NBER WORKING PAPER SERIES

EXPECTATIONS WITH ENDOGENOUS INFORMATION ACQUISITION: AN EXPERIMENTAL INVESTIGATION

Andreas Fuster
Ricardo Perez-Truglia
Mirko Wiederholt
Basit Zafar

Working Paper 24767
http://www.nber.org/papers/w24767

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2018, Revised March 2019

We would like to thank Christopher Carroll, Theresa Kuchler, Kristoffer Nimark, Luigi Pistaferri, Dan Silverman, Michael Weber, and seminar attendants at the Arizona State University, Bocconi, ECB, EIEF Rome, Federal Reserve Board, Federal Reserve Bank of New York, Goethe University, Harvard, IIES Stockholm, UC Berkeley, UCSC Milan, the GEM/BPP Workshop, CESifo Venice Summer Institute, UCSB TEC Conference, and the AEA Annual Meetings for useful comments. Felipe Montano Campos and Jacob French provided superb research assistance. We thank the UCLA Ziman Center for Real Estate's Rosalinde and Arthur Gilbert Program in Real Estate, Finance and Urban Economics for funding. Fuster and Zafar were employed at the Federal Reserve Bank of New York while much of this paper was written. The views presented here are those of the authors and do not necessarily reflect those of the Swiss National Bank, the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research.

Expectations with Endogenous Information Acquisition: An Experimental Investigation
Andreas Fuster, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar
NBER Working Paper No. 24767
June 2018, Revised March 2019
JEL No. C81,C93,D80,D83,D84,E27,E3

ABSTRACT

Information frictions play an important role in many theories of expectation formation and macroeconomic fluctuations. We use a survey experiment to generate direct evidence on how people acquire and process information. Participants can buy different information signals that could help them forecast the future median national home price. We use an incentive-compatible mechanism to elicit their willingness to pay for information, and introduce exogenous variation in the rewards for ex-post forecast accuracy. We find that participants put higher value on their preferred signal when rewards are higher, and incorporate the signal in their beliefs if they obtain it. However, they disagree on which signal to buy, and as a result, making information cheaper does not decrease the cross-sectional dispersion of expectations. We further document that numeracy and the revealed "taste" for accurate expectations are important correlates of heterogeneity in all stages of the expectation formation process. We provide a model with costly acquisition and processing of information, and show that it can match almost all of our empirical results. Our findings also have implications for the design of information interventions.

Andreas Fuster
Swiss National Bank
Boersenstrasse 15
CH-8022 Zurich
Switzerland
andreas.fuster@gmail.com

Ricardo Perez-Truglia
Anderson School of Management
University of California, Los Angeles
110 Westwood Plaza
Los Angeles, CA 90095
and NBER
ricardo.truglia@anderson.ucla.edu

Mirko Wiederholt
Professor of Economics
Sciences Po, Paris
28 rue des Saints Peres
75007 Paris
France
Mirko.Wiederholt@gmail.com

Basit Zafar
Department of Economics
Arizona State University
P.O. Box 879801
Tempe, AZ 85287
and NBER
basitak@gmail.com

An online appendix is available at http://www.nber.org/data-appendix/w24767

1 Introduction

Given the centrality of expectations in decision-making under uncertainty, consumer expectations have been the focus of much research, particularly in macroeconomics. Studies have found considerable dispersion in consumers' expectations (Mankiw, Reis, and Wolfers, 2003). The literature has theorized that this dispersion results from "rational inattention,"[1] which may arise due to the costs of acquiring information, as in the sticky information models of Mankiw and Reis (2002) and Reis (2006), or due to constraints on individuals' information processing capacity, as in Sims (2003) and Woodford (2003). However, there is little *direct* empirical micro evidence that shows how individuals acquire and process macroeconomic information in the real world.[2] In this paper, we present a survey experiment to study the causes and consequences of information acquisition and processing decisions.

We study information acquisition in the context of expectations about national home prices. Our interest in home prices stems from the fact that home price expectations play a prominent role in many accounts of the housing boom that occurred during the mid-2000s in the United States (e.g., Shiller, 2005; Glaeser and Nathanson, 2015). Home price expectations have been measured with survey data, and these survey measures have been shown to be associated with real behavior such as buying or making investments in a home (Armona, Fuster, and Zafar, 2017; Bailey et al., 2018). Given the prominence of housing in household portfolios, these decisions can have substantial welfare consequences.

We design a survey experiment to study the choice, valuation and use of information, and to investigate whether expectations across individuals converge upon provision of information in a setting where the acquisition of information is endogenized. The main survey was conducted in February 2017 on a nationwide sample of households broadly representative of the U.S. population, as part of a regular online survey on housing issues run by the Federal Reserve Bank of New York. The experimental design has four main stages. In the first stage, respondents report their expectations about the national median home price for the end of the year (their "prior belief"). In the second stage, which occurs much later in the survey, respondents are informed that their forecast will be re-elicited and incentivized: if it falls within 1% of the realized price, the respondent is eligible for a monetary reward. Half of the subjects are randomly assigned to a reward that pays \$100 with a probability of 10%, and the other half is assigned to a reward that pays \$10 with a probability of 10%.

Before the belief re-elicitation, respondents are given the opportunity to choose among different pieces of information that could be potentially useful for their forecasts: the average expert forecast of home price growth during 2017 (this forecast was +3.6% at the time of the survey), the national

[1]We use "rational inattention" in the broad sense of referring to all models where there is some trade-off between expectations incorporating all available information optimally and some cost of doing so.

[2]In a recent survey article, Gabaix (2017) makes the case for more experimental evidence on the determinants of attention, and the consequences of inattention.

home price change over the past one year (+6.8%), or the national home price change over the past ten years (-0.9%, or -0.1% annually). Respondents can also state that they do not want to see any information. These information pieces differ markedly in terms of their informativeness. For instance, one reasonable criterion, although certainly not the only one, is the information's ex-ante predictive power during the years leading up to the survey. Based on this criterion, the expert forecast is the most informative (RMSE of 2.8), followed by the past one-year change (RMSE of 3.2), and the ten-year change (RMSE of 7.9). This ranking of informativeness is consistent with findings from the real estate literature. For instance, the fact that past one-year price changes perform better than ten-year changes is consistent with the well-documented momentum in home prices over short horizons (Case and Shiller, 1989; Guren, 2016; Armona et al., 2017).

In the third stage, we elicit each respondent's willingness to pay (WTP) for their most preferred information type. We use a multiple-price-list variation of the Becker-DeGroot-Marschak (BDM) method: we ask individuals to choose either information or a payment between $0.01 and $5 in eleven scenarios. One scenario is then randomly chosen, and the corresponding choice is implemented. Finally, the survey concludes with the re-elicitation of home price expectations (the "posterior belief").

This experiment was designed to capture some important features of models of endogenous information acquisition and processing. On the one hand, sticky information models (e.g., Mankiw and Reis, 2002; Reis, 2006) propose that information frictions arise due to fixed costs of updating information sets, broadly interpreted as the costs of acquiring, absorbing and processing information. Conditional on paying this cost, agents have full-information rational expectations. On the other hand, noisy information models (e.g., Sims, 2003; Woodford, 2003; Maćkowiak and Wiederholt, 2009) highlight the importance of partial updating and noisy perceptions due to limited information-processing capacity. Our experimental design creates exogenous variation in the cost of acquiring information and exogenous variation in the reward to holding an accurate posterior. Moreover, individuals can choose between different pieces of information. To explore the sticky information mechanism, we study how variation in the fixed cost of acquiring information affects posterior beliefs. To explore the noisy information mechanism, we investigate how the size of the reward and numeracy/financial literacy (our proxy for the marginal cost of attention in noisy information models) affect the choice of the piece of information and, conditional on information being displayed, its processing.[3]

Our first result, with regards to preferences over pieces of information, indicates that individuals disagree on which piece of information to use: 45.5% chose the forecast of housing experts, 28% chose the past one-year home price change, and 22% chose the past ten-year home price change. The remaining 4.5% reported to prefer no information at all. Thus, less than half of the

[3]Note that the purpose of our research design is not to distinguish between rational updating and some alternatives that have been proposed in recent years, such as experience-based learning (Malmendier and Nagel, 2016), natural expectations (Fuster, Hebert, and Laibson, 2012), or diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2018).

sample chose the option that was most informative according to ex-ante predictive power. Some of this heterogeneity could be due to respondents using other criteria.[4] However, sophisticated respondents, as measured by their education or numeracy, were substantially more likely to choose the expert forecast than less sophisticated respondents. This finding suggests that at least part of the variation was due to cognitive limitations in identifying informative signals.

Second, we find that individuals put substantial value on information that can help them make more accurate forecasts, and also use the information they acquire. The average individual was willing to forego \$4.16 to see their preferred piece of information. This high valuation suggests that individuals value this information beyond the accuracy rewards provided in the survey. Furthermore, we find strong support for a basic prediction of models of endogenous information acquisition and processing: the average WTP is significantly higher in the \$100-reward condition than in the \$10-reward condition (\$4.58 and \$3.75, respectively). Interestingly, we also find that respondents that are more confident in their existing knowledge of the housing market, as indicated by low uncertainty in their prior forecast or them indicating that they have previously looked up house-price information, have significantly higher WTP for the available information.

We exploit the randomness in information provision (conditional on WTP) to study how the information acquired by the individuals affects their expectations. Our design measures whether individuals incorporate received information into their forecasts. Consistent with a genuine interest in information, individuals incorporate the information that they were willing to pay for into their forecast: they form posterior beliefs by putting, on average, 38% weight on the signal bought and 62% on their prior belief. As evidence of genuine learning, we show that the information provided in the main survey had a persistent effect in a follow-up survey conducted four months later. The rate of learning was similar across all three pieces of information, which confirms that the disagreement about the ranking of pieces of information was meaningful. However, we find patterns that at first sight run counter to the basic model of Bayesian updating. In particular, we find that individuals who had less uncertain prior beliefs put *more* weight on the purchased information. The same is true for more numerate individuals. As an alternative measure of effort spent on information processing and updating, we further use the time spent on the relevant survey screens. We find that this time spent increases in the possible reward and the revealed valuation for the information, but again decreases in a respondent's prior uncertainty.

Our final and main result is about the effect of endogenous information acquisition on the cross-sectional dispersion in beliefs. We measure how the effective price of information (which was randomly assigned) affects the dispersion of expectations. We find that a lower cost of information acquisition does not cause lower cross-sectional dispersion in expectations. To understand the reason, we divide respondents into groups, based on their preferred information piece. On the one

[4]This finding could be partly driven by the fact that some respondents distrust experts (Silverman, Slemrod, and Uler, 2014; Cavallo, Cruces, and Perez-Truglia, 2016; Cheng and Hsiaw, 2017). We provide direct evidence about this mechanism using an auxiliary survey.

hand, exposure to information tends to reduce the dispersion in posterior beliefs *within* a group. For example, among individuals who preferred the expert forecast (a signal of 3.6%), exposure to it results in their posterior beliefs becoming more compressed around 3.6%. On the other hand, exposure to information increases the dispersion in beliefs *across* these three groups, because each group acquires a different signal and the signals were far apart. These opposing effects are similar in magnitude, and thus end up canceling each other out.[5] Additionally, we show that this result is not an artifact of respondents being able to view only one piece of information – in an auxiliary survey with a similar set-up except that individuals are allowed to view multiple pieces of information simultaneously, we confirm that exposure to information does not cause lower cross-sectional dispersion. Finally, contrary to the prediction of sticky information models, dispersion in beliefs remains high within each group of individuals who acquire the same information.

We then show that most of our experimental findings can be explained by a model of rational inattention with some ex-ante heterogeneity. Agents in the model face a fixed cost of acquiring information and a variable cost of processing displayed information (i.e., paying more attention to displayed information is more costly). Agents choose whether to acquire information, which piece of information to acquire, and the amount of attention allocated to the displayed information. There are three dimensions of heterogeneity in the model. Agents have heterogeneous beliefs about the precision of different pieces of information. The taste for information (i.e., the incentive to have an accurate posterior) and the marginal cost of processing information may also differ across individuals. Under the assumption that prior uncertainty about home price growth is negatively correlated with taste for information in the cross section (i.e., respondents who enter the experiment with more precise beliefs value the information more), the model can match most of the experimental findings. For example, the model can match the seemingly puzzling finding that individuals with lower prior uncertainty have a higher willingness to pay for information, spend more time processing displayed information, and respond more to displayed information. Moreover, if low numeracy is a proxy for high costs of paying attention, the model is consistent with low-numeracy individuals responding less to displayed information.

In the benchmark version of the model, agents have heterogeneous prior beliefs over precisions and therefore select different pieces of information. In an extension, we assume that agents have homogeneous prior beliefs over precisions and also face the cost of processing information when selecting an information source. In that version of the model, the finding that the most informative piece of information is the modal choice and the finding that more numerate individuals are more likely to select this piece of information also arise endogenously.

This paper is related to various strands of literature. First and most importantly, it is related

[5]This finding has some parallels with the literature on media bias and political attitudes, according to which dispersion in beliefs can be persistent because voters self-select into different information sources (Mullainathan and Shleifer, 2005). The underlying mechanisms, however, are different: in the political economy literature the differences in information choices arise due to self-serving biases, while in our context the differences in choices seem to arise due to differences in preferences over information sources or cognitive limitations.

to a growing body of work on inattention models (Mankiw and Reis, 2002; Sims, 2003; Woodford, 2003; Reis, 2006; Gabaix, 2014).[6] We contribute to this literature by providing *direct* empirical micro evidence on how individuals acquire and process information. Our results are broadly consistent with these models and highlight that information processing constraints (as in noisy information models) are likely an important feature of consumer expectation formation.[7] This is apparent in our finding that cross-sectional dispersion of expectations does not decrease when the cost of information is lowered, and that individual information acquisition and processing are systematically related to numeracy. Our first main theoretical contribution is the combination of a fixed cost of information acquisition and a variable cost of information processing. Existing models do not combine these two features. Our second main theoretical contribution is that, in the extended model, agents face the cost of processing information twice – when they choose information (selecting an information source) and when they use that information (incorporating displayed information before reporting posterior beliefs). In the extended model, cognitive limitations in deciphering the informativeness of signals generate heterogeneity in information choice, and thereby, heterogeneity in beliefs.

This paper is also related to a literature on the sources of dispersion in consumer expectations. For example, Figure 1 shows the distribution of house price growth expectations among consumers and experts. Consistent with the evidence for other types of macroeconomic expectations (Mankiw et al., 2003; Cavallo, Cruces, and Perez-Truglia, 2017), this figure shows that expectations are very dispersed, and more so among consumers than among experts. Our findings shed light on the sources of this dispersion. Our evidence suggests that disagreement about what sources of information to rely on plays an important role in explaining this dispersion. As a result, even if the acquisition cost of information went down to zero, our findings imply that we would still observe substantial dispersion in consumers' expectations. This finding may explain why dispersion in expectations among consumers tends to be much larger than it is among experts even though the estimated information acquisition costs are not larger for consumers (Coibion and Gorodnichenko, 2012). Furthermore, constraints in information processing and heterogeneous tastes or incentives for accuracy further contribute to the dispersion, even among consumers who prefer the same information sources.

Our approach is related to a recent literature on information-provision experiments. Particularly relevant for our purposes are papers that employ random information provision in surveys to understand expectation formation in the context of inflation (Armantier et al., 2016; Cavallo et al., 2017; Coibion, Gorodnichenko and Kumar, 2015), housing (Armona et al., 2017), or GDP

[6]For a textbook treatment of the literature on information acquisition and processing, including inattention models, see Veldkamp (2011).

[7]In that sense, our conclusion is consistent with Coibion and Gorodnichenko (2012), who estimate the response of disagreement to shocks to distinguish between inattention models, and find results that are more consistent with noisy information than with sticky information models. Their setup is quite different since they use time series variation in survey data (opposed to experimental variation, as in our case).

growth (Roth and Wohlfart, 2018). These papers tend to find that when individuals are provided with official statistics, the dispersion in expectations substantially decreases, thereby providing suggestive evidence in favor of costly information acquisition models: once a piece of information is provided by the experimenter for free, the dispersion in expectations is reduced. However, information-provision experiments ignore a crucial aspect of the real world: individuals have to choose from multiple information sources, and *where* they look for information can be even more important than *how frequently* they look for information. Our findings indicate that, once respondents are allowed to choose information endogenously, reducing the cost of information may fail to reduce dispersion in expectations.[8]

Endogenous information acquisition may thus increase disagreement across individuals, even though the opinions of people who look at the same signals tend to converge. Disagreement is thought to be an important driver of trade in asset markets (e.g. Harrison and Kreps, 1978; Scheinkman and Xiong, 2003; Hong and Stein, 2007). In the context of the housing market, Bailey et al. (2017) show that counties with higher disagreement (which in their case is driven by different house price experiences of out-of-town friends) see higher trading volumes. Thus, to the extent that households vary in the information sources they rely on, as our evidence suggests, differing signals from these sources may have important consequences for housing market activity and prices.

Finally, our results have implications for the design of information interventions. A growing body of research shows that, in a wide range of contexts, providing individuals with accurate information can have substantial effects on their beliefs and decisions (e.g., Duflo and Saez, 2003; Allcott, 2011; Cruces, Perez-Truglia, and Tetaz, 2013; Wiswall and Zafar, 2015). One of the policy implications often drawn from this literature is that entities should make more information widely available and easily accessible. Our evidence suggests that this strategy may not be sufficient, because individuals may not know which of the different pieces of information to focus on. Our findings imply that these interventions should either be targeted (providing consumers with limited but relevant information) or that they should guide consumers to help them interpret and weigh the various pieces of information.

The rest of the paper proceeds as follows. Section 2 introduces the research design and survey, and provides an outline of the empirical analysis. Section 3 presents the results. Section 4 describes the theoretical model and discusses how its predictions compare to the experimental findings. The last section concludes.

[8]Endogenous information acquisition has been studied in other contexts, such as hiring decisions (Bartoš et al., 2016) and tax filing (Hoopes, Reck and Slemrod, 2015). Additionally, some laboratory experiments have been used to study demand for information in stylized settings (e.g., Gabaix et al., 2006). More recently, laboratory experiments have also been used to test various predictions of inattention models (Caplin, Dean, and Martin, 2011; Khaw, Stevens, and Woodford, 2017; Ambuehl, Ockenfels, and Stewart, 2018; Dean and Neligh, 2018).

2 Research Design

Our main survey module was embedded in the 2017 housing supplement of the Federal Reserve Bank of New York's Survey of Consumer Expectations (hereon, SCE Housing Survey). This survey has been fielded annually every February since 2014 and contains multiple blocks of questions, some of which distinguish between owners and renters.[9] Among other things, the survey asks about perceptions of past local home price changes, expectations for future local home price changes, and past and future intended housing-related behavior (e.g., buying a home, refinancing a mortgage). Respondents also provide information about their location and many other demographic variables.

The SCE Housing Survey is run under the Survey of Consumer Expectations, an internet-based survey of a rotating panel of approximately 1,400 household heads from across the United States. The survey, as its name suggests, elicits 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 in and out of the panel each month.[10] Active panel members who had participated in any SCE monthly survey in the prior eleven months were invited to participate in the housing module. Out of 1,489 household heads on the panel that were invited, 1,162 participated, implying a response rate of 78%. Item non-response is extremely uncommon and rarely exceeds 1% for any question. The total survey time for the median respondent was 37 minutes; we will later report time spent on specific questions analyzed here as a measure of effort spent on acquiring and processing information.

2.1 Survey Module

Appendix E provides screenshots of the relevant module. The broad organization of the module was as follows:

1. **Stage 1 - Prior Belief**: This stage elicits individuals' expectations of future national home price changes. Respondents were informed that, according to Zillow, the median price of a home in the United States was $193,800 as of December 2016.[11] The respondents were asked for a point forecast: "*What do you think the value of the typical home in the U.S. will be at the end of this year (in December 2017)?*" To prevent typos in the responses, the survey environment calculated and reported the implied percentage change after individuals entered the value. Individuals could confirm the number and proceed to the next screen, or revise

[9]See Armona et al. (2017) and https://www.newyorkfed.org/microeconomics/sce/housing#main.

[10]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. (2016) for additional information.

[11]They were then asked how the price changed over the prior one year (since December 2015) and the prior ten years (since December 2006). They also were asked to rate their recall confidence on a 5-point scale.

their guess. We refer to the response to this question as the respondent's "prior belief." The survey also elicited the respondents' probability distribution over outcomes around their own point estimate: specifically, they were asked to assign probabilities to five intervals of future year-end home price changes: more than 10% below their point forecast; between 10% and 1% below their forecast; within +/-1% of their forecast; between 1% and 10% above their forecast; and more than 10% above their forecast.

2. **Stage 2 - Information Preferences**: After answering a block of other housing-related questions for roughly 15 minutes, respondents entered the second stage. They were notified that the same questions about future national home prices that were asked earlier in the survey would be asked again, except this time their responses would be incentivized: "*This time, we will reward the accuracy of your forecast: you will have a chance of receiving $[X]. There is roughly a 10% chance that you will be eligible to receive this prize: we will select at random 60 out of about 600 people answering this question. Then, those respondents whose forecast is within 1% of the actual value of a typical U.S. home at the end of this year will receive $[X].*" We randomly assigned half of the respondents to X=$100 ("High Reward") and the other half to X=$10 ("Low Reward").

   Before providing their forecast, respondents were given an opportunity to see a potentially relevant piece of information: "*Before you report your forecast, you will have the opportunity to see only one of the following pieces of information that may help you with forecasting future year-ahead U.S. home prices. Please rank the following pieces of information on a 1 to 4 scale, where 1 is "Most Preferred" and 4 is the "Least Preferred*":

   - *Change in the value of a typical home in the U.S. over the last one year (2016).*
   - *Change in the value of a typical home in the U.S. over the last ten years (2007-2016).*
   - *Forecasts of a panel of housing experts about the change in U.S. home prices over this coming year (2017).*
   - *None of the above – I would not like to see any information.*"

   Respondents were asked to drag and drop each of their selected rankings into a table with labels from "1=Most Preferred" to "4=Least Preferred."

3. **Stage 3 - Valuation of Information**: This stage, which immediately followed the second stage, elicited the respondents' valuation, or "willingness to pay" (WTP), for their highest-ranked information type. Respondents who ranked "*None of the above*" as their most preferred information in Stage 2 skipped this stage. To assess WTP, we used the list price version of the BDM method (e.g., Andersen et al., 2006) with eleven scenarios. In each scenario, respondents chose between seeing their preferred piece of information (i.e., the one they ranked highest in Stage 2) or receiving extra money in addition to their compensation

for completing the survey. The amount of money offered in these scenarios was predetermined and varied in \$0.50 increments, from \$0.01 (in Scenario 1) to \$5 (in Scenario 11). Respondents were told that one of these eleven scenarios would be drawn at random and the decision in that randomly chosen scenario would be implemented.

4. **Stage 4 - Posterior Belief**: In this stage, the respondent may have seen their highest-ranked information choice, depending on the randomly chosen scenario in Stage 3 and their choice of whether to see the information in that scenario or not.[12] Year-ahead home price expectations (the point estimate and the subjective belief distribution) that were elicited in Stage 1 were re-elicited from all respondents. We used the Zillow Home Value Index (ZHVI) as the source for prices of the typical (median) home in the U.S. over the last one or ten years.[13] According to the ZHVI, U.S. home prices decreased by 0.1% per year on average (or 0.9% in total) over the ten years 2007-2016 and increased by 6.8% over the prior one year (2016). The Zillow Home Price Expectations Survey, a quarterly survey of about 100 economists, real estate experts, and market strategists, was the source for the expert forecast.[14] On average, experts forecasted an increase of 3.6% in home prices during 2017. Note that all these information sources are publicly available.

   A paragraph providing the information followed a similar structure in all three cases. The raw information was provided, followed by a naive projection of home prices in December 2017 based on the annual growth rate implied by the information. For instance, respondents who chose expert forecast were presented with "*The average forecast of a distinguished panel of housing market experts who participate in the Zillow Home Price Expectations Survey is that home values in the U.S. will increase by 3.6% over the next year. If home values were to increase at a pace of 3.6% next year, that would mean that the value of a typical home would be 200,777 dollars in December 2017.*" At the bottom of this same screen, expectations about year-end home prices were re-elicited. Respondents were reminded about their prior belief. As in Stage 1, both the point estimate and subjective belief distribution were elicited. We refer to the point estimate from this stage as the "posterior belief."

Afterwards, respondents were picked at random to be eligible for the incentive, as indicated in Stage 2, and eligible respondents were informed at the end of the survey that they would be paid the \$10 (or \$100) reward in case of a successful forecast (within 1% of the December 2017 ZHVI) in early 2018.[15] At the end of the survey, respondents were also asked whether they had used any

[12]In Stage 3, the scenarios 1-11 were picked with probabilities 0.15, 0.14, 0.13, 0.12, 0.11, 0.10, 0.09, 0.07, 0.05, 0.03, and 0.01, respectively.

[13]For more information on the construction of the ZHVI, see http://www.zillow.com/research/ zhvi-methodology-6032/ (accessed on December 8, 2017). We used the ZHVI as of December 2016.

[14]For details, see https://pulsenomics.com/Home-Price-Expectations.php. We used the average forecast as of the fourth quarter of 2016.

[15]Payments to those who qualified and met the reward criterion were made in March 2018. 14 respondents received a payout (half of them \$100, and the other half \$10).

external sources (such as Google or Zillow) when answering any question in the survey.

Four months after the initial survey, a short follow-up was fielded to active panelists in the June 2017 SCE monthly survey. As in Stages 1 and 4 of the main experiment, respondents were asked to report their expectations about year-end U.S. median home prices. We kept the identical frame of reference in the follow-up survey: we provided individuals with the median U.S. home price as of December 2016 and asked them to forecast the value in December 2017. Both the point estimate and subjective density were re-elicited. Of the 1,162 respondents who took the SCE Housing Survey, 762 were still in the panel in June and hence eligible to take the follow-up survey. Of those, 573 did so, implying a response rate of 75.2%.

An additional module was fielded in the 2018 SCE Housing Survey. Since the main purpose of that module is some robustness checks and because that sample has no overlap with the sample in the original study, we defer the details to Appendix B.

2.2 Discussion of the Experimental Design and Outline of Analysis

Our design tries to mimic real-world information acquisition and processing, albeit in a stylized setting. Before turning to the empirical analysis, it is useful to discuss the features of the experimental design and to outline how our analysis will proceed.

A key feature of our setup is that respondents are presented with three possible pieces of information, which they are asked to rank in terms of their preference, including a no-information option. Ideally, we want to test whether individuals have some reasonable idea or consensus about the usefulness of the information. However, no single criterion can measure informativeness. One reasonable metric of information usefulness is how well the source has historically predicted past year-ahead home price changes in the United States.

Let HPA_t denote the actual home price change during year t. Let HPA_t^F be the mean forecast of experts about home price changes for year t, HPA_{t-1} the annualized home price change over the past 1 year, and HPA_{t-10} the annualized home price change over the past 10 years. For each piece of information $I_t \in \{HPA_t^F,\ HPA_{t-1},\ HPA_{t-10}\}$, we define its informativeness as the root mean squared error (RMSE) of a model $HPA_t = I_t$.

To calculate the RMSE of each piece of information, we use the Zillow Home Value Index as the outcome (that is, as our measure of HPA_t), because it is the same outcome that we are asking the subjects to forecast in our survey. Using this data, the RMSE for experts' forecast is 2.8, the RMSE for past one-year changes is 3.2, and the RMSE for past ten-year changes is 7.9 when using the longest available series (the experts' forecast is available since 2010, and the ZHVI since 1996). Based on these results, the expert forecast has been the most informative in predicting year-ahead home price changes, followed by past one-year change, and then the ten-year change. This ranking remains the same when we use only data since 2010 for all three series (in this case, the one-year RMSE is 3.3, and the ten-year RMSE is 5.2). Using a longer home price index series from CoreLogic (starting in 1976), the ranking also remains consistent.[16]

This criterion for ranking the informativeness of the signals is consistent with existing literature. First, the fact that the forecasts are ranked highest is consistent with the view that experts use all available information in past home price changes optimally when providing a forecast.[17] Additionally, this criterion is consistent with the model of Carroll (2003), in which consumers periodically update their expectations based on reports of expert forecasts, which are assumed to be rational. Second, the higher ranking of past one-year home price change relative to past ten-year change is consistent with the well-documented momentum in home prices over short horizons (Case and Shiller, 1989; Guren, 2016; Armona et al., 2017). For instance, for the nominal CoreLogic national home price index from 1976–2017, the AR(1) coefficient of annual growth is 0.73 and highly statistically significant, with an R^2 of 0.57. This serial correlation is only slightly weaker if we calculate price growth in real terms (the coefficient falls to 0.66 but remains highly significant).[18] In contrast, regressing one-year growth on growth over the previous ten years yields a small and insignificant negative coefficient.

Although reasonable, our criterion is not the only one that can determine the usefulness of information. For example, according to the ZHVI, U.S. home prices increased by 6.5% during 2017. Thus, based on ex-post accuracy, using the past one-year change would have led to the most accurate expectation. By this same ex-post metric, however, it remains hard to rationalize picking home price change over the past ten years over either of the other two pieces of information.

Our design elicits beliefs about national home prices. An alternative would have been to elicit beliefs about local home prices instead, such as the median price in the zip code or county where the respondent resides. Beliefs about local prices have the advantage that they may be of more direct relevance to the household's decision-making. However, we decided to elicit national home price expectations for two main reasons. The first is related to data availability: while surveys like the Zillow survey ask experts for national price forecasts, to the best of our knowledge, there is no comparable data where the same experts are asked for price forecasts at a local level. Second, using national home prices simplifies the research design enormously, because we do not have to deal with local heterogeneity. For example, when two individuals report different expectations about national prices, that difference can be interpreted directly as disagreement, because they are forecasting the same random variable. This is not the case when individuals report their own local home price expectations, because individuals living in different locations would be making forecasts about different random variables. If we wanted to measure disagreement about local forecasts,

[16]Using the CoreLogic series, the RMSE is 4.6 for the average expert forecast (6 observations), 5.0 for the past one-year change (39 observations), and 7.8 for the past ten-year change (30 observations).

[17]This should be true at least for the consensus forecast, even though individual forecasters may have incentives to deviate for strategic reasons (e.g. Laster et al., 1999).

[18]It is also robust to using alternative home price indices, such as Case-Shiller. Further, momentum is similarly strong at a more local level: Armona et al. (2017) find that in a regression of one-year home price changes on lagged one-year home price changes at the zip code level, the average estimate (across the zip codes in the U.S.) is 0.53 (statistically significant with $p < 0.01$).

we would need to restrict the comparison to pairs of individuals living in the same zip code or county, which would require a massive number of observations in a nationally representative sample. Likewise, using national home prices keeps the interpretation of the heterogeneity in choices in our experiment clean. For example, given geographic variation in home price movements in the US, different information sources may have been more or less informative across locations, something that we would have had to account for if we had elicited local home price expectations.

2.2.1 Outline of Analysis

Our analysis will begin with a description of the sample and the distribution of prior beliefs in the initial stage of the survey module. Then, we will proceed in three stages. First, we will study what information sources people rank highest, and whether there is systematic heterogeneity. As discussed above, an ex-ante evaluation would prescribe that the expert forecast should be the preferred choice, and that ten-year past growth should be least preferred.

Second, we will investigate what determines people's valuation of the information, and how they use the information if they obtain it. In terms of valuation, rational inattention predicts that individuals invest more resources in acquiring and processing house-price-related information if they have a stronger incentive to have accurate expectations.[19] The randomization of the accuracy incentive in Stage 2 provides a direct test of this prediction. In addition, we will study how the valuation of information relates to other (non-randomly-allocated) proxies for stakes, such as a high likelihood of a purchase of a home in the near future, and to measures of confidence in one's existing knowledge of the housing market. We will further study how those respondents who get to see information use it. Here, we will exploit that our design generates random variation in the provision of information, since whether a respondent gets to see their preferred piece of information depends on their valuation but also on which scenario from Stage 3 is randomly picked. We will test whether respondents incorporate the signal into their posterior beliefs, as would be expected if individuals were willing to pay for the information. Rational updating also implies that individuals who have uncertain prior beliefs put more weight on the information they receive.[20] As an alternative metric for how people use information, we will further rely on the time spent on forming the posterior belief.

Third, we will investigate how a change in the price of information affects the cross-sectional dispersion of expectations. In a pure sticky information model (Reis, 2006), cross-sectional dispersion in beliefs arises only because some individuals update their information sets to perfect

[19] This would follow from most "rational inattention" models in the broad sense of Footnote 1. For example, in the sticky updating model of Reis (2006), agents are modeled as maximizing utility subject to constraints, which also include costly information. Increasing the payoff for more accurate expectations would lead more agents to incur the cost of acquiring housing-relevant information.

[20] Under Bayesian updating, the weight put on the signal is positively related to the uncertainty in the prior belief, and inversely related to the (perceived or actual) noise in the signal. As long as the perceived noise in the signal is independent of one's uncertainty in the prior belief, Bayesian updating predicts that individuals with more uncertain priors put more weight on the signal.

information, while other individuals do not update their information sets. If all individuals updated at a given point in time, there would be no cross-sectional dispersion in beliefs. When only some individuals update at a given point in time, there is no cross-sectional dispersion in beliefs among those who update. Thus, one might expect that, as the price of information decreases, more people update and the dispersion of expectations decreases. Since our design randomly varies the effective price of information (due to randomness in which scenario is picked), we can directly test this hypothesis.

2.3 Sample Characteristics

Of the 1,162 valid responses, we trim the sample by dropping 43 respondents: those with prior beliefs below the 2.5th percentile (an annual growth rate of -7.1%) or above the 97.5th percentile (an annual growth rate of 16.1%). These extreme beliefs may be the product of typos or lack of attention. As the prior belief was reported before the treatments, dropping these extreme prior beliefs should not contaminate the experimental analysis. Typos may also be present in posterior beliefs, but dropping individuals based on post-treatment outcomes could contaminate the experimental analysis. Instead, we winsorize the post-treatment outcomes using the same extreme values presented above (-7.1% and 16.1%).[21] In any case, we use graphical analysis whenever possible to certify that the results are not driven by outliers.

Column (1) of Table 1 shows characteristics of the sample for the main survey. Most dimensions in the sample align well with average demographic characteristics of the United States. For instance, the average age of our respondents is 50.8 years, and 47.4% are females, which is similar to the corresponding 45.5 years and 48.0% among U.S. household heads in the 2016 American Community Survey (ACS). Also, 74.8% of respondents in our sample are homeowners, somewhat higher than the national homeownership rate of 63.6% in the first quarter of 2017, according to the ACS. Our sample, however, has significantly higher education and income: 55.2% of our respondents have at least a bachelor's degree, compared to only 37% of U.S. household heads. Likewise, the median household income of respondents in the sample is $67,500, which is substantially higher than the U.S. 2016 median of $57,600. This may be partly due to different internet access and computer literacy across income and education groups in the U.S. population. The survey also included a battery of 5 questions taken from Lipkus, Samsa, and Rimer (2001) and Lusardi (2008) that measure the respondent's numeracy. On average, respondents answer 4 questions correctly. The rank correlation between education and numeracy in the sample is +0.31. Underscoring the fact that home prices are something that individuals actively think about, more than half of the sample – 56.3% – reports looking up home price information over the past 12 months.[22] The

[21]For the beliefs from the follow-up survey, we winsorize the values in the same way. Results are robust under alternative thresholds.

[22]Specifically, they are asked about how often, over the past 12 months, they have consulted websites or other sources that provide information on the estimated current value of their property or properties in their area.

average reported probability of moving and buying a different home over the next 3 years in the sample is 20 percent.

Columns (2) and (3) of the table show average characteristics for the subsamples assigned to the low- and high-reward treatments, respectively; in turn, columns (5) and (6) show the characteristics for the subsamples assigned to low and high realized prices of information. Columns (4) and (7) present p-values for the test of the null hypothesis that the characteristics are balanced across groups. The differences in pre-treatment characteristics are always small, and statistically insignificant in 31 out of the 34 tests. This is not surprising, because random assignment should preserve balance between the two groups. Additionally, the last row of Table 1 reports the response rate to the follow-up survey. There is no evidence for selective attrition: the response rate does not differ by reward or price treatments. Table A.1 provides additional information on how the follow-up sample compares with the initial sample. There is no evidence of selection in terms of who is invited to the follow-up, or that, conditional on being invited to the follow-up, the individuals who responded to this survey are significantly different from the ones who did not.

3 Empirical Analysis

Figure 2.a shows a histogram of the point estimates provided by respondents in the initial stage of the survey module, that is, prior to the information acquisition stages. In terms of the implied annual growth rates, the mean (median) value is 2.2% (1.7%), with substantial dispersion across respondents: the cross-sectional standard deviation of prior beliefs is 3.1%. To assess if individuals felt confident about their expectations, Figure 2.b shows the probability distribution of beliefs around the individual's own point estimate, averaged over all individuals. On average, individuals thought there was a 51 percent chance that the true price would fall within 1% of their guesses. Moreover, there was high dispersion in the degree of certainty. For example, 13% of the sample thought that there was a 90 percent chance or higher of year-end home prices being within 1% of their guess, and 16% of the sample thought that there was a 20 percent chance or lower.[23] We use the responses to the probability bins to measure prior uncertainty at the individual level, which we use in the analysis below. We fit the binned responses to a normal distribution for each individual and use the estimated standard deviation of the fitted distribution as a measure of individual-level uncertainty, with higher values corresponding to higher uncertainty.[24]

[23]Ex post, only 3.5% of respondents had a prior forecast within 1% of the realized ZHVI price as of December 2017, which was $206,300 (according to Zillow in January 2018), corresponding to realized growth over 2017 of 6.5%. For the posterior forecast, this fraction increased to 11.5%.

[24]For instance, consider an individual with a 2% house price growth point forecast who has an uncertainty (i.e., fitted standard deviation) of 1 percentage point. It means that the individual's 95% confidence interval for house price growth is [0.04%, 3.96%] ($= [2-(1*1.96), 2+(1*1.96)]$). In cases, where the respondent puts all mass in one bin or equal mass in two adjacent bins, a uniform distribution is fit. We are unable to fit a density for 4 individuals because of missing data. The uncertainty data is winsorized above the 98.5th percentile of the prior uncertainty distribution.

3.1 Ranking of Information Sources

What happens when individuals with uncertain prior beliefs are offered the chance to acquire information? Figure 3.a shows the ranking distribution for the different information types over the whole sample. Individuals disagreed on which of the three pieces of information they would want to see: 45.5% chose forecasts of housing experts, 28.1% chose the last-one-year home price change, 22.1% chose the last-ten-year home price change, and the remaining 4.3% preferred no information.[25] The past predictive power criterion indicated that expert forecast was most informative, followed by the last-one-year home price change and then the last-ten-year home price change. Thus, the popularity of the choice is increasing with its informativeness. However, this correlation is far from perfect: less than half of the sample chose the most informative signal (i.e., expert forecast).

This heterogeneity in the ranking of information could be driven by consumers' lack of knowledge about the relative informativeness of the signals or by respondents using different criteria to determine the informativeness of the signals. Systematic differences in ranking by education or numeracy of respondents, which are reasonable proxies for ability to filter signals, would suggest evidence of the former. Figure 3.b and 3.c thus break down the information choices by respondents' numeracy and education, respectively, and show that individuals with more education or with higher numeracy were substantially more likely to choose the "best" information: college graduates chose the expert forecast 50% of the time, compared with non-graduates who chose it 40% of the time (p-value<0.01).[26] They are also substantially less likely to choose not to see any information, or choose the past ten-year home price change, the least informative signal.

Table 2 further explores the heterogeneity and reports univariate relationships between the choice of information and various individual- and location-specific characteristics.[27] The dependent variables in columns (1)–(4) correspond to dummy variables indicating the highest ranked piece of information.[28] Besides numeracy and education of respondents, only a handful of variables are significant, suggesting that observable characteristics (at the individual or location level) cannot explain much of the heterogeneity in how individuals rank information. Homeowners and individuals who are likely to be active housing market participants, as measured by the reported probability of moving and buying a home within the next 3 years, are less likely to choose no information source (column 4). Conditional on choosing an information source, homeowners are 8.8 percentage points more likely to choose the past-one-year information (and less likely to choose

[25] The median respondent spent 2.17 minutes choosing between the information sources (and reading the associated instructions), with the 10th percentile at 1.23 minutes and 90th percentile at 4.85 minutes.

[26] Similarly, Burke and Manz (2014) find that respondents with higher levels of economic literacy choose more relevant information when forming inflation forecasts.

[27] The results are qualitatively similar using multivariate regressions, as reported in Appendix Table A.2. Fewer estimates are, however, statistically significant in the multivariate set-up due to the dependence between the covariates.

[28] The results are also robust if instead of a linear probability model we use a multinomial logit model.

the expert forecast), perhaps because they are curious to learn about how their housing wealth has evolved. On the other hand, active housing market participants are much more likely to choose the forward-looking signal: they are 16 percentage points more likely to choose the expert forecast (and significantly less likely to choose the past one-year information). Higher income respondents and white respondents are more likely to choose the expert forecast, although these coefficients are only marginally significant.

One might expect respondents who have high confidence in their perceptions of past home price changes to be more likely to choose the expert forecast (since they think they know the past realized growth already); however, that is not the case. Likewise, one might expect respondents residing in states with volatile housing prices (as measured by the standard deviation in monthly home prices over the past 24 months) to be less likely to choose past home price changes. We do not find evidence of that. Similarly, individuals who report looking up information about home prices in the past 12 months might already be well-informed and so less likely to choose an information signal. We in fact see the opposite, suggesting that such individuals have a positive taste for information. Along the same lines, the individuals who report looking up information during the survey are, relative to their counterparts, no less likely to choose no information.[29] Conditional on picking an information source, they are significantly more likely to choose the expert forecast – the information source that would arguably be the hardest to find. Notably, we see that being assigned to the High Reward treatment does not have any systematic impact on the ranking of information sources; this is also shown in Figure 3.d.

Since both the expert forecast and the past one-year home price change could be argued to be the most informative signals, in column (5), as an alternative outcome we study whether a respondent ranked either the expert forecast or the one-year realized growth as their top choice. Other than education and numeracy, we see little relation with observables.

The supplementary survey that was conducted in 2018 provides some additional insights, which are discussed in detail in Appendix B. First, we validate the finding that subjects disagree in terms of the information that they acquire, and that those disagreements are correlated with education and numeracy. Second, the supplementary survey included a couple of additional questions to explore the role of trust in experts as a driving factor for preferences over information sources. Overall levels of trust in the credibility of experts and their ability to forecast accurately is moderate, and we do find that less-educated respondents exhibit lower levels of trust in experts. However, while a relevant explanation, distrust of experts is not the main factor driving the information choices of our respondents: for instance, we find that these differences in trust can explain less than a quarter of the education gap in preferring experts.

We can summarize our first result as follows:

[29] At the end of the survey, we asked respondents whether they had searched for information online during the survey, explaining that doing so was permitted. 14% of the sample reported doing so. Interestingly, the search rate was not statistically different for respondents who were assigned the high reward treatment (14.7%) and those assigned the low reward treatment (13.3%); p-value=0.49).

Result 1*: The information with the highest ex-ante predictive power, expert forecast, is the modal choice. The information with the second highest ex-ante predictive power is the second most frequent choice. Considerable disagreement exists across households on the ranking of information. The ranking is systematically related to measures of respondent ability, which suggests that cognitive limitations in deciphering the informativeness of signals partially drives the heterogeneity.*

3.2 Valuation and Use of Information

3.2.1 Valuation of Information

Using responses to the eleven scenarios in Stage 3, we identify the range of an individual's valuation or willingness to pay (WTP). For example, if an individual chose information instead of any amount up to \$3 and then chose the money from \$3.5 on, it means that the individual's WTP must be in the range \$3 to \$3.5.[30] Around 5% of respondents provided inconsistent responses; for example, they chose information instead of \$3 but then chose \$2.5 instead of information. This inconsistency is within the range of other studies using this method for elicitation of WTP for information. For instance, the share of inconsistent respondents was about 2% in Allcott and Kessler (2015) and 15% in Cullen and Perez-Truglia (2017).

Figure 4.a shows the histogram of WTP for the whole sample (excluding inconsistent respondents). We find that individuals have significant WTP for their favorite information, with a median WTP between \$4.5 and \$5. An alternative estimate is given by means of an interval regression model. This is a maximum likelihood model that assumes that the latent WTP is normally distributed. The constant in this model is estimated to be \$4.16 (95% CI from 3.94 to 4.40). This coefficient can be interpreted as the mean WTP under the implicit assumption that WTP can take negative values; if we instead assume that the WTP must be non-negative, then the mean would be even higher. This is a fairly high WTP, given that the information we provide is public and readily available using a search tool like Google. This finding indicates that most individuals are either unaware of the availability of this information or they expect a high search cost. Also, the median WTP (\$4.5-\$5) is very high compared to the expected reward for perfect accuracy (\$1 for half of the sample and \$10 for the other half). This evidence suggests that individuals value the information beyond the context of the survey. They may want to use this information for real-world housing decisions. In this context, having incorrect expectations about house prices can translate into thousands of dollars in losses, relative to which the experimental incentive pales in comparison.[31]

[30]Individuals who ranked "no information" highest in Stage 2 are assigned a WTP in the interval [-infinity, 0].

[31]We can also compare the median WTP in our study (\$4.5-\$5) with the results from a few other papers that elicit WTP for information using similar methods. Those studies find lower valuations: \$0.40 for travel information (Khattak, Yim, and Prokopy, 2003), \$0.80 for food certification information (Angulo, Gil, and Tamburo, 2005), and \$3 for home energy reports (Allcott and Kessler, 2015). Hoffman (2016) finds that, in a field experiment, business

We next study whether WTP systematically varies with reward size and other factors. To directly test the effect of stakes, Figure 4.b compares the distribution of WTP between the two reward groups. This figure shows that, as expected, individuals in the higher-reward treatment are willing to pay significantly more. A Mann-Whitney-Wilcoxon (henceforth MWW) test indicates that the difference is significant at $p<0.01$.

To better understand the economic magnitude of this difference, we regress the WTP onto a constant and a High Reward dummy. The constant in this regression can be interpreted as the mean WTP for the low-reward condition ($10 with 10% probability). This average valuation is estimated to be $3.75 (95% CI from 3.50 to 4.10). The estimated coefficient on High Reward is 0.83 (95% CI from 0.34 to 1.32), indicating that, relative to the $10 reward, individuals assigned to the $100 reward are willing to pay an additional $0.83 for their favorite information (or 22% more). Note that the expected reward goes from $1 to $10, because the reward is given only with 10% probability. The $0.83 difference in WTP then implies that for each additional dollar of expected reward, the WTP for information goes up by 9.2 cents. Equivalently, under the assumption of risk neutrality, the average individual expects that the probability of being accurate (i.e., being within 1% of the realization) will increase by 9.2 percentage points, or 18% of the baseline probability, if they acquire the information.[32]

Column (1) in Table 3 shows the univariate relationship between the WTP and various correlates. Here, we discuss the more notable and interesting ones. The expected effect of past information acquisition effort on WTP is ambiguous. On the one hand, individuals who looked for information in the past may be willing to pay less for the information, because they already have good information. On the other hand, individuals who acquired more information in the past may have the highest revealed demand or "taste" for information and thus could be more willing to buy additional information. Our evidence suggests that the second channel dominates: individuals who looked for housing-related information in the past were willing to pay an additional 77 cents, relative to those who did not.

Likewise, we can study how the uncertainty in prior belief correlates with WTP. We again find evidence for the "selection channel": individuals with a one-standard deviation higher uncertainty

---

experts tend to underpay (overpay) for information when signals are informative (uninformative).

[32]To see this, note that one can write the WTP as:

$$WTP_i = U_{Info} + 0.1 \cdot Reward_i \cdot \left[P_i\left(Accurate|Info\right) - P_i\left(Accurate|NoInfo\right)\right] + \varepsilon_i.$$

The first term, U_{Info}, represents the expected real-world benefit from having the information (e.g., because one expects to make better choices when deciding whether to buy a house). The second term reflects the benefits of information from the survey reward, under the simplifying assumption that the respondent is risk-neutral for small amounts. We can infer the value of $P_i\left(Accurate|Info\right) - P_i\left(Accurate|NoInfo\right)$ from the regression of WTP_i on $Reward_i$, reported in column (1) of Table 3. The coefficient on the High Reward dummy of 0.83 indicates that increasing $0.1 \cdot Reward_i$ by 9 (i.e., $0.1 \cdot 100 - 0.1 \cdot 10$) increases the WTP by $0.83. Thus, increasing $0.1 \cdot Reward_i$ by 1 would increase the WTP by 0.092 ($= \frac{0.83}{9}$). This estimator implies that $P_i\left(Accurate|Info\right) - P_i\left(Accurate|NoInfo\right) = 0.092$. In the first stage, the average individual responded that there was a 51.3% chance that their guess is within 1% of the true price. We use this as an estimate of the average $P_i\left(Accurate|NoInfo\right)$. Thus, the 9.2 percentage point increase translates into a increase of 18% (=9.2/51.3) in the baseline probability.

in their prior beliefs were, on average, willing to pay \$0.28 less.[33] Similarly, individuals who are more confident in their perceptions of past home price growth are willing to pay more for information. In sum, the evidence strongly points to heterogeneity in the taste for information, and this channel will be a key element of the model presented in Section 4.

The expected effect of local volatility in home prices on WTP also is ambiguous. On the one hand, updating more often is valuable for respondents in volatile locations, and hence they should value information more. On the other hand, past changes in home prices are less informative. We have seen that respondents in these areas do not choose expert forecast more often. Here, we see that these respondents in fact value information more: increasing the home price volatility by 1 standard deviation increases the WTP by 25 cents.

Finally, we study whether WTP differs depending on the information source a respondent chose. We argued earlier that based on past performance, individuals should value experts' forecasts most highly and ten-year growth least highly. However, if individuals select a given information source because they erroneously believe it to be the most accurate/predictive one, then WTP would not differ by information source. Indeed, we find no evidence that individuals pay more for information that has higher ex-ante predictive power. Panels c and d of Figure 4 show how WTP for the expert forecast compares with that for past-one-year and past-ten-year home price changes, respectively. In a simple interval regression similar to the ones above, average WTP is highest for the ten-year information, followed by the expert forecast and the one-year information; the difference between ten-year and one-year information is significant at $p<0.05$ (while the coefficient on the expert forecast is not significantly different from either of the others).

The first two columns of Table A.3 show the correlates of the WTP in a multivariate framework. When we simultaneously control for other variables in column (2) of Table A.3, we see that the WTP is significantly higher for those respondents who chose the ten-year information. The other patterns are qualitatively similar to those shown in the first column of Table 3.

3.2.2 Use of Information

We now turn to the question of how people incorporate their preferred information, if they obtain it, in their posterior forecast. To do so, we first study the updating in the point forecast, looking at learning rates overall and across different sub-populations. Second, we use the time spent on reporting the posterior forecast as another measure of updating effort.

To study updating, we use a simple learning model that naturally separates learning from the signal shown from other sources of reversion toward the signal.[34] Let b^{prior} denote the mean of the prior belief, b^{signal} the signal, and $b^{posterior}$ the mean of the corresponding posterior belief.

[33]Note that the correlation of prior uncertainty with education/numeracy as well as with looking up housing-related information in the past is negative. This further suggests that the selection channel – of people genuinely interested in information having more precise priors and willing to pay more for information – being the dominating factor.

[34]Similar learning models are used in Cavallo et al. (2017).

When priors and signals are normally distributed, Bayesian learning implies that the mean of the posterior belief should be a weighted average between the signal and the mean of the prior belief:

$$b^{posterior} = \alpha \cdot b^{signal} + (1-\alpha) \cdot b^{prior}. \tag{1}$$

In a Bayesian framework, the weight parameter α increases in the uncertainty (i.e., the variance) of the prior and decreases in the uncertainty and noise in the signal. This parameter can take a value from 0 (individuals ignore the signal) to 1 (individuals fully adjust to the signal). Re-arranging this expression, we get:

$$b^{posterior} - b^{prior} = \alpha \cdot \left(b^{signal} - b^{prior}\right). \tag{2}$$

That is, the slope between the perception gaps ($b_i^{signal} - b_i^{prior}$) and revisions ($b_i^{posterior} - b_i^{prior}$) can be used to estimate the learning rate.[35] However, it is possible that individuals will revise their beliefs towards the signal even if they are not provided with it. For instance, consider someone who makes a typo when entering her prior belief and reports an estimate that differs significantly from the signals. If that person does not commit the typo again when reporting the posterior belief, it will look like she is reverting to the signal despite not being shown information. Also, it is possible that individuals think harder the second time they are asked about their home price expectation, especially since the posterior belief is incentivized but the prior belief is not. Similarly, 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 expectations even if they are not provided with any new information (see Zwane et al., 2011, for related evidence).

Recall that our design generates random variation in whether a respondent saw information. We can use this variation to separate true learning from mean-reversion. Denote by S_i a dummy that takes the value 1 if the individual was shown the signal, and let WTP_i be a set of dummies corresponding to the "threshold price" chosen by the individual in the scenarios. Conditional on this threshold, whether the individual received the information ($S_i = 1$) depends on the randomly chosen scenario. Thus, we use the following regression specification:

$$b_i^{posterior} - b_i^{prior} = \alpha \cdot \left(b_i^{signal} - b_i^{prior}\right) \cdot S_i + \beta \cdot \left(b_i^{signal} - b_i^{prior}\right) + WTP_i\delta + \varepsilon_i. \tag{3}$$

The parameter of interest is still α, which measures the true learning rate (i.e., the effect of being

[35] There is an alternative specification for this learning model. Consider the case when the information chosen is the past 10 year home price change. b^{signal} is the actual past 10 year change, and $b_i^{\hat{signal}}$ is i's prior belief about the past 10 year home price change, that was also elicited in the first stage of the survey. $b^{signal} - b_i^{\hat{signal}}$ is then the difference between the actual change and the perceived change. The revision in expectations can be regressed onto this metric (this kind of learning model has been used in Armantier et al., 2016, and Armona et al., 2017). We do not use this alternative model for two reasons. First, this alternative model cannot be estimated for one of the data sources, because we did not elicit the prior belief about the signal of experts. Second, when considered simultaneously in the regression analysis, our baseline model fits the data better than this alternative specification.

randomly shown information on the updates). β reflects the degree of spurious mean-reversion. Figure 5.a shows the results from this regression. The y-axis indicates the revision in the forecast (i.e., posterior belief minus prior belief). The x-axis shows the "gap" between the signal and the prior belief, interacted by the treatment assignment dummy. For instance, if the respondent had a prior belief of 1% and was shown the expert forecast (which was 3.6%), the x-axis would take the value of 2.6%. Intuitively, the x-axis shows the potential for revision, and the y-axis shows the actual revision. The slope of the line is 0.38, which is highly statistically significant (p-value<0.001) and economically substantial: the average individual puts 38% weight on the signal and 62% on their prior belief.[36]

One potential concern with survey experiments is that, instead of inducing genuine learning, the information provided in the experiment elicits spurious reactions—for instance, due to unconscious numerical anchoring (Tversky and Kahneman, 1974) or experimenter demand (Goffman, 1963). Following Cavallo et al. (2017) and Armona et al. (2017), we use the follow-up survey to address this potential concern: if the reaction to the information was completely spurious, then the experimental effects should not persist for months after the information provision. To do this, Figure 5.b reproduces Figure 5.a, but instead of using $b_i^{posterior} - b_i^{prior}$ as the y-axis, we use $b_i^{follow-up} - b_i^{prior}$, where $b_i^{follow-up}$ is the belief reported four months later. We see that the estimated slope in the follow-up (0.173) is smaller than in the main survey (0.380), but it is still economically meaningful and statistically significant at the 10% level. This suggests that a significant part of the reaction to the information was not spurious. Also, note that the slope is expected to be lower in the follow-up, because individuals may have been exposed to additional signals during the interim four months, thus gradually diluting the effect of the signal provided during our experiment.

We next study heterogeneity in learning rates. Figure 6.a investigates whether the learning rates differ across the three pieces of information. Ex ante, there is little reason for rates to differ: once respondents reveal their information preference, they should be equally responsive to it. This is confirmed in the figure. Panels b and c of Figure 6 investigate whether the learning rate differs by WTP for information or by the reward size. One might expect respondents who valued the information more to put greater weight on the signal; however, we do not find evidence for differential slopes. Similarly, while the slope in the high reward treatment is directionally stronger, the difference is far from statistically significant.

In contrast, we do find significant differences in updating depending on the uncertainty in respondents' prior belief. Under Bayesian updating, respondents who were more uncertain should put more weight on the signal (as long as the prior uncertainty is orthogonal to the perceived signal noise). However, Figure 6.d shows that respondents with higher prior uncertainty tend to update *less*. While this is surprising based on the standard Bayesian intuition, it is arguably consistent with the earlier result that respondents with higher prior uncertainty have a lower WTP for information. We will return to this in the theory in Section 4. Another result that the

[36] Appendix A presents additional details on the estimation.

theory will be able to rationalize is that high-numeracy respondents tend to update more strongly than low-numeracy respondents (panel e). In contrast, updating rates do not differ by respondent education (panel f).

As an alternative measure of how people allocate effort and use information, we rely on the time spent reporting the posterior belief. Column (2) of Table 3 uses the log time spent on the screen used to report the posterior belief as the dependent variable.[37] Due to the design of the survey, this variable includes the time spent looking at the information, and therefore respondents who get to see information would mechanically take more time reporting the posterior belief. Thus, in column (3), we restrict to the sample that actually got to see information. We see that individuals assigned to the higher reward spent significantly more time reporting their posterior belief: an additional 11% when comparing to the full sample (column 2), and 12% more when just looking at respondents who saw information (column 3). Similarly, higher WTP for information is associated with significantly more time spent on forming the posterior, even when we just look at respondents who saw information. These two variables thus do seem to correlate with effort spent on updating, even though they were not associated with differential updating above.

Reinforcing the earlier finding that lower-uncertainty individuals may be those who have a taste for information, and that they update more strongly in response to the information, we see that they also spend significantly more time reporting their posterior belief. This remains true in a multivariate version of the same analysis, shown in Table A.3. Higher numeracy or education are not associated with time spent on the posterior forecast. In terms of other variables, we see that higher-income individuals spend less time on the screen, while the 14% of participants who later report having looked up other sources during the survey spent substantially more time (about 30%) on the screen—this suggests that they looked up the information during this time.

We summarize the findings in this subsection as follows:

> ***Result 2****: Respondents put value on information that can help them form more accurate forecasts, and incorporate signals if they obtain them. The valuation of information increases with incentives for accuracy, but is lower for people with higher prior uncertainty. Furthermore, contrary to standard models of rational updating, individuals with higher prior uncertainty put less weight on the signal (and spend less time on forming their posterior belief).*

3.3 Information Acquisition and Dispersion of Expectations

In this subsection, we study how information acquisition affects dispersion in beliefs. We begin by investigating the effect of an exogenous reduction in the cost of information. In Stage 3, a scenario is picked at random. Thus, the experimental setup induces exogenous variation in the

[37]We winsorize this variable at the 1.5th percentile (0.46 minutes) and at the 98.5th percentile (18.79 minutes) of the distribution.

cost of information. We exploit this and compare how beliefs evolve when "low-price" (\$0.01–\$1.5) scenarios are picked at random, versus "high-price" (\$2–\$5) scenarios. Table 4 presents the results from this test. First of all, notice from the first row of the table that the lower cost of information did result in more information acquisition: the share of individuals acquiring information is 21 percentage points higher in the low-price group relative to the high-price group.

The rest of the rows from Table 4 show how beliefs evolved for the low- and high-price groups. As expected (due to the scenario being picked at random), the distribution of prior beliefs for the two groups is similar. At the final stage, due to the belief updating of those who saw the signal (as studied above), the mean forecast increased and uncertainty decreased. However, even though a significantly higher share of respondents in the low-price group saw a signal, the dispersion in beliefs remains similar across the two groups. In particular, we do not find evidence that the mean absolute deviation (MAD) is lower for the low-price group: in fact, it is slightly higher, at 2.21, than for the high-price group, which has a mean absolute deviation of 2.13 (the difference is not statistically significant at conventional levels; p-value=0.59).

We also study an additional measure of disagreement, defined as follows: for each respondent, we construct a 95% confidence interval for their forecast based on their point forecast along with the reported uncertainty.[38] We then form all possible pairs of respondents within a group (here, the low-price and high-price groups) and define a disagreement as occurring for a pair if the two respondents' constructed confidence intervals do not overlap. This measure thus reflects effects of information both on the dispersion in point forecasts and on respondents' uncertainty. In Table 4, we see that the fraction of disagreements roughly doubled from the prior stage to the posterior stage, primarily because respondents' uncertainty went down. However, we again see that disagreement is almost exactly at the same level for the group with a low cost of information, which was much more likely to obtain the signal, than for the group with a high cost of information.

How is it that more information does not induce higher consensus? Figures 7 and 8 explore this question. Figure 7 shows the distribution of prior beliefs for individuals who were not shown the information (Figure 7.a) versus individuals who were shown the information (Figure 7.b). Comparing the two indicates that these two groups started with similar distributions of beliefs. Figure 8 shows the comparison of posterior beliefs between individuals who were not shown information (Figure 8.a) versus individuals who were shown the information (Figure 8.b). Figure 8.a shows that, among individuals who were not shown information, the distribution of posterior beliefs is the same regardless of whether the individuals preferred the expert forecast, past-one-year home price change, or past-ten-year home price change.[39] In contrast, Figure 8.b shows that, for in-

[38]Note that our results are qualitatively unchanged if we use confidence intervals with different coverage, e.g. 90% or 50%.

[39]Consistent with the earlier discussion that subjects in the no-information group may have searched for information or thought harder about the question, a comparison of Figure 7.a versus Figure 8.a indicates that the distribution of beliefs changed from prior to posterior even for individuals who were not shown information ($p<0.01$, MWW test).

dividuals who saw the information, posterior beliefs were substantially different across the three information groups. In each group, posterior beliefs moved towards the values of the respective signals: that is, -0.1% for the ten-year price change, 3.6% for the expert forecast, and 6.8% for the one-year price change. Within a group, the revelation of information tended to decrease dispersion of expectations. However, because the three groups moved towards differing signals, the dispersion in beliefs across groups increased. The net effect of information acquisition on belief dispersion depends on the combination of these two channels, which end up canceling each other out.

Table 5 provides a more quantitative version of the previous graphical argument. The first two columns of Table 5 describe prior and posterior beliefs, respectively. It is worth remembering that whether or not a respondent sees information is endogenous to their WTP, which in turn may reflect other characteristics. Therefore, the comparison here is not as "clean" as the one in Table 4, which relies on experiment-induced variation in whether a respondent saw the information.

We are primarily interested in one feature of these beliefs: the dispersion, measured by the mean absolute deviation (MAD) across individuals. The first thing that we can corroborate is that, within information groups, information provision tended to reduce belief dispersion (but belief dispersion remained high). For instance, for individuals who preferred the forecast and were shown the information, the MAD decreased from 2.19 to 1.14 percentage points. In contrast, for individuals who preferred the expert forecast but did not get to see the information, the MAD in beliefs increased from 1.93 percentage points for prior beliefs to 2.38 for posterior beliefs. These qualitative differences are consistent inside the group of individuals who chose the 10-year information, while for those who chose the 1-year information, MAD increased even for those who saw the info (perhaps because the signal was rather "extreme"), though less so than for those who did not see the info.

Now we turn to the sample that pools the individuals across all three information sources. In this pooled sample, the group that saw the information did not see a decline in the mean absolute deviation of beliefs: this measure of dispersion is 2.04 percentage points for the prior beliefs and 2.05 percentage points for the posterior beliefs.[40] And since respondents became more confident in their forecast, disagreement substantially increases, from 11.6% to 20.8% of all pairs. Disagreement also increases similarly in the no-information group; the difference-in-differences across groups is not statistically significant. Regarding confidence in expectations, mean uncertainty in posterior beliefs is lower than that in prior beliefs for both groups (those who saw the information and those who did not). However, consistent with the notion that information should make individuals more certain, we see that uncertainty declined more for the group that was shown information (from 4.0 to 2.9 percentage points, or more than 1 percentage point) than for the group that was not shown information (from 4.3 to 3.8 percentage points).

The table also shows how the cross-sectional dispersion persists over time, by looking at the

[40] The mean absolute deviation in the pooled sample that does not see information does go up (from 2.17 to 2.64), with the difference statistically significant (p-value<0.01).

follow-up survey. It is not clear whether we should expect dispersion to persist – in the interim four months, individuals may have received various signals. Depending on the heterogeneity in these signals, the cross-sectional dispersion may go up or down. Additionally, because individuals are being asked about year-end home prices, some uncertainty may have resolved over the interim four months. The last column in Table 5 shows how these statistics evolved for the information-shown and not-shown groups. Comparing the follow-up belief with the posterior belief, the mean absolute deviation in expectations increases for both the information-shown and not-shown groups, though more so for the former group.

One potential concern is that the cross-sectional dispersion does not decrease when information is cheaper just because respondents could buy, at most, one of the three information pieces. Could allowing individuals to view multiple pieces of information reverse this result? To investigate this, we fielded a supplementary module in the 2018 SCE Housing Survey. Details of this module and the analysis are presented in Appendix B. In this supplementary study, respondents can choose between two pieces of information. Then, we randomize them into three groups: they get to observe either no information, their preferred information, or both pieces of information (for free). The comparison between no information and their preferred information is equivalent to the comparison from the main experiment (i.e., randomizing the price of the preferred information between zero or infinity). We corroborate the finding from the main survey: cross-sectional dispersion does not decline when subjects get to see either their preferred information or both pieces of information. We find that randomly providing two signals at the same time has effects similar to providing just one piece of information, and that cross-sectional dispersion (measured either by the MAD or the disagreement metric) does not go down.

The third set of findings can be summarized as follows:

> ***Result 3***: *A lower cost of information does not lead to a decrease in the cross-sectional dispersion of beliefs. This finding arises for two reasons. First, individuals choose to acquire different pieces of information and put significant weight on the acquired information. Second, within groups of individuals who acquire the same information, dispersion in beliefs remains high.*

4 A Model of Information Acquisition and Processing

In this section, we present a simple model that can match most of the experimental findings. The model is a combination of a sticky information model, as in Reis (2006), and a rational inattention model, as in Sims (2003). The model also features heterogeneous prior beliefs about the quality of different information sources.

4.1 Setup of the Model

Individual i has the prior belief that the fundamental θ is normally distributed with mean $\mu_\theta(i)$ and variance $\sigma_\theta^2(i)$, where the i indicates that the prior belief over the fundamental may differ across individuals. In the experiment, the fundamental is one-year future house price growth.

Individuals start with a common prior belief that each information source $j \in \{1, 2, 3\}$ is a noisy signal on the fundamental

$$x_j = \theta + \varepsilon_j,$$

where x_j is the displayed information and the noise ε_j is normally distributed with mean zero, but individuals may have heterogeneous prior beliefs over precisions. Individual i believes that the precisions of the three information sources equal $\tau_1(i) \equiv (1/\sigma_{\varepsilon,1}^2(i))$, $\tau_2(i) \equiv (1/\sigma_{\varepsilon,2}^2(i))$, and $\tau_3(i) \equiv (1/\sigma_{\varepsilon,3}^2(i))$, where the i indicates that the prior belief over precisions may differ across individuals. While the model presented here allows for heterogeneous priors about precisions, we show in Appendix D that the main predictions are the same even with a common prior belief about the precisions – moreover, the Appendix model can match additional features of the experimental data. Here we present the version of the model with heterogeneous beliefs just because it can provide the basic intuitions with less notation.

The timing of the model is as follows:

1. Individuals can acquire access to one of three information sources at cost c, or choose no information ("information acquisition").

2. The selected information is displayed and individuals choose the amount of attention to allocate to the displayed information ("information processing").

3. Individuals report posterior beliefs over the fundamental and receive a payoff.

If an individual acquires access to an information source, the information is displayed. Paying attention to this information is modeled as a noisy signal about the displayed information

$$s(i) = x_j + \psi(i),$$

where j is the information source that the individual selected, x_j is the displayed information, and $\psi(i)$ is noise that arises due to limited attention to the displayed information. That is, limited attention creates a noisy perception of the displayed information. The noise $\psi(i)$ is assumed to be normally distributed with mean zero and variance $\sigma_\psi^2(i)$. Paying more attention to the displayed information is formalized as a smaller variance of noise, $\sigma_\psi^2(i)$. Individuals choose the amount of attention allocated to the displayed information, i.e., they choose $\sigma_\psi^2(i)$.

Posterior beliefs follow from Bayesian updating. If individual i selected information source $j \in \{1, 2, 3\}$ and chose the variance of noise $\sigma_\psi^2(i)$, her posterior belief is given by combining her prior belief with the signal

$$s(i) = \theta + \varepsilon_j + \psi(i).$$

The posterior mean of the fundamental is

$$E[\theta|s(i)] = \mu_\theta(i) + \frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i) + \sigma_{\varepsilon,j}^2(i) + \sigma_\psi^2(i)} [\theta + \varepsilon_j + \psi(i) - \mu_\theta(i)].$$

The weight on the displayed information is an increasing function of the perceived precision of the selected information source and the attention allocated to the displayed information.[41] The posterior variance of the fundamental is

$$\sigma_{\theta|s}^2(i) = \frac{1}{\frac{1}{\sigma_\theta^2(i)} + \frac{1}{\sigma_{\varepsilon,j}^2(i) + \sigma_\psi^2(i)}}.$$

The payoff received by individual i at the end equals

$$-\phi(\theta - E[\theta|s(i)])^2,$$

where the parameter ϕ controls the incentive to have an accurate posterior.

The optimal information strategy of an individual consists of an acquisition strategy ($j \in \{1, 2, 3\}$ or no information) and an attention strategy ($\sigma_\psi^{-2}(i) \geq 0$) that maximize the expected payoff net of the costs of acquiring and processing information:

$$-\phi\sigma_{\theta|s}^2(i) - cz - d\left(\sigma_\psi^{-2}(i)\right),$$

where c is the cost of acquiring information, z is an indicator variable which takes the value one if information is acquired, and $d\left(\sigma_\psi^{-2}(i)\right)$ denotes the cost of paying attention to the displayed information. One can think of the cost c as the sticky information aspect of the model (Mankiw and Reis, 2002), because in micro-founded versions of sticky information models there is a fixed cost of acquiring information (Reis, 2006).

Following Sims (2003), the cost of paying attention to displayed information is assumed to be an increasing function f of the reduction in uncertainty about the displayed information:

$$d\left(\sigma_\psi^{-2}(i)\right) = f(H(x_j) - H(x_j|s(i))),$$

where $H(x_j)$ denotes the entropy of the displayed information and $H(x_j|s(i))$ denotes the condi-

[41]In Gabaix's (2014) model of sparsity, the weight on information is also an increasing function of attention to the information, as in Sims' (2003) model of rational inattention. One difference between these theories is that in Sims' (2003) model of rational inattention, the weight on information can be viewed as the optimal response to a noisy perception of the information. The noisy perception of information in turn helps to match heterogeneity in reported posterior beliefs among individuals who see the same information.

tional entropy of the displayed information given the signal on the displayed information. Since entropy is simply a measure of uncertainty, the argument of the function f measures the uncertainty reduction about the displayed information due to the signal on the displayed information. The entropy of a normally distributed random variable x with variance σ_x^2 equals a constant plus $\frac{1}{2}\ln\left(\sigma_x^2\right)$, and hence, the last equation reduces to

$$d\left(\sigma_\psi^{-2}(i)\right) = f\left(\frac{1}{2}\ln\left(\frac{\sigma_{x,j}^2(i)}{\sigma_{x,j|s}^2(i)}\right)\right) = f\left(\frac{1}{2}\ln\left(\frac{\sigma_{x,j}^2(i)}{\frac{\sigma_{x,j}^2(i)\sigma_\psi^2(i)}{\sigma_{x,j}^2(i)+\sigma_\psi^2(i)}}\right)\right) = f\left(\frac{1}{2}\ln\left(1+\frac{\sigma_{x,j}^2(i)}{\sigma_\psi^2(i)}\right)\right).$$

The cost of paying attention to displayed information is an increasing function of the signal-to-noise ratio in the signal about the displayed information. In the rational inattention literature following Sims (2003), it is quite common to assume that f is a linear function, in which case the last equation reduces to

$$d\left(\sigma_\psi^{-2}(i)\right) = \mu\frac{1}{2}\ln\left(1+\frac{\sigma_{x,j}^2(i)}{\sigma_\psi^2(i)}\right),$$

where $\mu > 0$ denotes the marginal cost of attention. All qualitative results presented in the following subsection also hold for any function f that is strictly increasing, convex, twice continuously differentiable, and has a non-zero derivative at zero.

4.2 Solution of the Model and Comparison to Experimental Findings

The following proposition characterizes the solution to the model.

Proposition 1 *Individual i's optimal information strategy is to select the information source of the highest perceived precision,*

$$j^*(i) = \underset{j\in\{1,2,3\}}{\arg\max}\left(\frac{1}{\sigma_{\varepsilon,j}^2(i)}\right),$$

and to choose the precision of the signal on the displayed information that maximizes the net benefit of paying attention to the displayed information. The optimal signal precision equals

$$\sigma_\psi^{-2}(i) = \max\left\{0, \left[\frac{2\phi\sigma_\theta^2(i)}{\mu}\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)} - 1\right]\frac{1}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}\right\}. \tag{4}$$

The resulting weight on the displayed information is given by

$$E\left[\theta|s(i)\right] = \mu_\theta(i) + \alpha(i)\left[\theta+\varepsilon_{j^*}+\psi(i)-\mu_\theta(i)\right],$$

where

$$\alpha(i) = \max\left\{0, \frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)} - \frac{\mu}{2\phi\sigma_\theta^2(i)}\right\}. \tag{5}$$

The willingness to pay for access to the selected information source equals

$$WTP(i) = \max\left\{0, \phi\sigma_\theta^2(i)\left(\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)} - \frac{\mu}{2\phi\sigma_\theta^2(i)}\right) - \frac{\mu}{2}\ln\left(\frac{2\phi\sigma_\theta^2(i)}{\mu}\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}\right)\right\}. \tag{6}$$

The index i on the left-hand side of equations (4)-(6) indicates that the optimal signal precision, the weight on the displayed information, and the willingness to pay differ across individuals. When $\mu = 0$ the weight on the displayed information equals the perfect attention weight of $\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}$. When the marginal cost of attention is strictly positive, the optimal amount of attention and the weight on the displayed information are increasing in $\frac{\phi\sigma_\theta^2(i)}{\mu}$, holding constant the perfect attention weight $\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}$. The max operator in equations (4)-(6) indicates that for $\frac{\mu}{2\phi\sigma_\theta^2(i)} \geq \frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}$, individual i finds it optimal to pay no attention to the displayed information.

The proof of Proposition 1 is in Appendix C. Before turning to a comparison of model predictions and experimental findings, we make three assumptions about heterogeneity.

Assumption 1: There is cross-sectional heterogeneity in $\arg\max_{j\in\{1,2,3\}}\left(\frac{1}{\sigma_{\varepsilon,j}^2(i)}\right)$.

Assumption 2: There is cross-sectional heterogeneity in the marginal cost of attention μ, and the marginal cost of attention μ is negatively correlated with numeracy in the cross section. This assumption seems natural since individuals with higher numeracy presumably find it less costly to pay attention to quantitative information; all displayed information in the experiment is quantitative.

Assumption 3: There is cross-sectional heterogeneity in the taste for information ϕ, and the precision of the prior $\sigma_\theta^{-2}(i)$ is positively correlated with the taste for information ϕ in the cross section. This assumption also seems natural because experimental subjects had the possibility to acquire and process information already before the experiment.

Finally, we compare model predictions and experimental findings. First, due to the possibility of a corner solution in the allocation of attention, the model can match the experimental finding that some individuals choose to acquire no information even before learning the cost of information acquisition. There is no point in acquiring information that one will not pay attention to anyway. Thus, under Assumption 2, the model can also match the experimental finding that individuals with lower numeracy select the alternative "I would not like to see any information" more frequently.

Second, with heterogeneous priors over precisions (Assumption 1), the model can match the experimental finding that individuals select different information sources. Furthermore, the fact that average willingness to pay and average weight on the displayed information do not differ across the group of individuals who select the expert forecast and the group of individuals who select the last-one-year home price change suggests that $\max_{j\in\{1,2,3\}}\left(\frac{1}{\sigma_{\varepsilon,j}^2(i)}\right)$ does not differ systematically across these two groups. That is, individuals rank information sources differently but think equally highly of their preferred information source.

Third, the model can match the experimental finding that individuals in the higher-reward treatment have a higher willingness to pay and thus are more likely to acquire the information. Furthermore, under Assumptions 2-3 the model can match the finding that willingness to pay is strictly increasing in numeracy and in the precision of the prior.

Fourth, with costly attention ($\mu > 0$), the model can match the experimental finding that numeracy and reward matter *conditional* on information being displayed. Under Assumption 2 the model can match the experimental finding that individuals with higher numeracy react more to the displayed information. Under Assumption 3 the model can match the experimental finding that individuals with higher precision of the prior spend more time on processing the displayed information and react more to the displayed information. Bayesian updating implies that agents with higher precision of the prior put a smaller weight on the signal. Endogenous attention implies that agents with a stronger taste for information pay more attention to the displayed information, implying a larger weight on the signal. Hence, if the precision of the prior and ϕ are positively correlated in the cross section, individuals with higher precision of the prior may end up putting a larger weight on the signal. Furthermore, the model can match the experimental finding that individuals in the higher-reward treatment spend more time on processing the displayed information.

Fifth, Bayesian updating under rational information acquisition and rational inattention implies that individuals with strictly positive willingness to pay put a strictly positive weight on the signal. A strictly positive willingness to pay implies that the preferred information source is perceived to be useful and the individual plans to pay attention to the displayed information, which in turn implies that the weight on the signal is strictly positive. Hence, the posterior beliefs of individuals who acquire access to different information sources should diverge, which matches the experimental finding that the posterior beliefs of individuals who select different information sources move towards different signals. Finally, within a group, access to the information source decreases dispersion in beliefs if and only if the fact that individuals put weight on the same displayed information dominates the fact that there is individual-specific noise in the signal on the displayed information and the fact that individuals may have heterogeneous signal weights. Thus, the model can also match the experimental finding that dispersion in beliefs falls within some groups and increases within other groups once information is displayed.[42]

Table 6 summarizes the predictions of the following versions of the model: (i) common prior over precisions and costless attention, (ii) heterogeneous priors over precisions and costless attention, (iii) heterogeneous priors over precisions and costly attention. These predictions hold for any level of the cost of information acquisition, $c \geq 0$. In sum, the special case of the model with a common

[42]In the literature on noisy rational expectations models of financial markets following Grossman and Stiglitz (1980), there also exist models in which heterogeneity in beliefs arises because agents select different pieces of information. However, in that literature, agents select different pieces of information because equilibrium prices partially reveal information, while here agents select different pieces of information because of heterogeneity in beliefs over precisions.

prior over precisions and costless attention can match only a small subset of the findings. The benchmark version of the model with heterogeneous priors over precisions and costly attention can match almost all of the findings.

In Appendix D, we solve a model where we replace the assumption of heterogeneous priors over precisions by the assumption of a common prior over precisions and the assumption that individuals can process information about the quality of information sources before selecting an information source. Heterogeneity in beliefs over precisions arises ex post due to idiosyncratic information-processing mistakes. In that model, additional features of the experimental data arise endogenously. In particular, the highest-precision information source is endogenously the modal choice, and under Assumption 2, high numeracy individuals are endogenously more likely to select the highest-precision information source.

5 Discussion and Conclusion

Using an innovative experimental setup that makes the information acquisition process endogenous, this paper attempts to understand the role of information frictions in explaining the heterogeneity in consumers' expectations. Our survey respondents exhibit substantial demand for information and, consistent with rational inattention, the demand for information increases with incentives for accuracy. Although information acquisition costs do seem to matter, our findings indicate that the main drivers of heterogeneity in consumer expectations are constraints on information processing. Our respondents disagree on what information is most informative, with less sophisticated ones less likely to choose "informative" signals. Importantly, we see that the cross-sectional variance of the expectations distribution does not decrease when information is more cheaply available, due to the disagreement in the choice of information sources.

Disagreement in (inflation) expectations has been shown to vary over time, and the levels tend to be larger among consumers than among experts. Mankiw et al. (2003) try to explain these findings through the lens of a sticky-information model, where some people form expectations based on outdated information. While their model can fit the survey data better than models of full-information rational expectations, it is unable to match some features of the data, such as the positive relationship between the level of inflation and disagreement in expectations, or the higher level of disagreement during recessions. Our findings suggest an alternative potential explanation for these patterns: consumers all update at regular frequencies but simply look at different information. Future work that tries to understand the dynamics of information acquisition and expectation heterogeneity might benefit from collecting high-frequency data on the information sources that consumers are paying attention to.

Our results also underscore the crucial role of numeracy in the formation of expectations. Numeracy and cognitive abilities have been found to matter for the accuracy of (inflation) expectations and for the link between expectations and behavior (Armantier et al., 2015; D'Acunto et al., 2019). As to why that happens has been less clear. We are able to provide direct insights into these relationships, and find that numeracy matters at all stages of belief formation: (1) whether to consider acquiring information at all; (2) what information to acquire; (3) the valuation of the information; and (4) the weight that is put on the acquired signal. Under the plausible assumption that the cost of attention is negatively correlated with numeracy in the cross section, the theoretical model we outline in Section 4 can rationalize these empirical findings.

Our results suggest that heterogeneity in the "taste" for accurate expectations – perhaps driven by differences in current or future expected benefits from accuracy – is another important factor in the formation of expectations. We find that people who start out with more precise beliefs report a higher willingness to pay for information, spend more time on processing information, and react more to the displayed information. Future work that tries to understand the determinants of this heterogeneity in taste for accuracy would be useful for understanding consumer expectations.

Besides their implications for modeling expectation formation, our findings have some direct policy implications. There is a debate in the literature about the optimal level of information disclosure by government agencies such as central banks and statistics agencies. For instance, most government agencies have the choice of releasing data such as official statistics on inflation, unemployment, and gross domestic product, among others. Some authors have argued that information disclosure is optimal (Hellwig, 2005), whereas others argue that it can be harmful (Morris and Shin, 2002). These models always assume that individuals process all the available information optimally. Our evidence indicates that this assumption may need to be revisited. Instead, our findings imply that it is especially important for the government (and, for that matter, non-government) entities to disclose the information in a careful manner. Policy makers may want to act paternalistically by either disclosing only the most relevant information, or by guiding consumers on how to interpret and weigh all the available information.

References

[1] Allcott, H. (2011). "Social Norms and Energy Conservation." Journal of Public Economics, Vol. 95(9-10), pp. 1082-1095.

[2] Allcott, H., and Kessler, J.B. (2015). "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," NBER Working Paper No. 21671.

[3] Ambuehl, S., Ockenfels, A., and Stewart, C. (2018). "Attention and Selection Effects," Working paper.

[4] Andersen, S., Harrison, G.W., Lau, M.I., and Rutström, E.E. (2006). "Elicitation Using Multiple Price List Formats". Experimental Economics, Vol. 9(4), pp. 383-405.

[5] Angulo, A.M., Gil, J.M., and Tamburo, L. (2005). "Food Safety and Consumers' Willingness to Pay for Labelled Beef in Spain". Journal of Food Products Marketing, Vol. 11(3), pp. 89-105.

[6] Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W., and Zafar, B. (2015). "Inflation Expectations and Behavior: Do Survey Respondents Act on their Beliefs?" International Economic Review, Vol. 56(10), pp. 505-536.

[7] Armantier, O., Nelson, S., Topa, G., van der Klaauw, W. and Zafar, B. (2016). "The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment," Review of Economics and Statistics, Vol. 98(3), pp. 503-523.

[8] Armona, L., Fuster, A., and Zafar, B. (2017). "Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment," Review of Economic Studies, forthcoming.

[9] Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2017). "Social Networks and Housing Markets," NBER Working Paper No. 22258.

[10] Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2018). "The Economic Effects of Social Networks: Evidence from the Housing Market," Journal of Political Economy, forthcoming.

[11] Bartoš, V., Bauer, M., Chytilová, J., and Matějka, F. (2016). "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition," American Economic Review, Vol. 106(6), pp. 1437-1475.

[12] Burke, M., and Manz, M. (2014). "Economic Literacy and Inflation Expectations: Evidence from an Economic Experiment," Journal of Money, Credit, and Banking, Vol. 46(7), pp. 1421-1456.

[13] Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). "Diagnostic Expectations and Credit Cycles," Journal of Finance, Vol. 73(1), pp. 199-227.

[14] Carroll, C. (2003). "Macroeconomic Expectations of Households and Professional Forecasters," Quarterly Journal of Economics, Vol. 118(1), pp. 269-298.

[15] Caplin, A., and Dean, M. (2015). "Revealed Preference, Rational Inattention, and Costly Information Acquisition," American Economic Review, Vol. 105(7), pp. 2183-2203.

[16] Caplin, A., Dean, M., and Martin, D. (2011). "Search and Satisficing," American Economic Review, Vol. 101(7), pp. 2899-2922.

[17] Case, K., and Shiller, R. (1989). "The Efficiency of the Market for Single-Family Homes," American Economic Review, Vol. 79(1), pp. 125–137.

[18] Cavallo, A., Cruces, G., and Perez-Truglia, R. (2016). "Learning from Potentially Biased Statistics," Brookings Papers on Economic Activity, Spring 2016, pp. 59–108.

[19] Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). "Inflation Expectations, Learning and Supermarket Prices: Evidence from Survey Experiments," American Economic Journal: Macroeconomics, Vol. 9(3), pp. 1–35.

[20] Cheng, I., and Hsiaw, A. (2017). "Distrust in Experts and the Origins of Disagreement." Working Paper.

[21] Coibion, O., and Gorodnichenko, Y. (2012). "What Can Survey Forecasts Tell Us About Informational Rigidities?" Journal of Political Economy, Vol. 120, pp. 116-159.

[22] Coibion, O., Gorodnichenko, Y., and Kumar, S. (2015). "How Do Firms Form Their Expectations? New Survey Evidence," NBER Working Paper No. 21092.

[23] Cover, T. and H. Thomas (1991). "Elements of Information Theory," Wiley-Interscience, New York, NY, USA.

[24] Cruces, G., Perez-Truglia, R., and Tetaz, M. (2013). "Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment," Journal of Public Economics, Vol. 98, pp. 100-112.

[25] Cullen, Z., and Perez-Truglia, R. (2017). "How Much Does Your Boss Make? The Effects of Salary Comparisons," Working Paper.

[26] D'Acunto, F., Hoang, D., Paloviita, M., and Weber, M. (2019). "IQ, Expectations, and Choice." NBER Working Paper 25496.

[27] Dean, M., and Neligh, N. (2018). "Experimental Tests of Rational Inattention," Working paper.

[28] Duflo, E., and Saez, E. (2003). "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," Quarterly Journal of Economics, Vol. 118(3), pp. 815-842.

[29] Fuster, A., Hebert, B., and Laibson, D. (2012). "Natural Expectations, Macroeconomic Dynamics, and Asset Pricing," NBER Macroeconomics Annual, Vol. 26(1), pp. 1-48.

[30] Gabaix, X. (2014). "A Sparsity-based Model of Bounded Rationality," Quarterly Journal of Economics, Vol. 129(4), pp. 1661-1710.

[31] Gabaix, X. (2017). "Behavioral Inattention," NBER Working Paper No. 24096.

[32] Gabaix, X., Laibson, D., Moloche, G., and Weinberg, S. (2006). "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," American Economic Review, Vol. 96(4), pp.1043-1068.

[33] Glaeser, E., and Nathanson, C. (2015). "Housing Bubbles," in Handbook of Regional and Urban Economics, Volume 5. Eds. Gilles Duranton, Vernon Henderson, and William Strange. pp. 701-751.

[34] Goffman, E. (1963). Stigma: Notes on the Management of Spoiled Identity. Englewood Cliffs, NJ: Prentice-Hall.

[35] Grossman, S.J., and Stiglitz, J.E. (1980). "On the Impossibility of Informationally Efficient Markets," American Economic Review, Vol. 70(3), pp. 393-408.

[36] Guren, A. (2016). "House Price Momentum and Strategic Complementarity," Journal of Political Economy, forthcoming

[37] Harrison, J.M., and Kreps, D.M. (1978). “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations,” Quarterly Journal of Economics, Vol. 92(2), pp. 323–336.

[38] Hellwig, C. (2005). “Heterogeneous Information and the Welfare Effects of Public Information Disclosures.” UCLA Economics Working Paper No. 283.

[39] Hoffman, M. (2016). “How is Information Valued? Evidence from Framed Field Experiments.” The Economic Journal, Vol. 126(595), pp. 1884-1911.

[40] Hong, H., and Stein, J.C. (2007). “Disagreement and the Stock Market,” Journal of Economic Perspectives, Vol. 21(2), pp. 109–128.

[41] Hoopes, J., Reck, D., and Slemrod, J. (2015). “Taxpayer Search for Information: Implications for Rational Attention,” American Economic Journal: Economic Policy, Vol. 7(3), pp. 177-208.

[42] Khattak, A.J., Yim, Y., and Prokopy, L.S. (2003). “Willingness to Pay for Travel Information,” Transportation Research Part C: Emerging Technologies, Vol. 11(2), pp. 137-159.

[43] Khaw, M.W., Stevens, L., and Woodford, M. (2017). “Discrete Adjustment to a Changing Environment: Experimental Evidence,” Journal of Monetary Economics, Vol. 91, pp. 88-103.

[44] Laster, D., Bennett, P., and Geoum, I.S. (1999). “Rational Bias in Macroeconomic Forecasts,” The Quarterly Journal of Economics, Vol. 114(1), pp. 293-318.

[45] Lipkus, I., Samsa, G., and Rimer, B. (2001). “General Performance on a Numeracy Scale among Highly Educated Samples,” Medical Decision Making, Vol. 21, pp. 37–44.

[46] Lusardi, A. (2008). “Financial Literacy: An Essential Tool for Informed Consumer Choice?” NBER Working Paper No. 14084.

[47] Maćkowiak, B., and Wiederholt, M. (2009). “Optimal Sticky Prices under Rational Inattention,” American Economic Review, Vol. 99(3), pp. 769-803.

[48] Malmendier, U., and Nagel, S. (2016). “Learning from Inflation Experiences,” Quarterly Journal of Economics, Vol. 131(1), pp. 53-87.

[49] Mankiw, N.G., and Reis, R. (2002). “Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve,” Quarterly Journal of Economics, Vol. 117(4), pp. 1295-1328.

[50] Mankiw, N.G., Reis, R., and Wolfers, J. (2003). “Disagreement About Inflation Expectations.” In NBER Macroeconomics Annual 2003, ed. by M. Gertler, and K. Rogoff.

[51] Matějka, F., and McKay, A. (2015). “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” American Economic Review, Vol. 105(1), pp. 272-298.

[52] Milgrom, P., and Shannon, C. (1994). “Monotone Comparative Statics,” Econometrica, Vol. 62(1), pp. 157-180.

[53] Morris, S., and Shin, H.S. (2002). “The Social Value of Public Information,” American Economic Review, Vol. 92 (5), pp. 1521–34.

[54] Mullainathan, S., and Shleifer, A. (2005). “The Market for News,” American Economic Review, Vol. 95(4), pp. 1031-1053.

[55] Reis, R. (2006). “Inattentive Consumers,” Journal of Monetary Economics, Vol. 53(8), pp. 1761-1800.

[56] Roth, C., and Wohlfart, J. (2018). “How Do Expectations About the Macroeconomy Affect Personal Expectations and Behavior?” Working Paper.

[57] Scheinkman, J.A., and Xiong, W. (2003). “Overconfidence and Speculative Bubbles,” Journal of Political Economy, Vol. 111(6), pp. 183–1220.

[58] Shiller, R. (2005). Irrational Exuberance. Second Edition, Princeton: Princeton University Press.

[59] Silverman, D., Slemrod, J., and Uler, N. (2014). “Distinguishing the Role of Authority ’In’ and Authority ’To’,” Journal of Public Economics, Vol. 113, pp. 32-42.

[60] Sims, C. (2003). “Implications of Rational Inattention,” Journal of Monetary Economics, Vol. 50(3), pp. 665-690.

[61] Tversky, A., and Kahneman, D. (1974). “Judgment under Uncertainty: Heuristics and Biases.” Science 185 (4157), pp. 1124–31.

[62] Veldkamp, L. (2011). Information Choice in Macroeconomics and Finance. Princeton: Princeton University Press.

[63] Wiswall, M., and Zafar, B. (2015). “Determinants of College Major Choices: Identification from an Information Experiment.” Review of Economic Studies, Vol. 82(2), pp. 791-824.

[64] Woodford, M. (2003). “Imperfect Common Knowledge and the Effects of Monetary Policy,” In Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps, edited by Aghion, P., Frydman, R., Stiglitz, J., and Woodford, M. Princeton, NJ: Princeton Univ. Press.

[65] Zwane, A., Zinman, J., Van Dusen, E., Pariente, W., Null, C., Miguel, E., Kremer, M., Karlan, D., Hornbeck, R., Gine, X., Duflo, E., Devoto, F., Crepon, B., and Banerjee, A. (2011). “Being Surveyed Can Change Later Behavior and Related Parameter Estimates,” Proceedings of the National Academy of Sciences, Vol. 108, pp. 1821–1826.

Figure 1: Comparison of House Price Expectations: Households versus Experts

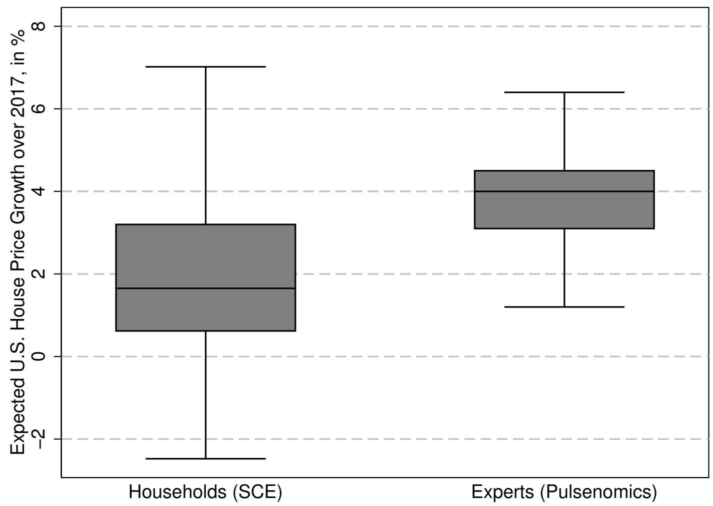

Notes: Box plots of the distribution of expected change (in percentage points) in the median house price value (Zillow Home Value Index) from December 2016 to December 2017. The left plot corresponds to the responses in the Survey of Consumer Expectations collected in February 2017 (N=1,119). The right plot corresponds to the responses of experts from the Pulsenomics panel in 2016:Q4 (N=105).

Figure 2: Prior Beliefs: Expectations about Median House Price

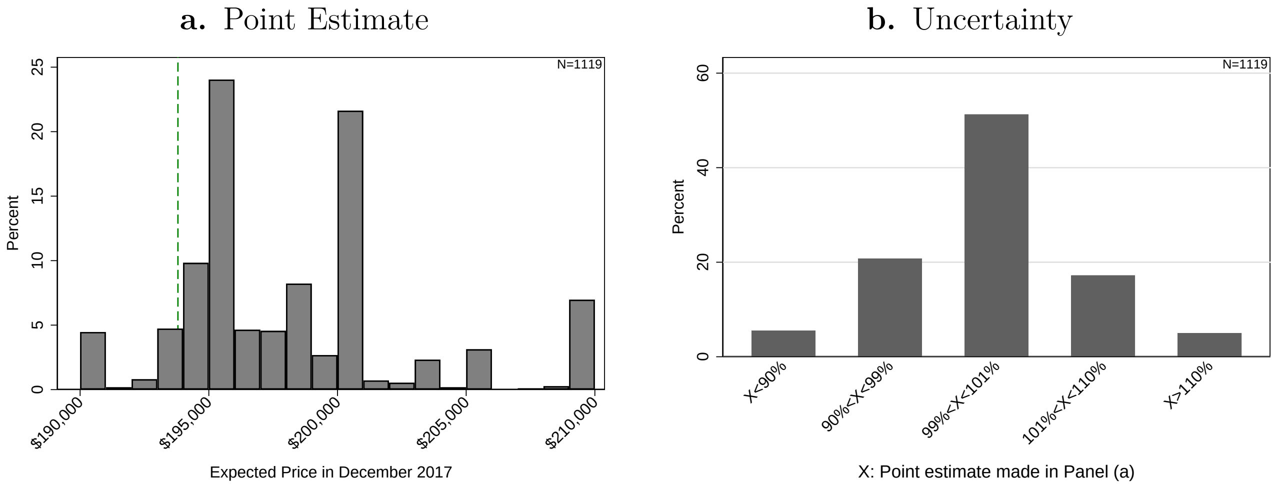

Notes: Panel (a) shows the distribution of the expected value of the typical home in the U.S. at the end of 2017 (as of February 2017, when the survey took place). The vertical dotted line corresponds to the median house value in U.S. in December 2016 according to the Zillow Home Value Index (this value was shown to respondents). The histogram is censored at $190,000 and $210,000. Panel (b) corresponds to the distribution of the confidence about the forecast made in Panel (a) by individuals.

Figure 3: Type of Information Most Preferred

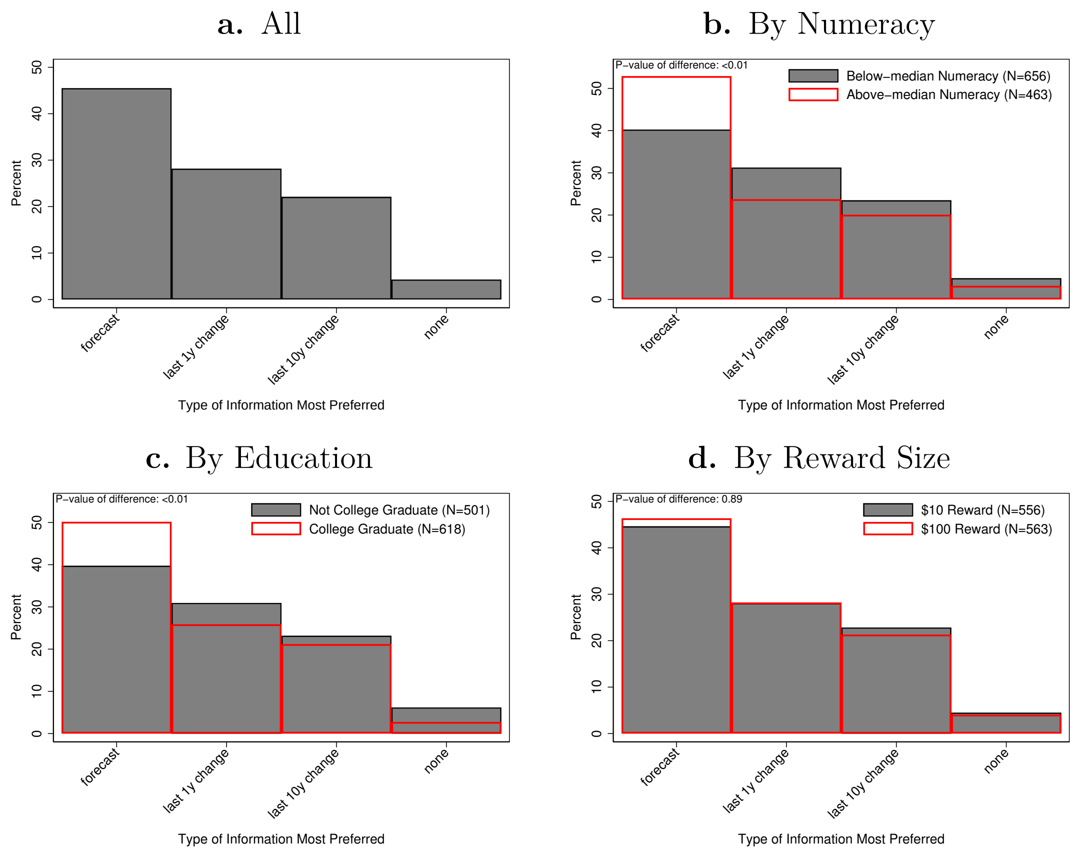

Notes: Panel (a) shows the distribution of the type of information most preferred by individuals that may help them with forecasting future year-ahead U.S. home prices; Panel (b) provides the same information according to the level of numeracy, panel (c) according to the level of education, and panel (d) according to the size of the reward. P-value of difference (in last three panels) tests for whether the distribution of most preferred information source differs by that characteristic.

Figure 4: Willingness to Pay for Favorite Information

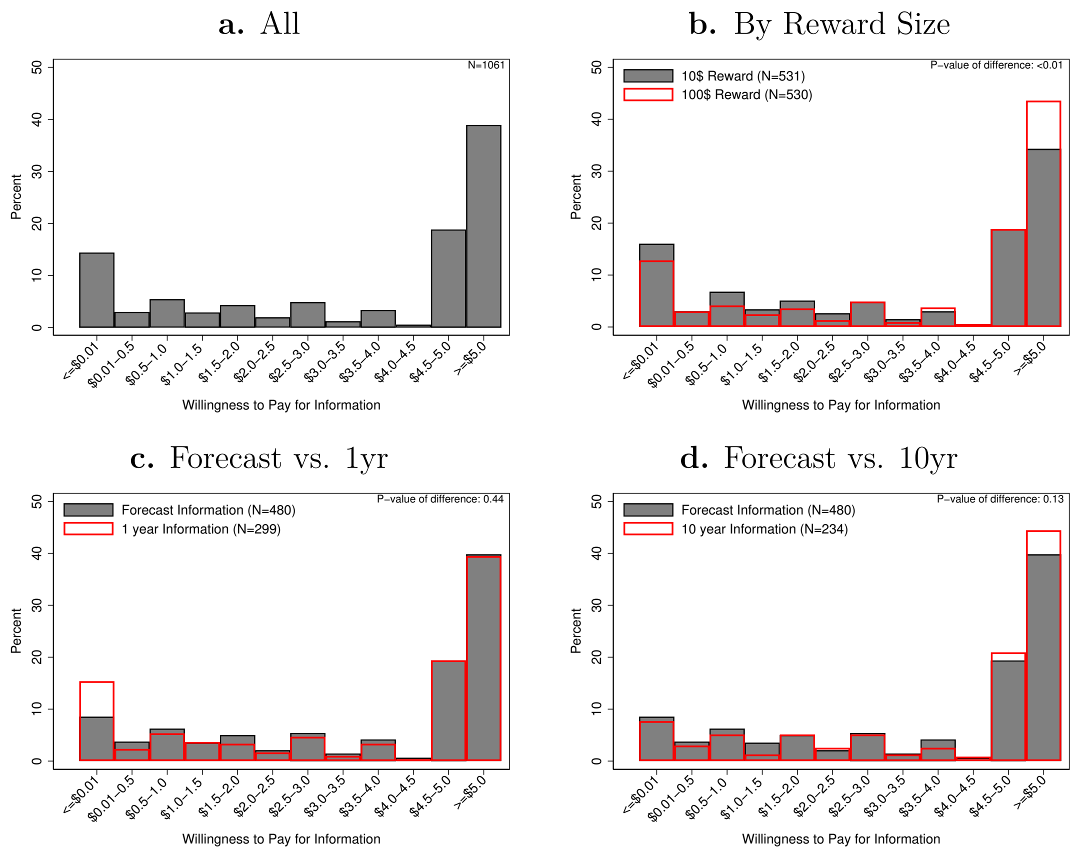

Notes: This sample does not include respondents with non-monotonic choices across the BDM scenarios. Panel (a) shows the distribution of willingness to pay (WTP) for favorite information in the whole sample. Panel (b) shows the distribution of WTP for information according the size of the reward. Panel (c) compares the distribution of WTP between individuals who preferred forecasts information and individuals who preferred information over the last one year. Panel (d) compares the distribution of WTP between individuals who preferred forecasts information and individuals who preferred information over the last ten years. P-value of difference refers to the Mann-Whitney-Wilcoxon test of the equality of two distributions.

Figure 5: Learning Rates from Information

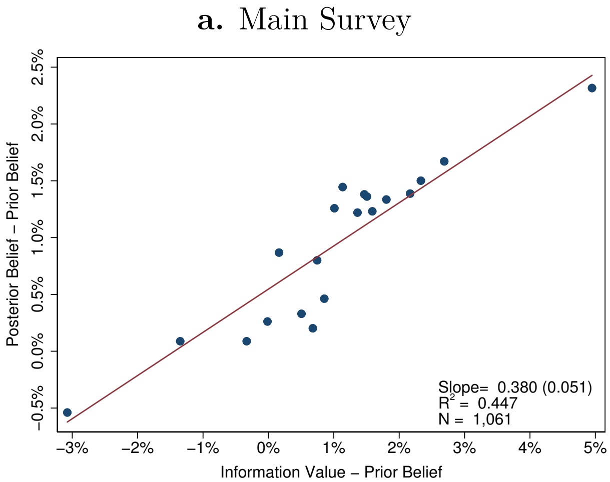

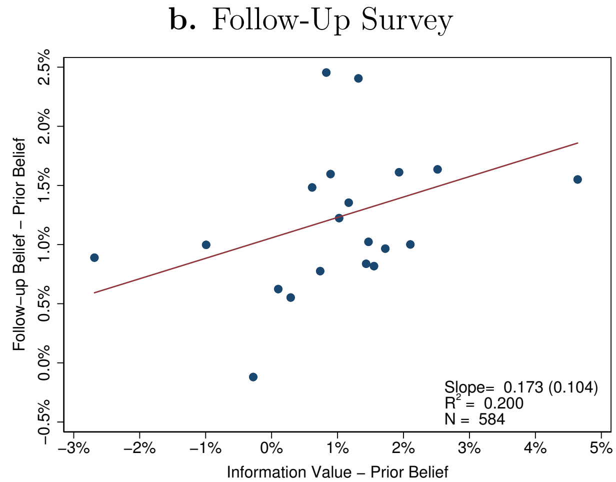

Notes: Learning rates are estimated using equation (3) from Section 3.2.2. The graphs show a binned-scatter plot based on 20 bins. Slopes, robust standard errors (in parentheses) and R^2 are based on a linear regression of the belief update (i.e., posterior belief minus the prior belief) on the signal gap (i.e., signal value minus the prior belief) interacted by a dummy that takes the value 1 if the individual was shown the signal. The regressions control for the signal gap (i.e., without the interaction), dummies for willingness to pay and the prior belief. Panel a. presents the results for the main survey (i.e., the dependent variable is the belief update during the main survey in February 2017) and panel b. presents the results for the follow-up survey (the dependent variable is the difference between the posterior belief from the follow-up survey and the prior belief from the main survey).

Figure 6: Learning Rates: Heterogeneity

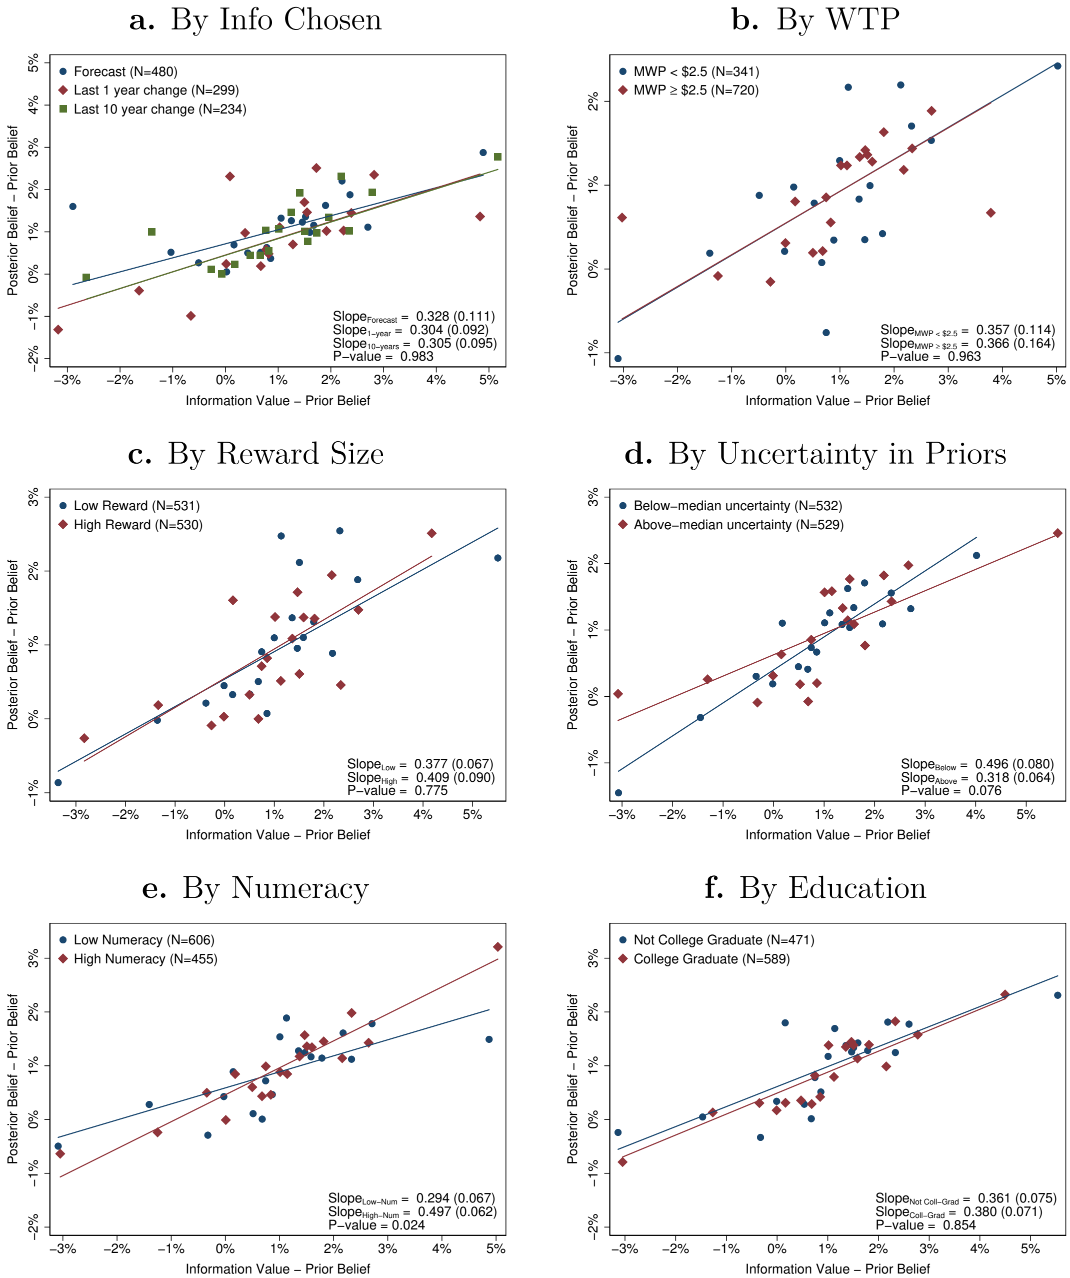

Notes: Same as in Figure 5, except that now the different panels estimate separate learning rates depending on (a) the chosen information (i.e., forecast, last 1-year change, and last 10-year change); (b) WTP (i.e., above and below the median WTP); (c) reward size (high or low reward); (d) uncertainty in prior beliefs (i.e., above and below the median uncertainty); (e) numeracy (i.e., above or below median numeracy); and (f) respondent education (at least a college degree or not).

Figure 7: Prior Beliefs: Individuals Who Will not be Shown Information vs. Individuals Who Will

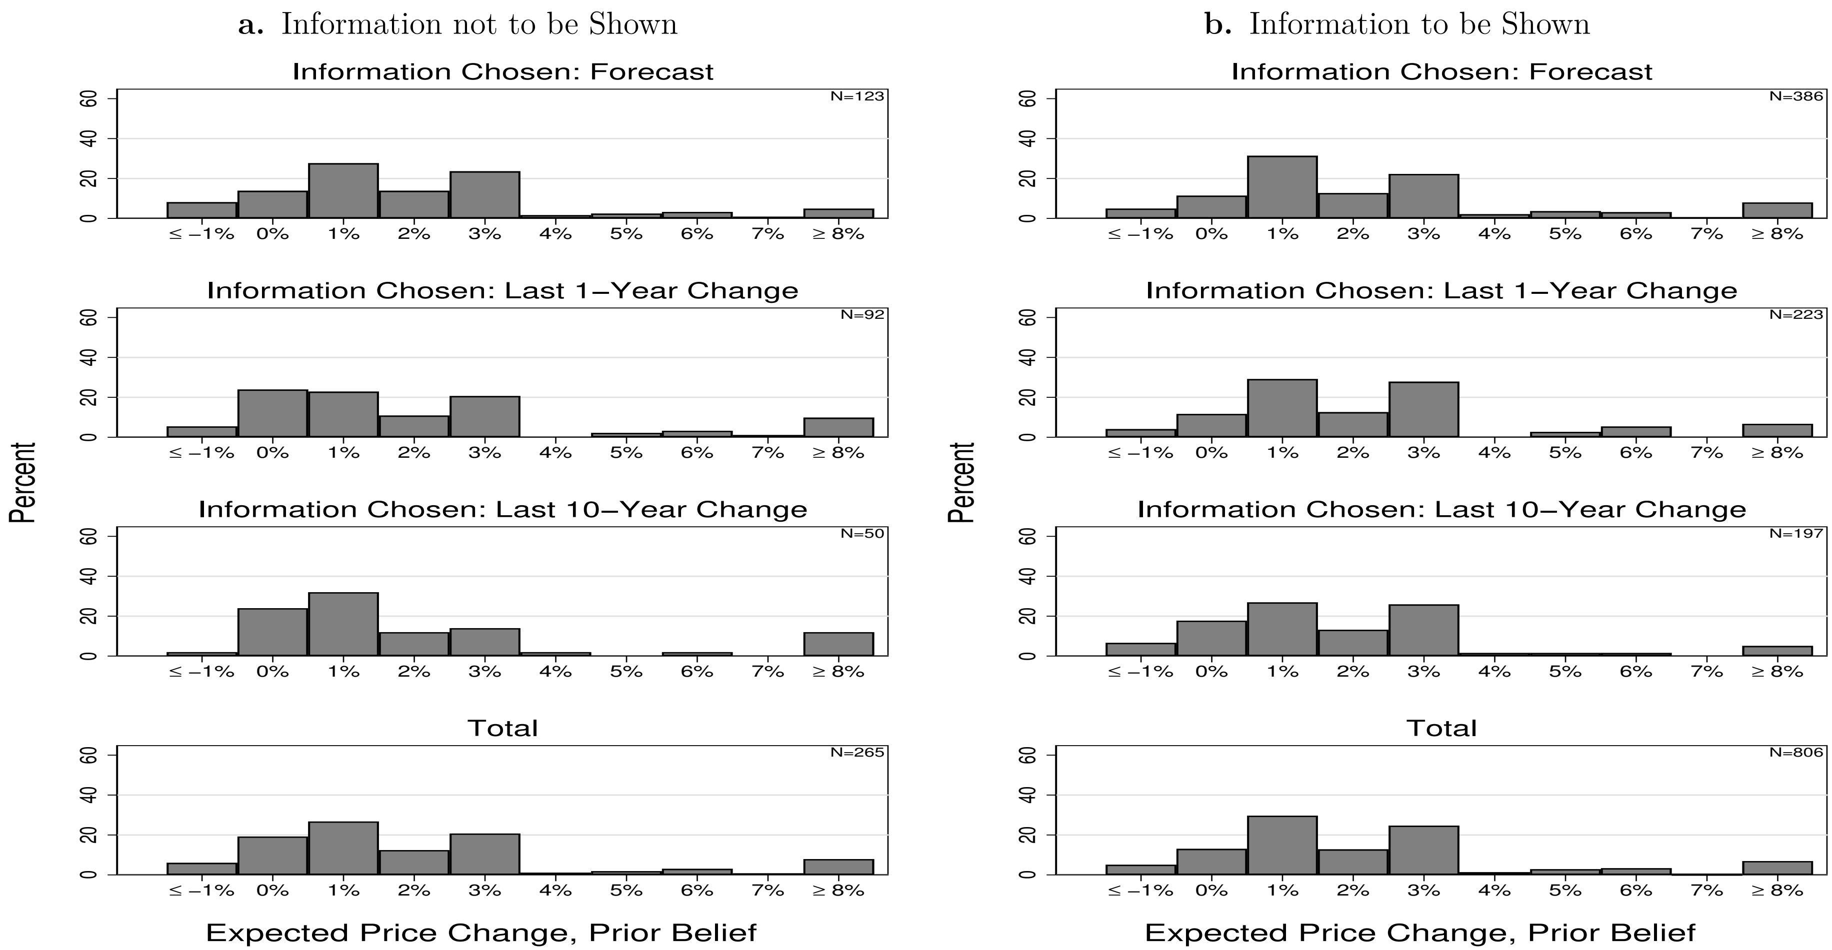

Notes: The distribution of the prior beliefs according the type of information most preferred (this sample does not include respondents who chose "None" as their most favorite information source). Panel (a) shows the distribution when individuals will not be shown information. Panel (b) shows the distribution when individuals will be shown information.

Figure 8: Posterior Beliefs: Individuals Who Were Shown Information vs. Individuals Who Were Not

a. Information Not Shown **b.** Information Shown

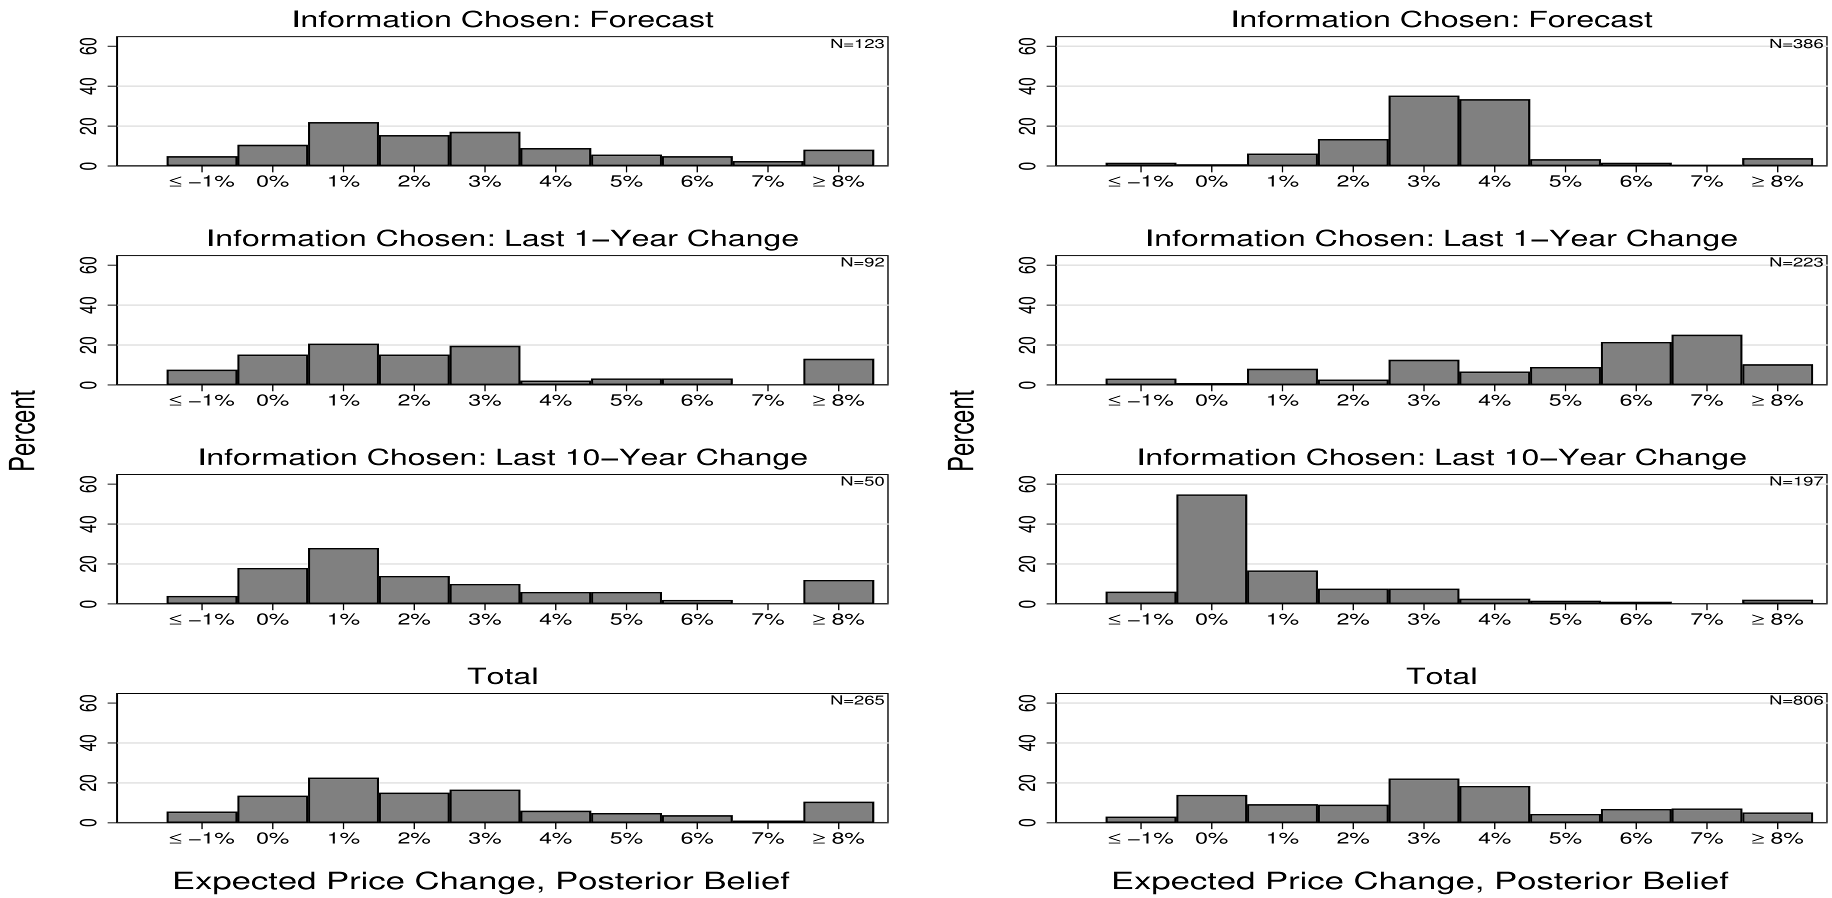

Notes: The distribution of the posterior beliefs according the type of information most preferred (this sample does not include respondents who chose “None” as their most favorite information source). Panel (a) shows the distribution of individuals who were not shown the information. Panel (b) shows the distribution of individuals who were shown the information.

Table 1: Descriptive Statistics and Randomization Balance by Reward Size

	All (1)	Low Reward (2)	High Reward (3)	F-Test P-value (4)	Low Price (5)	High Price (6)	F-Test P-value (7)
Prior Belief (1,000s)	198.1	198.2	197.9	0.374	198.1	198.2	0.662
	(5.969)	(6.095)	(5.843)		(6.022)	(5.986)	
Prior Belief (% change)	0.022	0.023	0.021	0.374	0.022	0.023	0.662
	(0.031)	(0.031)	(0.030)		(0.031)	(0.031)	
Income > 60,000 (0/1)	0.553	0.574	0.532	0.164	0.583	0.544	0.201
	(0.497)	(0.495)	(0.499)		(0.494)	(0.499)	
College Graduate (0/1)	0.552	0.550	0.554	0.877	0.577	0.543	0.275
	(0.498)	(0.498)	(0.497)		(0.495)	(0.499)	
Age	50.83	51.18	50.48	0.450	50.71	50.76	0.965
	(15.458)	(15.637)	(15.286)		(15.743)	(15.262)	
Female (0/1)	0.474	0.467	0.481	0.641	0.454	0.493	0.197
	(0.500)	(0.499)	(0.500)		(0.498)	(0.500)	
Married (0/1)	0.634	0.656	0.611	0.115	0.636	0.644	0.790
	(0.482)	(0.475)	(0.488)		(0.482)	(0.479)	
White (0/1)	0.813	0.788	0.837	0.039	0.806	0.826	0.383
	(0.390)	(0.409)	(0.370)		(0.396)	(0.379)	
Homeowner (0/1)	0.748	0.752	0.744	0.771	0.757	0.746	0.689
	(0.434)	(0.432)	(0.437)		(0.429)	(0.436)	
Numeracy (0-5)	4.013	4.005	4.020	0.824	4.069	4.000	0.278
	(1.062)	(1.096)	(1.029)		(1.034)	(1.053)	
Uncertainty in Prior Belief (Normalized)	0.041	0.044	0.038	0.065	0.042	0.041	0.881
	(0.049)	(0.051)	(0.045)		(0.048)	(0.049)	
Median House Value in State (1,000s)	225.235	226.613	223.872	0.674	233.383	218.550	0.026
	(108.080)	(107.852)	(108.384)		(114.023)	(102.627)	
House Value Volatility in State (Normalized)	0.037	0.037	0.036	0.316	0.037	0.037	0.699
	(0.015)	(0.015)	(0.015)		(0.015)	(0.015)	
Looked for Info in Past (0/1)	0.563	0.561	0.565	0.901	0.570	0.569	0.967
	(0.496)	(0.497)	(0.496)		(0.495)	(0.496)	
Conf. in Past Recall (1-5)	2.873	2.875	2.871	0.937	2.866	2.883	0.744
	(0.847)	(0.855)	(0.840)		(0.846)	(0.838)	
Probability Move and Buy in 3yr	0.200	0.196	0.205	0.572	0.204	0.204	0.967
	(0.280)	(0.275)	(0.285)		(0.278)	(0.287)	
Resp. Follow-Up Survey (0/1)	0.552	0.550	0.554	0.898	0.545	0.569	0.438
	(0.497)	(0.498)	(0.497)		(0.498)	(0.496)	
Observations	1,119	556	563		563	508	

Notes: Individual characteristics obtained from main survey. Column (1) corresponds to all respondents; columns (2) and (3) correspond to treatment groups for reward size treatment; columns (5) and (6) correspond to the price treatments (Low Price correspond to scenarios 1-4, while High-Price corresponds to scenarios 5-11). Column (4) and (7) present p-values for the test of the null hypothesis that the mean characteristic is equal to the corresponding pair of treatment groups. All variables constructed from the survey data. Uncertainty is the standard deviation derived from the individual-level subjective density (this variable is winsorized above the 98.5th percentile), normalized by the home price level at the end of 2016. House price volatility in state is the standard deviation of median home prices in the state of residence over the last 2 years, ending in December 2016, normalized by the average home price over the past 2 years.

Table 2: Factors Associated with Information Choice

	Indicator: chose...				
	Forecast	1yr	10yr	None	Forecast or 1yr
	(1)	(2)	(3)	(4)	(5)
High Reward (0/1)	0.018	0.002	-0.015	-0.004	0.019
	(0.030)	(0.027)	(0.025)	(0.012)	(0.026)
Income > 60,000 (0/1)	0.056*	-0.018	0.008	-0.045***	0.038
	(0.030)	(0.027)	(0.025)	(0.013)	(0.027)
College Graduate (0/1)	0.105***	-0.052*	-0.019	-0.034***	0.054**
	(0.030)	(0.027)	(0.025)	(0.013)	(0.027)
Age	-0.001	0.003***	-0.002**	0.000	0.002*
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
Female (0/1)	0.016	-0.013	-0.012	0.008	0.004
	(0.030)	(0.027)	(0.025)	(0.012)	(0.026)
Married (0/1)	-0.025	0.009	0.040	-0.025*	-0.016
	(0.031)	(0.028)	(0.025)	(0.013)	(0.027)
White (0/1)	0.071*	-0.032	-0.022	-0.018	0.040
	(0.038)	(0.035)	(0.033)	(0.017)	(0.035)
Numeracy (0-5)	0.063***	-0.042***	-0.001	-0.020***	0.021*
	(0.014)	(0.013)	(0.011)	(0.007)	(0.012)
Uncertainty in Prior Belief (Std)	-0.001	0.009	0.000	-0.008	0.008
	(0.015)	(0.014)	(0.012)	(0.006)	(0.013)
Median House Value in State (Std)	0.027*	-0.010	-0.008	-0.010**	0.017
	(0.015)	(0.013)	(0.012)	(0.005)	(0.013)
House Value Volatility in State (Std)	-0.001	-0.006	0.007	0.000	-0.008
	(0.015)	(0.013)	(0.013)	(0.006)	(0.014)
Looked for Info in Past (0/1)	0.009	0.024	-0.007	-0.026**	0.033
	(0.030)	(0.027)	(0.025)	(0.013)	(0.027)
Homeowner (0/1)	-0.058*	0.088***	-0.003	-0.028*	0.030
	(0.034)	(0.029)	(0.029)	(0.016)	(0.031)
Conf. in Past Recall (1-5)	-0.025	0.025	0.001	-0.001	-0.000
	(0.018)	(0.016)	(0.015)	(0.008)	(0.016)
Prob Move and Buy Home in 3 Years	0.158***	-0.094**	-0.020	-0.044**	0.064
	(0.053)	(0.047)	(0.044)	(0.018)	(0.046)
Look at Info During Survey (0/1)	0.130***	-0.113***	-0.012	-0.005	0.018
	(0.043)	(0.034)	(0.035)	(0.017)	(0.037)
Mean	0.45	0.28	0.22	0.04	0.74
Observations	1119	1119	1119	1119	1119

Notes: Heteroskedasticity-robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Each coefficient comes from a separate univariate regression. In columns (1) through (5), OLS regression are estimated using a dummy variable (=1) if the individual preferred the forecast information, 1 year information, 10 years information, None, and forecast or 1 year information as the dependent variable, respectively. Variables with a (Std) are standardized to have a mean of zero and standard deviation of 1.

Table 3: Factors Associated with WTP and Response Times

	Willingness To Pay	Log Min Posterior Belief	Log Min Posterior Belief | See Info
	(1)	(2)	(3)
High Reward (0/1)	0.828***	0.107**	0.124**
	(0.250)	(0.043)	(0.049)
Income > 60,000 (0/1)	0.862***	-0.086**	-0.176***
	(0.259)	(0.044)	(0.050)
College Graduate (0/1)	0.398	-0.002	-0.040
	(0.257)	(0.043)	(0.049)
Age	0.031***	0.007***	0.007***
	(0.008)	(0.001)	(0.002)
Female (0/1)	-0.289	0.022	0.067
	(0.254)	(0.043)	(0.049)
Married (0/1)	0.445*	-0.036	-0.118**
	(0.268)	(0.045)	(0.052)
White (0/1)	0.300	-0.045	-0.026
	(0.350)	(0.058)	(0.066)
Numeracy (0-5)	0.244*	0.029	0.009
	(0.126)	(0.021)	(0.025)
Uncertainty in Prior Belief (Std)	-0.276**	-0.078***	-0.069***
	(0.136)	(0.021)	(0.025)
Median House Value in State (Std)	0.254**	-0.018	-0.032
	(0.126)	(0.020)	(0.022)
House Value Volatility in State (Std)	0.249**	0.009	-0.012
	(0.125)	(0.022)	(0.024)
Looked for Info in Past (0/1)	0.773***	0.040	-0.025
	(0.256)	(0.043)	(0.050)
Homeowner (0/1)	0.906***	0.110**	0.055
	(0.293)	(0.051)	(0.059)
Conf. in Past Recall (1-5)	0.288*	-0.016	-0.047
	(0.154)	(0.027)	(0.033)
Prob Move and Buy Home in 3 Years	0.172	-0.003	-0.050
	(0.437)	(0.074)	(0.086)
Look at Info During Survey (0/1)	0.067	0.320***	0.263***
	(0.351)	(0.064)	(0.077)
WTP		0.104***	0.056***
		(0.010)	(0.018)
Mean	4.16	0.65	0.77
Observations	1061	1119	806

Notes: Heteroskedasticity-robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Column (1) reports results for interval regressions with willingness to pay as the dependent variable (this sample does not include respondents with non-monotonic choices across the BDM scenarios). Columns (2)-(3) report OLS regressions where the dependent variables is the log minutes spent on reporting the posterior beliefs (this variable is winsorized at the top and bottom 1.5th percentile).

Table 4: Cost of Information and Dispersion of Expectations

		Low Price (1)	**High Price** (2)	P-value Diff (3)
Obtained Signal (%)		86.19 (1.057)	65.41 (1.545)	0.00
Expectations:				
Prior	Mean	2.15 (0.133)	2.22 (0.137)	0.74
	MAD	2.06 (0.098)	2.04 (0.100)	0.88
	Uncertainty	4.16 (0.147)	4.09 (0.161)	0.83
	Disagreement (%)	10.18 (0.95)	11.87 (1.07)	0.24
Posterior	Mean	3.24 (0.141)	3.02 (0.143)	0.26
	MAD	2.21 (0.104)	2.13 (0.104)	0.59
	Uncertainty	3.03 (0.132)	3.24 (0.145)	0.45
	Disagreement (%)	20.45 (1.34)	20.39 (1.43)	0.98
Observations		536	477	

Notes: This sample does not include respondents who chose "None" as their favorite information source and respondents with non-monotonic choices across the BDM scenarios. The group Low-Price corresponds to individuals randomly assigned to scenarios 1-4 (corresponding to prices from \$0.01 to\$1.5), while the group High-Price corresponds to individuals randomly assigned to scenarios 5-11 (corresponding to prices from \$2 to \$5). The average level, the dispersion (measured as mean absolute deviation, MAD), the uncertainty, and the fraction of disagreements within group is presented for the prior and posterior belief. The prior belief refers to the expected change for year-end home prices reported in Stage 1. Posterior belief refers to the expected home price change reported in Stage 4. To measure uncertainty at the individual level, we fit these binned responses to a normal distribution for each individual (or to a uniform distribution if the respondent puts all mass in one bin or equal mass in two adjacent bins), and use the estimated standard deviation of the fitted distribution as a measure of individual-level uncertainty, with higher values denoting higher uncertainty. A disagreement is defined as non-overlap of two respondents' constructed 95% confidence interval; the table reports the fraction of all pairwise meetings that would result in a disagreement so defined. Columns (1) and (2) present the information for individuals who were randomly assigned to the Low and High Price respectively. Column (3) presents p-value for the test of the null hypothesis that the mean characteristic is equal across (1) and (2). Numbers in parentheses in each cell are standard errors.

Table 5: Effect of Information-Acquisition on the Distribution of Expectations

		Baseline Sample		Follow-Up Sample
		Prior	Posterior	Follow-Up
		(1)	(2)	(3)
Information Shown				
All	Mean	2.27 (0.106)	3.28 (0.107)	3.36 (0.191)
N=806 (450)	MAD	2.04 (0.077)	2.05 (0.078)	2.73 (0.141)
	Uncertainty	4.02 (0.117)	2.95 (0.104)	3.28 (0.152)
	Disagreem. (%)	11.59 (0.83)	20.77 (1.10)	20.28 (1.47)
Forecast	Mean	2.41 (0.164)	3.38 (0.124)	3.72 (0.282)
N=386 (205)	MAD	2.19 (0.121)	1.14 (0.109)	2.80 (0.203)
	Uncertainty	3.99 (0.167)	2.97 (0.149)	3.33 (0.211)
	Disagreem. (%)	11.46 (1.17)	7.84 (1.05)	17.56 (2.01)
1 Year Change	Mean	2.42 (0.198)	5.17 (0.209)	3.77 (0.389)
N=223 (131)	MAD	2.01 (0.145)	2.25 (0.145)	3.14 (0.275)
	Uncertainty	3.85 (0.239)	3.48 (0.234)	3.60 (0.300)
	Disagreem. (%)	15.25 (1.89)	18.33 (2.09)	22.22 (2.82)
10 Year Change	Mean	1.82 (0.179)	0.92 (0.164)	2.23 (0.317)
N=197 (114)	MAD	1.79 (0.125)	1.35 (0.132)	2.15 (0.244)
	Uncertainty	4.27 (0.226)	2.28 (0.162)	2.83 (0.322)
	Disagreem. (%)	7.98 (1.29)	12.01 (1.73)	21.30 (3.05)
Information Not Shown				
All	Mean	2.07 (0.185)	2.66 (0.225)	3.20 (0.346)
N=313 (168)	MAD	2.17 (0.139)	2.64 (0.168)	2.91 (0.263)
	Uncertainty	4.32 (0.211)	3.78 (0.205)	4.02 (0.320)
	Disagreem. (%)	9.61 (1.15)	18.64 (1.75)	17.15 (2.28)
Forecast	Mean	1.97 (0.247)	2.99 (0.311)	2.60 (0.372)
N=123 (75)	MAD	1.93 (0.175)	2.38 (0.225)	2.27 (0.263)
	Uncertainty	4.41 (0.336)	3.49 (0.295)	2.90 (0.397)
	Disagreem. (%)	9.64 (1.76)	18.26 (2.75)	20.86 (3.62)
1 Year Change	Mean	2.32 (0.403)	2.56 (0.475)	4.08 (0.812)
N=92 (45)	MAD	2.61 (0.296)	2.97 (0.358)	3.91 (0.558)
	Uncertainty	5.23 (0.424)	4.84 (0.468)	5.43 (0.713)
	Disagreem. (%)	9.65 (2.23)	18.37 (3.18)	10.40 (3.20)
10 Year Change	Mean	2.29 (0.549)	2.60 (0.484)	3.17 (0.898)
N=50 (26)	MAD	2.55 (0.411)	2.48 (0.331)	3.20 (0.630)
	Uncertainty	3.47 (0.469)	3.39 (0.419)	3.76 (0.799)
	Disagreem. (%)	8.33 (2.55)	17.88 (4.14)	13.23 (6.80)

Notes: The average level, the dispersion, the uncertainty, and the fraction of disagreements within group is presented for the prior, posterior, and follow-up belief conditional on seeing the information and the most-preferred information source. The prior belief refers to the expected change for year-end home prices reported in Stage 1. Posterior belief refers to the expected home price change reported in Stage 4. Follow-up belief refers to the expected year-end home price change reported in the follow-up survey (4 months after the main survey). See notes to Table 4 for additional details on definitions of various measures. The first number in N corresponds to the number of observations in the main survey. The number in parentheses corresponds to the number of observations in the follow-up survey. In columns (1) and (2) we present the results for the main sample. In columns (3), the sample includes individuals who were invited and responded to the follow-up survey. Numbers in parentheses in each cell are standard errors.

Table 6: Empirical Findings and Model Predictions

	All individuals choose the same information source?	Relationship between prior precision and learning rate?	Is numeracy and reward relevant? (conditionally on info displayed)
Data	No	Positive	Yes
Model			
Common prior over precisions	Yes	Negative	No
Heterogeneous priors over precisions	No	Negative	No
Heterogeneous priors over precisions & rational inattention	No	Positive or Negative	Yes

Notes: This table refers to the theory presented in Section 4.