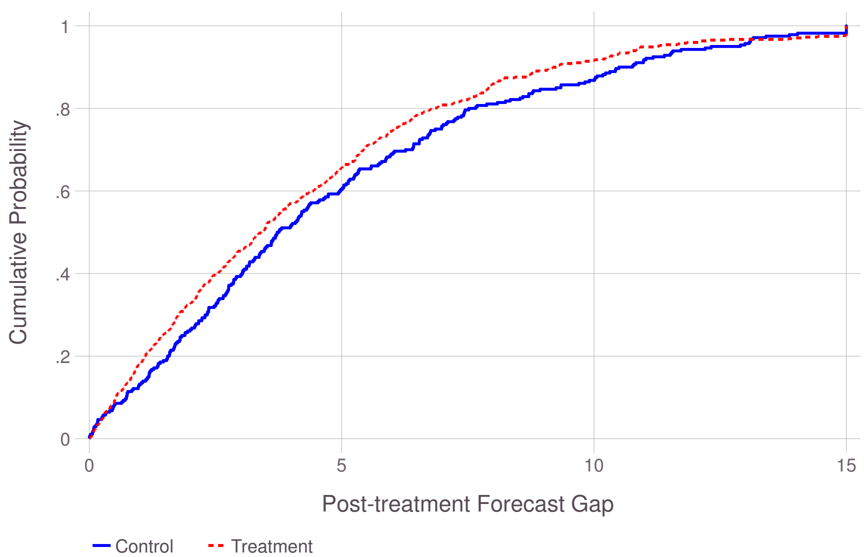

Data winsorized at [0,15]
Only respondents with Zillow zipcode data available are considered.

Table 1: Sample Characteristics

	Full Sample	Control	T1	T5	P-value[a]	Follow-up[b]
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	1020	336	345	339		729
Age (in years)	50.4	51.1	50.7	49.4	0.301	51.5
	(15.3)	(15.7)	(15.1)	(15.1)		(15.2)
Male	54.3%	57.4%	58.0%	47.5%	0.008	53.9%
White & Non-hispanic	78.8%	78.6%	79.7%	78.2%	0.887	80.1%
Married	67.8%	65.8%	69.9%	67.8%	0.523	66.3%
Homeowner	74.3%	75.0%	73.6%	74.3%	0.919	75.4%
Tenure in town/city (in yrs)	19.1	19.3	19.6	18.3	0.548	19.5
	(16.5)	(16.4)	(17.1)	(16.0)		(16.8)
Bachelor's Degree or More	55.3%	54.5%	54.2%	57.2%	0.680	55.0%
HH Income < $60,000	52.9%	51.5%	53.0%	54.3%	0.768	51.9%
HH Income < $30,000	20.0%	19.9%	20.3%	19.8%	0.985	20.2%
High Numeracy[c]	73.5%	75.9%	73.3%	71.4%	0.413	74.2%
Employed	66.5%	66.1%	66.4%	67.0%	0.970	65.6%
Unemployed	3.6%	4.8%	3.5%	2.7%	0.338	3.4%
Not in the Labor Force	28.9%	28.3%	30.1%	28.3%	0.828	30.2%
Census region location:						
Northeast	15.8%	14.0%	16.5%	16.8%	0.542	16.5%
Midwest	21.4%	22.3%	20.0%	21.8%	0.738	21.3%
South	38.4%	37.2%	40.3%	37.8%	0.677	37.4%
West	24.4%	26.5%	23.2%	23.6%	0.553	24.8%

Means of continuous variables reported. Standard deviations in parentheses for continuous variables.
[a] P-value of one-way ANOVA test of equality of each row variable across the three groups (Control, T1, T5).
[b] Follow-up is the sample that participates in the follow-up survey. Tests of equality of means or proportions between full sample (column 1) and follow-up sample fail to reject the null hypothesis of no differences (i.e. $p > 0.1$ for all variables.)
[c] High Numeracy indicates correctly answered four or more of five survey questions testing respondent's numeracy.

Table 2: Dependence in Actual Zip-code-level Home Price Changes (CoreLogic Data, 1976-2015)

	Survey Sample[a]			National Sample[b]		
	Estimates	Percent Positive[c]	Percent Negative[d]	Estimates	Percent Positive	Percent Negative
A. 1 year home price growth on lagged 1 year growth	0.53	91.2%	0.0%	0.53	89.7%	0.0%
	(0.14)			(0.14)		
	[0.55]			[0.56]		
B. 1 year home price growth on lagged 5 year growth	0.14	15.5%	2.3%	0.15	16.0%	1.2%
	(0.23)			(0.23)		
	[0.14]			[0.15]		
C. 2-5 year home price growth on lagged 1 year growth	0.03	8.1%	1.7%	0.03	8.8%	1.2%
	(0.12)			(0.12)		
	[0.02]			[0.02]		
D. 2-5 year home price growth on lagged 5 year growth	-0.38	3.2%	51.3%	-0.38	2.7%	49.9%
	(0.38)			(0.37)		
	[-0.40]			[-0.39]		

Table shows regression estimates of home price change dependence on previous changes. Mean coefficient across zip codes shown in first cell; standard deviation across zip codes shown in parentheses; median in square brackets.
Number of observations per zip code that these estimates are based on: A. 38 observations B. 34 observations C. 34 observations D. 30 observations.
[a] Consists of the sample of coefficients corresponding to each respondent's zip code, if covered by CoreLogic (N=753).
[b] Consists of the sample of all zip codes in the United States covered by CoreLogic. (N=7133).
[c] Indicates percent of local home price change coefficients statistically significantly positive at the 5% level, based on Newey-West standard errors (A: 1 lag; B: 5 lags; C: 5 lags; D: 10 lags).
[d] Indicates percent of local home price change coefficients statistically significantly negative at the 5% level, based on Newey-West standard errors (A: 1 lag; B: 5 lags; C: 5 lags; D: 10 lags).

Table 3: Home Price Perceptions and Expectations

	Full Sample	Control	T1	T5	P-value[a]	L-frame	C-frame	P-value[b]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Perceptions								
Past 1yr HP change	3.81	3.76	3.42	4.27	0.120	4.12	3.53	0.083
	(5.43)	(5.19)	(5.40)	(5.66)		(5.57)	(5.28)	
	[4.88]	[4.75]	[4.70]	[5.18]		[4.94]	[4.83]	
Past 5yr HP change	1.53	1.60	1.43	1.57	0.756	2.12	1.00	0.000
(annualized)	(3.04)	(2.85)	(3.10)	(3.18)		(3.44)	(2.52)	
	[2.71]	[2.58]	[2.70]	[2.84]		[3.37]	[2.10]	
Confidence in past HP change [c]	3.19	3.18	3.17	3.22	0.744	3.20	3.18	0.776
	(0.88)	(0.91)	(0.83)	(0.90)		(0.86)	(0.90)	
Panel B: Perception Gaps (=Zillow House Price Growth − Perceived House Price Growth)								
1yr Perception Gap	1.38	1.42	1.43	1.28	0.954	1.13	1.60	0.287
	(7.02)	(7.17)	(7.13)	(6.77)		(7.31)	(6.74)	
	[5.44]	[5.57]	[5.43]	[5.33]		[5.58]	[5.32]	
5yr Perception Gap	-0.51	-0.80	-0.36	-0.39	0.286	-1.42	0.30	0.000
(annualized)	(4.07)	(3.93)	(4.09)	(4.19)		(4.39)	(3.57)	
	[3.00]	[2.81]	[3.07]	[3.10]		[3.39]	[2.64]	
Panel C: Expectations								
Baseline 1yr exp. HP change	3.51	3.39	3.41	3.73	0.427	3.64	3.39	0.297
	(3.83)	(3.72)	(3.79)	(3.96)		(4.02)	(3.64)	
	[3.79]	[3.69]	[3.70]	[3.98]		[3.84]	[3.75]	
Baseline 5yr exp. HP change	11.01	9.99	11.68	11.33	0.059	14.11	8.21	0.000
	(9.77)	(9.10)	(10.14)	(9.96)		(10.84)	(7.68)	
	[11.37]	[10.59]	[11.91]	[11.58]		[14.39]	[8.63]	
Baseline 2-5yr exp. HP change	1.70	1.51	1.88	1.72	0.022	2.36	1.11	0.000
(annualized)	(1.76)	(1.61)	(1.87)	(1.78)		(1.91)	(1.37)	
	[1.81]	[1.68]	[1.96]	[1.80]		[2.43]	[1.25]	
1yr prior uncertainty[d]	5.55	5.70	5.38	5.56	0.214	5.27	5.80	0.000
	(2.38)	(2.59)	(2.23)	(2.29)		(2.16)	(2.53)	
5yr prior uncertainty	14.06	14.64	14.93	12.60	0.397	14.54	13.61	0.543
	(24.24)	(25.77)	(29.37)	(15.34)		(27.55)	(20.80)	
Panel D: Updates								
1yr forecast update	0.29	0.15	0.48	0.24	0.516	0.16	0.41	0.328
	(3.94)	(3.34)	(4.17)	(4.25)		(4.15)	(3.74)	
	[2.39]	[1.97]	[2.55]	[2.64]		[2.45]	[2.34]	
5yr forecast update	-0.12	-0.27	-0.25	0.18	0.715	-0.11	-0.12	0.981
	(8.01)	(6.25)	(8.37)	(9.12)		(9.26)	(6.69)	
	[4.80]	[3.78]	[4.85]	[5.78]		[5.64]	[4.06]	
2-5yr forecast update	-0.10	-0.11	-0.17	-0.03	0.484	-0.07	-0.14	0.438
(annualized)	(1.47)	(1.27)	(1.48)	(1.65)		(1.67)	(1.27)	
	[0.94]	[0.84]	[0.94]	[1.05]		[1.07]	[0.82]	
% update 1 year forecast	61.91%	56.89%	64.72%	64.01%	0.069	56.85%	66.48%	0.002
% update 5 year forecast	66.50%	60.48%	64.43%	74.56%	0.000	62.58%	70.04%	0.012

Perceptions, perception gaps, expectations, and updates all reported in percentage points. Mean reported in each cell. Standard deviation in parentheses. Mean absolute value in square brackets.

[a] P-value of one-way ANOVA test of equality of each row variable across the three groups (Control, T1, T5).

[b] P-value of one-way ANOVA test of equality of each row variable across the two framings (L-frame, C-frame).

[c] Confidence in recall of past home price changes. Question is asked on a 1-5 scale, where 5 is very confident and 1 is not at all confident.

[d] Prior uncertainty is the standard deviation of a beta distribution fit to the respondent's HP change expectations prior to treatment for the relevant horizons.

Table 4: Correlates of Perception Gaps

	Abs. Perception Gap[a]	
	1yr	5yr annualized
	(1)	(2)
Male	-0.23	0.32*
	(0.31)	(0.17)
Lived in current town/city for 15+ years	0.01	0.04
	(0.33)	(0.18)
Checked housing websites[b]	-0.29	-0.06
	(0.33)	(0.19)
Confident in recalled price change[c]	-0.45	-0.17
	(0.31)	(0.18)
Likely to buy or sell home in future[d]	0.19	-0.10
	(0.40)	(0.22)
White	-0.13	-0.34
	(0.38)	(0.24)
Age $<$ 50	0.02	-0.10
	(0.33)	(0.20)
Income $\geq$ \$75,000	-0.63*	-0.20
	(0.33)	(0.20)
Bachelor's Degree or More	-0.56*	-0.17
	(0.33)	(0.19)
Homeowner	0.24	-0.35
	(0.35)	(0.22)
Married	-0.56	-0.11
	(0.35)	(0.21)
Employed	0.64*	0.38*
	(0.35)	(0.20)
High Numeracy[e]	-0.92**	-0.10
	(0.36)	(0.20)
T1	-0.16	0.28
	(0.36)	(0.21)
T5	-0.26	0.34
	(0.35)	(0.22)
C-frame	-0.38	-0.78***
	(0.29)	(0.18)
Volatile Local Home Market [f]	1.26***	0.78***
	(0.33)	(0.21)
Constant	6.61***	3.40***
	(0.76)	(0.46)
Observations	1018	1017
R-Squared	0.053	0.063
Joint sig of covariates[g]	0	0
Mean of dep. variable	5.44	3.00

OLS estimates reported. Robust standard errors in parentheses. Significant at *p<0.10, **p<0.05, ***p<0.01. One percentage point is denoted as 1.

[a] The gap between the perceived and actual zip code home price change in absolute magnitude. All gaps annualized.

[b] Dummy that equals 1 if respondent reports consulting websites about home prices in past twelve months.

[c] Dummy that equals 1 if respondent reports being confident in their recall of past home price changes (i.e. answers 4 or more on a 1-5 scale, where 5 is very confident).

[d] Dummy that equals 1 if probability of buying home in three years is $\geq$ 50% or probability of selling home in one year is $\geq$ 50%.

[e] Dummy that equals 1 if respondent correctly answered four or more of five questions measuring numeracy.

[f] Dummy that equals 1 if zipcode home price volatility over the past five years, as measured by Corelogic's HPI, is above the sample median.

[g] F-test on equality of all covariates to zero (excluding constant). P-value shown.

Table 5: Relationship between Home Price Expectations and Perceptions in Baseline Stage

	1 Year Expectations				2-5 Year Expectations [a]			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Past 1 Year Perceptions	0.262***	0.226***			0.058***	0.047***		
	(0.029)	(0.028)			(0.012)	(0.012)		
Past 5 Year Perceptions[a]			0.213***	0.202***			0.094***	0.080***
			(0.051)	(0.047)			(0.024)	(0.023)
Constant	2.563***	1.306	3.189***	3.208	2.120***	2.272***	2.161***	1.696**
	(0.222)	(1.143)	(0.210)	(2.449)	(0.097)	(0.469)	(0.103)	(0.755)
Observations	1020	1020	1019	1019	1020	1020	1019	1019
R-Squared	0.138	0.236	0.029	0.166	0.158	0.226	0.151	0.224
Control for Fundamentals[b]	No	Yes	No	Yes	No	Yes	No	Yes

OLS estimates reported. Regression also includes a C-frame dummy. Robust standard errors in parentheses. Significant at *p<0.10, **p<0.05, ***p<0.01.

[a] Annualized.

[b] Fundamentals include measures of respondent expectations of general inflation, mortgage rate changes, rent inflation, future economic conditions, and future credit availability.

Table 6: Home Price Expectation Revisions

	Home Price Expectation Revisions at horizon:					
	1 year	2-5 years	1 year	2-5 years	1 year	2-5 years
	(1)	(2)	(3)	(4)	(5)	(6)
T1 (β_1)	0.02	-0.12	0.08	-0.17	-0.01	-0.13
	(0.29)	(0.11)	(0.30)	(0.12)	(0.29)	(0.11)
T5 (β_2)	0.10	0.10	0.09	0.09	0.17	0.10
	(0.29)	(0.11)	(0.30)	(0.12)	(0.30)	(0.12)
1yr Perception Gap[a] (β_3)	0.00	0.00	0.01	0.00	0.00	0.00
	(0.03)	(0.01)	(0.03)	(0.01)	(0.03)	(0.01)
5yr Perception Gap (β_4)	0.05	0.00	0.05	-0.00	0.05	-0.00
	(0.05)	(0.02)	(0.05)	(0.02)	(0.05)	(0.02)
T1 * 1yr Perception Gap (β_5)	0.20***	0.04**	0.19***	0.05***	0.19***	0.04***
	(0.04)	(0.02)	(0.05)	(0.02)	(0.04)	(0.02)
T5 * 5yr Perception Gap (β_6)	0.07	0.05*	0.07	0.05*	0.08	0.06**
	(0.08)	(0.03)	(0.08)	(0.03)	(0.08)	(0.03)
T1 * 1yr Percept. Gap * Confidence[b]			0.02	-0.04*		
			(0.05)	(0.02)		
T5 * 5yr Percept. Gap * Confidence			-0.04	-0.00		
			(0.10)	(0.04)		
T1 * 1yr Percept. Gap * Prior Uncertainty[c]					-0.05	-0.01*
					(0.05)	(0.01)
T5 * 5yr Percept. Gap * Prior Uncertainty					-0.03	-0.04
					(0.11)	(0.07)
C-frame (β_7)	0.14	-0.11	0.14	-0.10	0.15	-0.12
	(0.25)	(0.10)	(0.26)	(0.10)	(0.26)	(0.10)
Constant (β_0)	0.11	-0.05	0.08	-0.05	0.07	-0.04
	(0.24)	(0.10)	(0.24)	(0.10)	(0.24)	(0.10)
Observations	1015	1013	1015	1013	1010	1007
R-Squared	0.056	0.024	0.058	0.032	0.061	0.044
Joint sig of covariates[d]	0	.052	0	.155	0	0
Mean of dep. variable	0.29	-0.10	0.29	-0.10	0.27	-0.10
SD of dep. variable	3.94	1.47	3.94	1.47	3.92	1.47

OLS estimates reported. Robust standard errors in parentheses. Significant at *p<0.10, **p<0.05, ***p<0.01. One percentage point is denoted as 1. 5 year perception gap and 2-5 year home price change expectations are annualized. The specifications in columns (3)-(6) includes the uninteracted variables and the relevant single two-way interactions corresponding to the triple interaction coefficients shown.

[a] Perception gap $\alpha_{i,t} = \pi_{i,t} - \widehat{\pi}_{i,t}$, the difference between Zillow information and respondent i's price change perceptions over the past t years.

[b] Confidence in recall of past home price changes. Question is asked on a 1-5 scale, where 5 is very confident. The Confidence question is centered to be from -2 to 2, so that 0 corresponds to "somewhat confident."

[c] For 1yr [2-5yr] E(ΔHP) updates, this is the 1yr [5yr] standard deviation of ΔHP expectations prior to the treatment. The prior ΔHP uncertainty measures are standardized to have mean zero and standard deviation 1.

[d] F-test on equality of all covariatesto zero (excluding constant). P-value shown.

Table 7: Allowing For Possible Measurement Error in Perception Gaps

	Baseline (OLS)[b]		Information Shown IV[c]		County ΔHP IV[d]	
	(1)	(2)	(3)	(4)	(5)	(6)
	1 year	2-5 year	1 year	2-5 year	1 year	2-5 year
T1 * 1yr Perception Gap	0.20***	0.04**	0.31***	0.05**	0.29**	0.01
	(0.04)	(0.02)	(0.07)	(0.02)	(0.14)	(0.04)
T5 * 5yr Perception Gap	0.07	0.05*	0.42***	0.15***	0.60***	0.12
	(0.08)	(0.03)	(0.12)	(0.05)	(0.21)	(0.09)
1yr Perception Gap[a]	0.00	0.00	-0.02	-0.01	-0.02	0.01
	(0.03)	(0.01)	(0.04)	(0.01)	(0.07)	(0.03)
5yr Perception Gap	0.05	0.00	0.07	0.01	0.15	-0.01
	(0.05)	(0.02)	(0.08)	(0.03)	(0.12)	(0.05)
T1	0.02	-0.12	-0.16	-0.14	-0.24	-0.14
	(0.29)	(0.11)	(0.29)	(0.11)	(0.37)	(0.14)
T5	0.10	0.10	0.23	0.14	0.29	0.15
	(0.29)	(0.11)	(0.31)	(0.12)	(0.38)	(0.14)
C-frame	0.14	-0.11	-0.09	-0.17*	-0.34	-0.13
	(0.25)	(0.10)	(0.28)	(0.10)	(0.32)	(0.11)
Constant	0.11	-0.05	0.27	0.00	0.26	-0.08
	(0.24)	(0.10)	(0.25)	(0.10)	(0.33)	(0.12)
Observations	1015	1013	1015	1013	1015	1013
Joint Under-identification Test Statistic[e]			88.334	87.675	67.276	66.328
F-stat for T1 * 1yr Perception Gap of excluded instruments[f]			649.79	646.26	86.38	85.65
F-stat for T5 * 5yr Perception Gap of excluded instruments			356.65	351.63	83.87	82.86
F-stat for 1yr Perception Gap of excluded instruments			631.92	627.64	82.13	81.86
F-stat for 5yr Perception Gap of excluded instruments			267.7	266.45	52.82	52.56

Dependent variable: home price expectation revisions at 1-year or 2-5-year horizons.
OLS and IV estimates reported. Robust standard errors in parentheses. Significant at *p<0.10, **p<0.05, ***p<0.01.
One percentage point is denoted as 1. 5 year perception gap and 2-5 year home price change expectations are annualized.

[a] Perception gap $\alpha_{i,t} = \pi_{i,t} - \widehat{\pi}_{i,t}$, the difference between Zillow information and respondent i's price change perceptions over the past t years.

[b] Standard OLS Regression, as in Table 6, for comparison

[c] Instrumental Variables regression, where we instrument for 1yr Perception Gap, 5yr Perception Gap, T1*1yr Perception Gap, and T5*5yr Perception Gap with the information from Zillow that was shown (or would have been shown) at the treatment stage of the experiment. Specifically, the instruments are 1yr ΔHP information, 5yr ΔHP information, T1*(1yr ΔHP information), and T5*(5yr ΔHP information).

[d] Instrumental Variables regression, where we instrument for 1yr Perception Gap, 5yr Perception Gap, T1*1yr Perception Gap, and T5*5yr Perception Gap with the county-level ΔHP for the past 1 and 5 years. Specifically, the instruments are 1yr ΔHP at county level, 5yr ΔHP at county level,T1*(1yr ΔHP at county level), and T5*(5yr ΔHP at county level). We also add as control dummies for missing 1/5yr country ΔHP, and the relevant treatment indicator interactions (not shown).

[e] Displays the χ^2 statistic of a Kleibergen-Paap Lagrange Multiplier joint test for under-identification of IV regression with heteroskedastic-robust standard errrors. It is distributed under the null as a χ^2 random variable with 1 degree of freedom. Critical value at the 0.1% level is 10.828.

[f] Displays the F-test statistic of a Sanderson-Windmeijer multivariate F-test for weak identification of excluded instruments for a given endogenous regressor.

Table 8: Heterogeneity in Updating at both Horizons

		1 year Home Price Expectations				5 year Home Price Expectations			
	Sample Size	Non Updater[a]	P-value[b]	Extrapolator/ Mean Revert	P-value[c]	Non Updater	P-value[b]	Extrapolator/ Mean Revert	P-value[c]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Full Sample (T1+T5)	684	0.83		1.68		0.77		1.72	
T1	345	0.82	0.847	2.03	0.062	0.90	0.004	1.86	0.446
T5	339	0.83		1.40		0.64		1.60	
Panel A: Individual Characteristics									
$<$ 40 years old	202	0.90	0.483	2.56	0.045	0.88	0.178	2.45	0.070
40-60 years old	272	0.72		1.50		0.86		1.43	
$\geq$ 60 years old	210	0.90		1.38		0.63		1.58	
Has Lived In Locality $<$ 15 years	332	0.87	0.563	1.90	0.247	0.78	0.869	1.77	0.736
Has Lived In Locality 15+ years	352	0.78		1.51		0.77		1.66	
No Negative Housing Market Experience	621	0.84	0.551	1.73	0.360	0.77	0.653	1.86	0.014
Foreclosure, Short Sale, or Underwater Mortgage	63	0.66		1.30		0.86		0.88	
Pr(Buy) and Pr(Sell) Home $<$ 50%	397	0.84	0.970	1.50	0.198	0.76	0.778	1.70	0.880
Pr(Buy) or Pr(Sell) Home $\geq$ 50%	287	0.81		1.94		0.80		1.75	
Panel B: Location Characteristics									
Below-Median Supply Elasticity[d]	218	0.75	0.308	2.23	0.012	0.75	0.295	2.06	0.057
Above-Median Supply Elasticity	222	0.92		1.19		0.94		1.31	
Above-Median One-year Momentum [e]	315	0.93	0.341	1.84	0.464	0.74	0.457	1.76	0.422
Below-Median One-year Momentum	304	0.79		1.58		0.85		1.50	
Above-Median Long-term Mean Reversion[f]	310	0.83	0.673	1.90	0.319	0.78	0.913	2.15	0.007
Below-Median Long-term Mean Reversion	309	0.88		1.54		0.80		1.25	
Above Median House Price Growth Since Recession[g]	307	0.93	0.394	1.73	0.885	0.71	0.235	1.79	0.344
Below Median House Price Growth Since Recession	312	0.78		1.68		0.89		1.48	
Above Median Peak/Trough Ratio[h]	309	0.90	0.646	1.79	0.649	0.66	0.028	1.79	0.344
Below Median Peak/Trough Ratio	310	0.81		1.63		0.97		1.48	

[a] The proportion of respondents in the treatment groups who do not update their home price expectations, relative to the proportion in the control group.
[b] Significance test of differential effect of being treated between groups on the probability of updating. Uses a difference-in-difference approach by testing significance of regression coefficient for the interaction(s) of treatment indicator and group indicator(s).
[c] Significance test of difference in distribution of extrapolators and mean reverters among those who update, between groups. Reports p-value from Pearson χ_2 test.
[d] Housing supply elasticity measured at the MSA level according to Saiz (2010).
[e] Momentum measured by dependence of county-level 1 year home price appreciation on the previous year's home price appreciation.
[f] Mean reversion measured by dependence of county-level 2 to 5 year home price appreciation on the previous 5 year's home price appreciation. Above-median here means strong mean reversion (i.e., a relatively more negative regression coefficient).
[g] House Price Growth Since Recession defined as the growth rate of CoreLogic's HPI from January 2012 to January 2015 for a respondent's county.
[h] Peak/Trough Ratio defined as the maximum value in a county's CoreLogic HPI prior to January 2007 divided by the minimum value in the HPI after the peak. I.e. a higher ratio means a county saw a larger decline in house prices during the housing market downturn.

Table 9: Persistence in Impact of Information

Dependent Variable: Revision in HP expectations	1 year		2-5 year	
	Final - Baseline	Followup - Baseline	Final - Baseline	Followup - Baseline
	(1)	(2)	(3)	(4)
T1 (β_1)	-0.26	-0.09	-0.11	-0.19
	(0.34)	(0.40)	(0.12)	(0.17)
T5 (β_2)	-0.38	-1.16***	0.11	-0.24
	(0.33)	(0.44)	(0.13)	(0.19)
1yr Perception Gap (β_3)	-0.01	-0.03	0.00	0.01
	(0.03)	(0.04)	(0.01)	(0.02)
5yr Perception Gap (β_4)	0.05	0.05	0.01	0.06**
	(0.06)	(0.06)	(0.02)	(0.03)
T1 * 1yr Perception Gap (β_5)	0.18***	0.13**	0.03**	0.01
	(0.05)	(0.07)	(0.02)	(0.03)
T5 * 5yr Perception Gap (β_6)	0.02	-0.09	0.04	-0.04
	(0.11)	(0.12)	(0.03)	(0.05)
C-frame (β_7)	0.32	0.29	-0.21**	-0.17
	(0.30)	(0.36)	(0.11)	(0.15)
Constant (β_0)	0.15	0.96***	0.05	0.38**
	(0.27)	(0.34)	(0.11)	(0.16)
Observations	691	691	681	681
R-Squared	0.048	0.028	0.026	0.022
Joint sig of covariates[a]	0	.033	.049	.173
Mean of dep. variable	0.16	0.71	-0.05	0.14

OLS estimates reported. Robust standard errors in parentheses. Significant at *p<0.10, **p<0.05, ***p<0.01. One percentage point is denoted as 1. Sample restricted for all columns to those who answered both the original and followup survey.

[a] F-test on equality of all covariates to zero (excluding constant). P-value shown.

Table 10: Investment in Housing Fund and Expectations

Dependent Variable: Housing fund share (on a 0-100 scale)	Baseline		Post-treatment (Final)		Revision (Final-Baseline)			Revision (IV Regression[i])	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Baseline 1-year HP Expectation	0.82***	0.52*							
	(0.29)	(0.29)							
Pr(Decrease in HP next year)[a]		-0.14***							
		(0.04)							
Final 1-year HP Expectation			1.84***	1.19***					
			(0.28)	(0.21)					
1-yr Δ HP Exp. revision					0.83***	0.84***	1.14***	3.67***	3.95***
					(0.23)	(0.21)	(0.25)	(1.42)	(1.39)
1-yr Δ HP Exp. revision * Control							-1.30***		
							(0.42)		
Baseline Share in Housing Fund				0.75***		-0.22***	-0.21***		-0.23***
				(0.04)		(0.04)	(0.04)		(0.04)
T1	2.54	2.99	2.57	0.75	0.22	0.95	0.65	-0.91	-0.17
	(2.56)	(2.56)	(2.55)	(1.72)	(1.84)	(1.72)	(1.72)	(2.06)	(1.98)
T5	1.50	1.86	0.44	-0.66	-0.70	-0.18	-0.43	-1.05	-0.51
	(2.60)	(2.59)	(2.56)	(1.63)	(1.75)	(1.64)	(1.64)	(1.98)	(1.92)
Homeowner with zero equity[b]	-5.01	-4.82	-0.62	3.00	3.71	2.17	2.48	3.94	2.41
	(3.12)	(3.08)	(3.26)	(2.24)	(2.37)	(2.25)	(2.24)	(2.59)	(2.54)
Confident in recalled price change[c]	4.57*	3.82	1.48	-1.72	-2.75*	-1.38	-1.23	-2.75	-1.48
	(2.37)	(2.36)	(2.31)	(1.53)	(1.67)	(1.54)	(1.54)	(1.81)	(1.73)
Above-median risk aversion[d]	-7.28***	-7.35***	-6.56***	-1.55	0.29	-1.68	-1.72	0.88	-1.06
	(2.13)	(2.11)	(2.10)	(1.47)	(1.58)	(1.48)	(1.47)	(1.72)	(1.66)
Checked housing websites[e]	7.94***	8.08***	10.41***	4.76***	2.73	5.00***	4.83***	2.74	5.15***
	(2.41)	(2.40)	(2.37)	(1.68)	(1.77)	(1.70)	(1.69)	(1.94)	(1.91)
C-frame	2.42	0.76	2.13	0.64	-0.05	0.39	0.44	-1.30	-0.68
	(2.10)	(2.16)	(2.06)	(1.43)	(1.56)	(1.45)	(1.44)	(1.74)	(1.65)
Constant	-29.07	-21.07	18.24	37.48**	47.79***	37.72**	38.52***	49.59***	40.72**
	(19.68)	(20.38)	(19.28)	(14.56)	(15.96)	(14.90)	(14.85)	(17.67)	(17.05)
Demographics[f]	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Initial Corner Sols. Controlled[g]	No	No	No	Yes	No	Yes	Yes	No	Yes
Observations	1018	1013	1013	1013	1013	1013	1013	1012	1012
R-Squared	0.12	0.13	0.14	0.60	0.05	0.19	0.19		
Joint sig of covariates[h]	0	0	0	0	.016	0	0	.217	0
Mean of dep. variable	53.86	53.91	58.89	58.89	4.98	4.98	4.98	4.99	4.99
1st-stage F test statistic (IV only)								9.72	9.49

OLS estimates reported. Robust standard errors in parentheses. Significant at *p<0.10, **p<0.05, ***p<0.01. One percentage point is denoted as 1.

[a] The probability (on a 0-100 scale) that respondent assigns to year-ahead home prices decreasing.

[b] Because we also control for equity, which all homeowners have a value for, the interpretation of an isolated Homeowner variable is Homeowners with zero equity.

[c] Dummy that equals 1 if respondent reports being confident in their recall of past home price changes (i.e. answers 4 or more on a 1-5 scale, where 5 is very confident).

[d] Dummy that equals 1 if respondent reports a 4 or less (on 1-10 scale) to question about willingness to take risks in financial matters, where 10 is very willing.

[e] Dummy that equals 1 if respondent reports consulting websites about home prices in past 12 months.

[f] Includes binary indicators for owning a home, numeracy, ethnicity, gender, marital status, education, labor force status, and census region. Additionally, all regressions include controls for age, age^2, and logs of household income, equity in home, liquid savings, and personal debt.

[g] Includes Indicators for whether respondent assigned 0% or 100% to the housing fund in the baseline.

[h] F-test on equality of all covariates to zero (excluding constant). P-value shown.

[i] IV regression using treatment variables (perception gap times treatment indicators) as excluded instruments. First stage identical to column (1) of table 6. As a result, uninteracted perception gaps are included as additional controls in columns (7) and (8).