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

We thank two anonymous referees, Rudi Bachmann, Christopher Carroll, Olivier Coibion (the editor), Theresa Kuchler, Kristoffer Nimark, Luigi Pistaferri, Dan Silverman, Michael Weber, as well as participants at numerous seminars and conferences for helpful 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, or the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Andreas Fuster, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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 in September 2020
JEL No. C81,C93,D80,D83,D84,E27,E3

ABSTRACT

Information frictions play an important role in many theories of expectation formation. We use a survey experiment to generate direct evidence on how people select, acquire and process information. Participants can buy different information signals that could help them forecast future national home prices. We elicit their willingness to pay for information, and introduce exogenous variation in the cost of information. We find that participants put substantial value on their preferred signal and, when acquired, incorporate the signal in their beliefs. However, they disagree on which signal to buy. 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 process. We provide a model with costly acquisition and processing of information, which can match most of our empirical results.

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

Ricardo Perez-Truglia
Haas School of Business
University of California, Berkeley
545 Student Services Building #1900
Berkeley, CA 94720-1900
and NBER
ricardotruglia@berkeley.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
University of Michigan
611 Tappan Street
Ann Arbor, Michigan 48109
and NBER
basitak@gmail.com

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

1 Introduction

Consumer expectations play a central role in modern macroeconomics and finance. As a result, a growing literature studies survey data on expectations. One of the robust findings in this literature is the large heterogeneity in consumer expectations (e.g., Mankiw, Reis, and Wolfers, 2003; Coibion and Gorodnichenko, 2012). This disagreement is important for a number of reasons. Heterogeneity in expectations drives trade in asset markets (e.g., Hong and Stein, 2007), explains variation in household portfolios (Meeuwis et al., 2018; Giglio et al., 2019), reduces the effectiveness of policy announcements (Angeletos and Lian, 2018), and causes inefficiencies in allocations (Angeletos and Pavan, 2009). For these reasons, the study of subjective expectations has increasingly gained attention from policymakers (e.g., Bernanke, 2007).

The literature has explored various approaches to depart from full-information rational expectations and generate such heterogeneity. In existing models, the heterogeneity in expectations emerges at three different stages of belief formation: information selection, information acquisition, and information processing. In early noisy rational expectations models of asset markets (e.g., Grossman and Stiglitz, 1980) and in sticky information models (e.g., Mankiw and Reis, 2002; Reis, 2006), heterogeneity in beliefs emerges because *only some agents acquire information* in a given period. In other noisy rational expectations models, heterogeneity in beliefs arises because *agents who acquire information select different pieces of information* (e.g., Van Nieuwerburgh and Veldkamp, 2009) or agents are simply assumed to see different pieces of information (Lucas, 1972). Finally, there is a family of models where the heterogeneity in expectations emerges at the information processing stage, because *those who acquire the same information arrive at different posteriors*: these include models with idiosyncratic noise in signals or heterogeneity in signal precision (e.g., Woodford, 2003; Sims, 2003), with heterogeneity in behavioral biases (e.g., Bordalo et al., 2019) or with heterogeneous forecasting functions (e.g., Brock and Hommes, 1997).

The goal of this paper is to use experimental methods to measure how each of these three stages of belief formation give rise to heterogeneity in expectations. To this end, we conduct an experiment in the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE). The SCE is conducted on a nationwide sample of households, broadly representative of the U.S. population. In our experiment, respondents can select, acquire, and process information. Moreover, we create exogenous variation in the information acquisition costs, which allows to measure the effects of these costs on the heterogeneity in expectations.

The experimental design is inspired by the recent literature on information-provision experiments in surveys. For example, Armantier et al. (2016), Cavallo, Cruces, and Perez-Truglia (2017), Binder and Rodrigue (2018), and Coibion, Gorodnichenko, and Weber (2019) use this method to study household inflation expectations, while Armona, Fuster, and Zafar (2019) and Roth and Wohlfart (2019) apply it to home price expectations and GDP growth expectations, respectively.[1] These studies provide a random subset of respondents with a piece of information and measure the corresponding effects on expectations. In the real world, however, individuals are rarely offered hand-picked information by academics. Instead, individuals have to find the information on their own, having to choose from multiple information sources and incurring acquisition costs. In our experiment, we study a more realistic information acquisition environment, in which agents have to acquire the information at a cost and have a choice of the information sources they want to see.

The main survey underlying our study was conducted in February 2017 as part of an annual SCE special module on the topic of housing. While the experimental design can be applied to many contexts, we provide an application to the context of home price expectations. This choice is motivated by 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; Gennaioli and Shleifer, 2018; Kaplan, Mitman, and Violante, 2019). Furthermore, individuals' home price expectations measured in surveys have been shown to be associated with real behavior such as buying or making investments in a home (Armona et al., 2019; Bailey et al., 2018). Given the prominence of housing in household portfolios, these decisions can

[1] Coibion, Gorodnichenko, and Kumar (2018) and Coibion, Gorodnichenko, and Ropele (2019) study firms' expectations (and their effect on behavior) in a similar manner. The former paper also investigates which macroeconomic variables firms track, and how firms acquire information about economic conditions.

have substantial welfare consequences.

The experimental design has four 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 respondents are randomly assigned to a reward that pays $100 with a probability of 10%. The other half is assigned to a reward that pays $10 with a probability of 10%. This variation in rewards generates exogenous variation in the benefits from having accurate expectations, which we use to test basic predictions of models of endogenous information acquisition. Before the belief re-elicitation, respondents are given the opportunity to choose among different pieces of information that could be useful for their forecasts: they can choose to see the national home price change over the past one year (+6.8% at the time of the survey), the national home price change over the past ten years (-0.1% annually), or the average expert forecast of home price growth during 2017 (+3.6%). Respondents can also state that they do not want to see any information.

In the third stage, we elicit each respondent's willingness to pay (WTP) for their most preferred piece of information. We use an incentive-compatible "multiple price list" method. We ask respondents 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. This creates exogenous variation in the cost of acquiring information. In the fourth stage, respondents get to see their preferred piece of information, if they drew a scenario in which they chose to see information. The survey concludes with the re-elicitation of home price expectations of all respondents (the "posterior belief").

Our headline empirical result is that in our setting with endogenous information acquisition, lowering the cost of information does not cause lower cross-sectional dispersion in expectations. To understand why, it is useful to go through the findings in the various stages of the experiment.

The first finding is that respondents differ in their preferred piece of information: 45.5% chose the forecast of housing experts, 28% chose the past one-year home price change, 22% chose the past ten-year home price change, and the remaining 4.5% reported to prefer no information at all. The information sources that we offered in our experiment vary markedly in terms of their informativeness. One reasonable way to rank them, although certainly not the only one, is based on their ex-ante predictive power during the years leading up to the survey. Based on this criterion, the expert forecast is the most informative, closely followed by the past one-year change, and then the ten-year change. We find that sophisticated respondents, as measured by their education or numeracy, are substantially more likely to choose the expert forecast than less sophisticated respondents. This finding suggests that at least part of the variation is due to heterogeneity in cognitive abilities that are helpful in identifying informative signals.[2]

The second finding is that respondents are willing to pay for the information and, once they acquire it, they incorporate it into their expectations. The average respondent was willing to forego $4.16 to see their preferred piece of information, which is larger than the expected accuracy reward and thus suggests that respondents want to acquire information beyond their interest in our experiment. We find strong support for a basic prediction of models of endogenous information acquisition: the average willingness to pay is about 20% higher in the $100-reward condition than in the $10-reward condition.[3] As the demand for information decreases in the cost, a lower acquisition cost results in more respondents acquiring information. We also find that respondents incorporate the information they acquire in their posterior belief – indeed, the information learned persists months after the information was acquired.

Jointly, these two findings explain why lowering the cost of information did not cause lower cross-sectional dispersion in expectations. On the one hand, more access to information tends to reduce the dispersion in posterior beliefs *within* a group that obtains the same signal. For example,

[2] It could certainly be the case that some respondents use other criteria. For instance, some respondents may distrust experts (e.g., Cavallo, Cruces, and Perez-Truglia, 2016). In a supplementary survey, we explore the role of trust in experts as a driving factor for preferences over pieces of information. We do find that less-educated respondents exhibit lower levels of trust in experts. However, these differences in trust can explain less than a quarter of the education gap in preferring experts.

[3] As a measure of effort spent on information processing and updating, we 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.

the beliefs of respondents who chose and obtained the expert forecast (a signal of 3.6%) become more compressed around 3.6%. On the other hand, more access to information tends to increase the dispersion in beliefs *across* the three groups, because each group acquires a different signal and the signals were far apart. These two opposing effects are similar in magnitude, and thus end up canceling each other out.[4]

In addition to the main result discussed above, we document a number of novel insights about the three stages of belief formation. We find that numeracy matters at several stages of belief formation. High-numeracy individuals are less likely to prefer not to see any signal (and more likely to choose the expert forecast), have a higher willingness to pay for their preferred information, and incorporate their preferred information more into their beliefs (conditional on seeing it). Prior literature has documented a positive relationship between cognitive abilities and accuracy of expectations (D'Acunto et al., 2019). Our paper provides direct evidence on why such patterns may exist. Furthermore, we 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 a significantly *higher* WTP for the information (contrary to what one might have hypothesized). With regards to updating, we find patterns that, at first sight, run counter to the basic Bayesian model: individuals who have less uncertain prior beliefs put *more* weight on the purchased information and spend more time on processing the information.

To interpret our experimental findings, we introduce a novel rational inattention model with ex-ante heterogeneous agents. The agents have a choice among different pieces of information. They face a fixed cost of acquiring information and a variable cost of processing information (i.e., paying more attention to displayed information is costlier). Agents differ in terms of the marginal cost of processing information and the incentive to have an accurate posterior. Under the assumption that this incentive is negatively correlated with prior uncertainty (i.e., respondents who enter the experiment with more precise beliefs tend to value information more), the model can match all

[4]Moreover, using a supplementary survey on a completely separate sample of respondents, we show that this result persists even when individuals are allowed to see more than one piece of information at a time.

experimental findings that at first seem puzzling: respondents with lower prior uncertainty have a higher willingness to pay for information, spend more time processing the information, and revise their beliefs more. Similarly, if low numeracy is a proxy for a high marginal cost of processing information, the model is consistent with low-numeracy respondents having lower willingness to pay for information and them putting less weight on purchased information. Finally, agents in the model rank pieces of information differently because they have heterogeneous beliefs about the precision of different pieces of information.[5]

Our paper makes three main contributions. First, it shows that disagreement about what sources of information to rely on can play an important role in explaining the dispersion in consumer expectations. Our findings imply that even if the cost of acquiring information fell to zero, we would still observe substantial dispersion in expectations. Disagreement about information sources may also explain why dispersion in expectations among consumers tends to be much larger than it is among experts (who are likely more similar in terms of information they look at), even though the estimated information acquisition costs are not larger for consumers (Coibion and Gorodnichenko, 2012). Our results further suggest that the findings from randomized information experiments can be misleading if the information acquisition process is not endogenized.[6] More broadly, our results can reconcile two important trends over the past decades: the costs of information acquisition have gone down and yet polarization of views in many areas has gone up. As costs of information acquisition decrease, a larger fraction of the population acquires information – but if agents choose to see different pieces of information, polarization of views may increase.

Second, we provide direct empirical micro evidence to inform theoretical work on inattention models. By measuring choice, valuation, and use of information, we collect data at each stage of

[5]In the simplest version of the model that illustrates the main channels, this heterogeneity in beliefs over precisions is simply taken as given. In an extension, agents have a common prior over precisions, and heterogeneous beliefs arise endogenously due to cognitive limitations at the stage where agents select pieces of information.

[6]Endogenous information acquisition has been studied in other contexts, such as hiring decisions (Bartoš et al., 2016), tax filing (Hoopes, Reck and Slemrod, 2015) and salary negotiations (Cullen and Perez-Truglia, 2017). Additionally, some laboratory experiments have been used to study demand for information in stylized settings (e.g., Gabaix et al., 2006). Burke and Manz (2014), in a laboratory experiment, investigate the demand for different types of information in the context of inflation expectations. Specifically, respondents are asked to forecast inflation and can choose up to three free information sources from a menu with seven items. Similar to us, they find an important role for economic literacy/numeracy: economically literate respondents choose better sources and use them more effectively.

the expectation formation process. We find that the first stage and the third stage are particularly important. A model that combines a choice of which information to look at and an endogenous information-processing margin is needed to match the patterns in the data. Existing models typically do not combine these two features, and hence may only offer an incomplete picture of the expectation formation process.

Third, we show that an individual's prior uncertainty is correlated with behavior at each stage of the expectation formation process. We interpret this as evidence of a selection effect: some individuals care more about having accurate beliefs on house prices. They likely acquired and processed more information on house prices before the experiment (hence, the low prior uncertainty), and they acquire and process more information during the experiment. This finding has one immediate implication: reduced-form approaches that investigate the relationship between prior uncertainty and expectation updating may not be informative tests of Bayesian updating. However, we note that existing evidence on the relationship between prior uncertainty and updating is somewhat mixed.[7] One speculative explanation, which should be examined in future work, is that prior uncertainty is more likely to be endogenous in cases where the variables over which expectations are being formed are of first-order importance to the individual (housing or labor market).

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

2 Research Design

2.1 Survey Module

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). Appendix

[7] In the context of inflation expectations, respondents with higher prior uncertainty tend to react more to the provided information (Armantier et al., 2016; Cavallo et al., 2017; Coibion, Gorodnichenko and Kumar, 2018), while in the context of house price and labor market expectations, Armona et al. (2019) and Conlon et al. (2018) do not find that to be the case. In our supplementary survey, we also replicate our finding that respondents with higher prior uncertainty update less when receiving information.

E provides screenshots of the relevant module. The survey module proceeded as follows:

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 U.S. was $193,800 as of December 2016.[8] The respondents were then 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 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.

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*

[8]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.

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."

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 multiple price list version of the Becker-DeGroot-Marschak (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.

Stage 4 - Posterior Belief: In this stage, the respondent may be shown 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.[9] 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 refer to the point estimate from this stage as the "posterior belief."

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.[10] 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

[9]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.

[10]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.

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.[11] 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.*" Respondents were also reminded about their prior belief. At the bottom of this same screen, expectations about year-end home prices were re-elicited.

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.[12] At the end of the survey, respondents were also asked whether they had used any external sources (such as Google or Zillow) when answering any question in the survey.

2.2 Information Sources

Our design tries to mimic real-world information acquisition and processing, albeit in a stylized setting. 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. This feature of the design was meant to test whether individuals have some reasonable idea or consensus about the usefulness of the information. However, no single criterion can measure

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

[12] 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).

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. [13]

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. 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, 2018; Armona et al., 2019). 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. In contrast, regressing one-year growth on growth over the previous ten years yields a small and insignificant negative coefficient.

[13]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.

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 forecast. 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.

Finally, note that our design elicits beliefs about national home prices, as opposed to local home prices. Beliefs about local prices 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: (1) 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; and (2) using national home prices simplifies the research design enormously, because we do not have to deal with local heterogeneity.[14]

2.3 Survey Implementation

The SCE Housing Survey has been fielded annually every February since 2014 and contains multiple blocks of housing-related questions, some of which distinguish between owners and renters.[15] This survey is run under the Survey of Consumer Expectations, an internet-based survey of a rotating panel of about 1,400 household heads from across the U.S. Respondents participate in the panel for up to twelve months, with a roughly equal number rotating in and out of the panel each month. Active panel members who had participated in any monthly survey in the prior eleven months were invited to participate in the housing module. Out of 1,489 household heads 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 median total survey time was 37 minutes.

Of the 1,162 valid responses, we trim the sample by dropping 43 respondents: those with prior

[14]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.

[15]See Armona et al. (2019) and https://www.newyorkfed.org/microeconomics/sce/housing#main.

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%).[16]

Four months after the initial survey, a short follow-up was fielded to active panelists in the June 2017 monthly SCE. As in Stages 1 and 4 of the main survey, respondents were asked to report their expectations about year-end U.S. home prices. We kept the identical frame of reference: 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. 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

Appendix A.1 provides additional details about the survey implementation and descriptive statistics about the respondents. This Appendix shows that, consistent with successful treatment randomization, the subjects' pre-treatment characteristics are balanced across the treatment groups: only 3 of the 34 pairwise comparisons (by either price of the information or incentive randomization) are significant at the 10% level, as would be expected by chance. Furthermore, respondents to the baseline and follow-up surveys are similar in observable characteristics, and there is no evidence that the treatment assignment affected the response rates to the follow-up survey.

3 Empirical Analysis

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

3.1 Prior Beliefs and Uncertainty

The mean (median) home price expectation in the initial stage 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, we look at the probability distribution of beliefs around the individual's point estimate. The average respondent believed that there was a 51 percent chance that the true price would fall within 1% of their guesses; however, there was high dispersion in the degree of certainty. We use the responses to the probability bins to measure prior uncertainty at the individual level: 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 (for more details, see Appendix A.2).

3.2 Ranking of Information Sources

What happens when individuals with uncertain prior beliefs are offered the chance to acquire information? Figure 1.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 the expert forecast, 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. The past predictive power criterion indicated that the expert forecast was most informative, followed by the one-year home price change and then the 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), and more than a quarter chose no information or a signal that in the past had little predictive power.

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.[17] Figure 1.b and 1.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). 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 1 further explores the heterogeneity and reports bivariate relationships between the choice of information and various individual- and location-specific characteristics. The dependent variables in columns (1)–(4) correspond to dummy variables indicating the highest ranked piece of information.[18] 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 the expert forecast), while active housing market participants are significantly more likely to choose the forward-looking signal. Individuals who report looking up information during the survey are significantly more likely to choose the expert forecast – the information source that would arguably be the hardest to find.[19] We also see that individuals who report having looked up information about home prices in the past 12 months (and hence are arguably better informed) are in fact less likely to choose no information, suggesting that such individuals have a positive taste for information. Notably, being assigned to the high-reward treatment does not have any systematic impact on the

[17]Our numeracy measure is based on five questions; see Appendix A.3 for details.

[18]The results are qualitatively similar using multivariate regressions, as reported in Appendix Table A.3. Fewer estimates are, however, statistically significant in the multivariate set-up due to the dependence between the covariates. The results are also robust if instead of a linear probability model we use a multinomial logit model.

[19]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.

ranking of information sources.

Since both the expert forecast and the past one-year home price change could be argued to be the most informative signal, column (5) studies the alternative outcome of whether a respondent ranked either the expert forecast or the one-year change 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 replicate 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.

3.3 Valuation and Use of Information

3.3.1 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.[20] Around 5% of respondents provided inconsistent responses; for example, they chose information instead of \$3 but then chose \$2.5 instead of information.[21]

Figure 2.a shows the histogram of WTP for the whole sample (excluding inconsistent respon-

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

[21]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).

dents). 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, which 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 is a fairly high WTP, given that the information we provide is public and can be found using a search tool like Google. It is also 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 are either unaware of the availability of this information or they expect a high search cost, and that they value the information beyond the context of the survey.[22] They may want to use this information for real-world housing decisions, where having incorrect expectations can translate into thousands of dollars in losses, relative to which the experimental incentive pales in comparison. Note also that the WTP is elicited after respondents had made a choice regarding an information source in the prior stage. This may have led to an "expectation-based" endowment effect and result in the WTP being higher than it would have otherwise been (e.g., Ericson and Fuster, 2011).[23]

We next study whether WTP systematically varies with reward size and other factors. To directly test the effect of stakes, Figure 2.b compares the distribution of WTP between the two reward groups. This figure shows that, as predicted by rational inattention models, 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, which can be interpreted as the mean WTP for the low-reward condition, is estimated to be \$3.75 (95% CI from 3.50 to 4.10).

[22] We can also compare the median WTP in our study with those from other studies 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).

[23] Another possibility is that the goal of winning the forecasting "contest" (although it was not explicitly framed that way) led to higher WTP. Men have been found to put a higher value on being the winner of a competitive game (D'Acunto, 2020); however, we do not find significant gender differences in WTP (see Table 2, discussed below). It could also be that the high WTP are driven by a misunderstanding of the incentives or the elicitation method, but the fact that higher numeracy respondents tend to have higher WTP (Table 2) does not support that theory.

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.

Column (1) in Table 2 shows the bivariate relationship between the WTP and various other correlates. Here, we discuss the more notable and interesting ones. First, consider the effect of past information acquisition effort on WTP. The expected relationship 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.

Looking at the relationship between the uncertainty in prior beliefs and WTP, we again find evidence for the "selection channel": individuals with a one-standard deviation higher uncertainty in their prior beliefs were, on average, willing to pay $0.28 less. 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.

Numeracy and education, which we found to be strongly associated with the ranking of information sources in the previous subsection, are positively associated with WTP, though the relationship is only significant (at $p < 0.1$) for numeracy. Finally, we study whether WTP differs depending on the information source a respondent chose. If individuals select a given information source because they (sometimes 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 (less) for information that has higher (lower) ex-ante predictive power.

The first two columns of Table A.4 show the correlates of the WTP in a multivariate framework. Generally, the patterns are qualitatively similar to those shown in the first column of Table 2.

3.3.2 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, and then use the time spent on reporting the posterior forecast as an additional 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 (see, for example, Cavallo et al., 2017). Let b^{prior} denote the mean of the prior belief, b^{signal} the signal, and $b^{posterior}$ the mean of the corresponding posterior belief. 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:

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

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. Re-arranging this expression, we get:

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

That is, the slope between the gaps ($b_i^{signal} - b_i^{prior}$) and revisions ($b_i^{posterior} - b_i^{prior}$) can be used to estimate the learning rate. However, it is possible that individuals will revise their beliefs towards the signal even if they are not provided with it.[24] To account for this and separate true learning

[24]For instance, consider someone who makes a typo when entering her prior belief and reports an estimate that differs significantly from the signal. 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.

from spurious reversion to the signal, we can exploit the random variation generated by our design in whether a respondent saw information. 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:

$$b_i^{posterior} - b_i^{prior} = \alpha \cdot \left(b_i^{signal} - b_i^{prior}\right) \cdot S_i + \beta \cdot \left(b_i^{signal} - b_i^{prior}\right) + WTP_i \delta + \varepsilon_i. \quad (1)$$

The parameter of interest is still α, which measures the true learning rate (i.e., the effect of being randomly shown information). β reflects the degree of spurious mean-reversion. Figure 3.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%. The slope of the line, $\hat{\alpha} = 0.38$, is highly statistically significant ($p < 0.001$) and economically substantial: the average individual puts 38% weight on the signal and 62% on their prior belief.[25]

We next study heterogeneity in learning rates. Figure 3.b 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 c and d of Figure 3 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

[25] Appendix A.5 presents additional details on the estimation.

signal noise). However, Figure 3.e shows that respondents with higher prior uncertainty tend to update *less*. This result is also replicated in the supplementary study described in Appendix B (see Figure B.1). 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 theory will be able to rationalize is that high-numeracy respondents tend to update more strongly than low-numeracy respondents (panel f).[26]

One potential concern with survey experiments is that, instead of inducing genuine learning, the information provided affects forecasts due to experimenter demand (Goffman, 1963) or unconscious numerical anchoring (Tversky and Kahneman, 1974). Recent survey studies find modest or no experimenter demand effects even when they are induced (see de Quidt, Haushofer, and Roth, 2018, and Mummolo and Person, 2019), and it seems highly implausible that demand effects could explain the cross-sectional patterns above. Similarly, the heterogeneity in updating based on ex-ante uncertainty and numeracy that we find is arguably inconsistent with a simple anchoring explanation, which would predict no such heterogeneity. Finally, following Cavallo et al. (2017) and Armona et al. (2019), we further use the follow-up survey to address this concern. We estimate a version of equation (1) with $b_i^{follow\text{–}up} - b_i^{prior}$ as dependent variable, where $b_i^{follow\text{–}up}$ is the belief reported four months later. The estimated slope, 0.173, is smaller than in the main survey (0.380), but remains economically meaningful and statistically significant at the 10% level. This suggests that a significant part of the reaction to the information was not spurious, since otherwise it should not persist for months after the information provision.[27]

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 2 uses the log time spent on the screen used to report the posterior belief as the dependent variable.[28] Due to the design of the

[26] In contrast, updating rates do not differ by respondent education (not shown).

[27] The lower slope in the follow-up is expected, 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.

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

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.4, and is again also robust to alternative measures of uncertainty. Higher numeracy or education are not associated with time spent on the posterior forecast. We also see that 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.

3.4 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 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 3 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 3 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 MAD 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. 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 3, 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? Figure 4 shows the distribution of beliefs for individuals who were not shown information (Figure 4.a) versus individuals who were shown information (Figure 4.b); the translucid red bars correspond to the posterior beliefs while the gray bars correspond to the prior beliefs. Comparison of the distribution of prior beliefs between the two panels shows that individuals who were and were not shown information are similar; this is as expected since whether information is shown is quasi-random. Figure 4.a shows that, among individuals who were not shown information, the distribution of posterior beliefs is the same regardless of which information source the individuals preferred.[29] In contrast, Figure 4.b shows that, for individuals 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 (as can be seen in the last row of the figure).

Table 4 provides a more quantitative version of the previous graphical argument. The first two columns of Table 4 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 3, 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. We first corroborate that, within information groups, information provision tended to reduce belief dispersion (though 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 full sample, pooling the individuals across all three information sources.

[29]However, the distribution of beliefs changed from prior to posterior even for individuals who were not shown information (p<0.01, MWW test), suggesting that the re-elicitation by itself (and the other housing-related questions they answered between the two stages) may have led respondents to think differently about future price growth.

In this pooled sample, the group that saw the information did not see a decline in the MAD of beliefs: this measure of dispersion is 2.04 percentage points for the prior beliefs and 2.05 percentage points for the posterior beliefs.[30] 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).

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 could choose between two pieces of information. Then, we randomized them into three groups: they got 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 to be either 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 decrease.

[30]The MAD 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).

4 A Model of Information Acquisition and Processing

In this section, we present a model that can match most of the experimental findings. Agents in the model have a choice among different pieces of information. They face a fixed cost of acquiring information and a variable cost of processing information.

4.1 Model Set-up

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 purchased 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.

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

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

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

where x_j is the displayed information and ε_j is normally distributed noise with mean zero. Individuals may have heterogeneous priors 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 prior beliefs just because it can provide the basic intuitions with less notation.

Paying attention to displayed information is modeled as a noisy signal about the displayed information, as in Sims (2003),

$$s\left(i\right)=x_{j}+\psi\left(i\right),$$

where $s\left(i\right)$ is the signal on the displayed information, j is the information source that the individual selected, and $\psi\left(i\right)$ 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 $\psi\left(i\right)$ is assumed to be normally distributed with mean zero and variance $\sigma_{\psi}^{2}\left(i\right)$. Paying more attention to the displayed information is formalized as a smaller variance of noise, $\sigma_{\psi}^{2}\left(i\right)$. Individuals choose the amount of attention allocated to the displayed information, i.e., they choose $\sigma_{\psi}^{2}\left(i\right)$.

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

$$s\left(i\right)=\theta+\varepsilon_{j}+\psi\left(i\right).$$

The posterior mean of the fundamental is

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

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.[31] The posterior variance of the fundamental is

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

[31]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.

The payoff received by individual i at the end equals

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

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 ($\sigma_\psi^{-2}\left(i\right) \geq 0$) that maximize the expected payoff net of the costs of acquiring and processing information:

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

where c is the cost of acquiring information, z is an indicator variable which takes the value one if information is acquired, and $d\left(\sigma_\psi^{-2}\left(i\right)\right)$ 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 the displayed information is assumed to be an increasing function f of the reduction in uncertainty about the displayed information:

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

where $H\left(x_j\right)$ denotes the entropy of the displayed information and $H\left(x_j|s\left(i\right)\right)$ denotes the conditional 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

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

In the rational inattention literature following Sims (2003), it is common to assume that f is a linear function, in which case the last equation reduces to

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

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. Note that the cost of paying attention is an increasing function of the signal-to-noise ratio in the signal about the displayed information.

4.2 Solution of the Model and Comparison to Experimental Findings

The following proposition gives the closed-form solution of the model.

Proposition 1 *Individual i considers acquiring the information source with the highest perceived precision, $j^*(i) = \underset{j\in\{1,2,3\}}{\arg\max}\,\tau_j(i)$. For any information source, the optimal amount of attention to displayed information is determined by the incentive to have an accurate posterior, ϕ, the marginal cost of attention, μ, prior uncertainty, $\sigma_\theta^2(i)$, and the perceived precision of the information source, $\tau_j(i) = \left(1/\sigma_{\varepsilon,j}^2(i)\right)$:*

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

The willingness to pay for access to the preferred information source equals the difference between benefit and cost of paying attention at the optimal amount of attention:

$$WTP(i) = \max\left\{0, \phi\sigma_\theta^2(i)\left(\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i) + \sigma_{\varepsilon,j^*}^2(i)} - \frac{\mu}{2\phi\sigma_\theta^2(i)}\right) - \frac{\mu}{2}\ln\left(\frac{2\phi\sigma_\theta^2(i)}{\mu}\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i) + \sigma_{\varepsilon,j^*}^2(i)}\right)\right\}. \tag{3}$$

If $WTP(i) \geq c$, individual i acquires access to the preferred information source, the amount of attention to the displayed information is given by equation (2) with $j = j^*(i)$, and the weight on the displayed information is given by

$$E[\theta|s(i)] = \alpha(i)(\theta + \varepsilon_{j^*} + \psi(i)) + (1 - \alpha(i))\mu_\theta(i),$$

with

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

The proof of Proposition 1 is in Appendix C. The term $\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}$ in equation (4) is the full-attention Bayesian weight on the displayed information. The variable $\alpha(i)$ is the rational-inattention Bayesian weight on the displayed information. The difference between the two is governed by the cost and benefit of paying attention, as captured by the term $\frac{\mu}{\phi\sigma_\theta^2(i)}$. The willingness to pay for access to the preferred information source equals the gain from a more accurate posterior, $\phi\sigma_\theta^2\left(i\right)\alpha\left(i\right)$, minus the cost of paying attention to the displayed information at the optimal amount of attention to the displayed information. In this simplest version of the model, the choice of the information source is simply governed by the prior beliefs over precisions.

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 $\underset{j\in\{1,2,3\}}{\arg\max}\,\tau_j\left(i\right)$.

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 more numerate individuals 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 $\sigma_\theta^{-2}\left(i\right)$ is positively correlated with the taste for information ϕ in the cross section. This assumption 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.

Ranking of the four options. The model can match the finding that some individuals choose the option "*I would not like to see any information*" before even learning the cost of information acquisition. There is simply no point in acquiring information that one will not pay attention to anyway – the corner solution in equations (2)-(4).[32] The model is also consistent with low-numeracy individuals selecting this option more frequently, because their cost of processing information is higher.[33]

[32]The existence of corner solutions in rational inattention models has been emphasized at least since Maćkowiak and Wiederholt (2009) and is due to the fact that the marginal cost of attention is non-zero at zero. Thus, the fixed cost of acquiring information c is not required for our model to predict that some agents will not acquire information; we instead add it in order to be able to study how variation in this cost will affect information choice and disagreement.

[33]Note that numeracy and prior uncertainty are significantly negatively correlated in our data (Pearson correlation coefficient of -0.21). This is in line with Assumption 2 given that individuals had the possibility to acquire and process information in their daily life, prior to entering the experimental setting.

Valuation of information. The model matches the finding that individuals in the high-reward treatment have a higher willingness to pay ($WTP\left(i\right)$ is increasing in ϕ). The model can also match the finding that willingness to pay for information is increasing in the precision of the prior (rather than decreasing in the precision of the prior, as one may have expected). The intuition is explained below. In addition, the model matches the finding that willingness to pay for information is increasing in numeracy. High-numeracy individuals anticipate that processing information is cheaper for them and that they will incorporate the displayed information more into their beliefs. Moreover, the finding that the average willingness to pay does not differ across the group of individuals who select the expert forecast and the group of individuals who select the one-year home price change can be easily interpreted in terms of the model: individuals rank information sources differently but think equally highly of their preferred information source ($arg \max_{j\in\{1,2,3\}} \tau_j\left(i\right)$ differs across these groups but $\max_{j\in\{1,2,3\}} \tau_j\left(i\right)$ does not differ systematically across these groups).

Use of information. The model predicts that individuals in the high-reward treatment exert more effort on processing information .[34] Furthermore, under Assumption 3, the model can match the finding that individuals with a higher precision of the prior put *more* weight on the information. Standard Bayesian updating implies that the weight on the information is decreasing in the precision of the prior, contrary to our experimental findings. Formally, the full-attention Bayesian weight $\frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)}$ is strictly decreasing in $\sigma_\theta^{-2}\left(i\right)$. However, the rational-inattention Bayesian weight $\alpha\left(i\right) = \frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)+\sigma_\psi^2(i)} = \frac{\sigma_\theta^2(i)}{\sigma_\theta^2(i)+\sigma_{\varepsilon,j^*}^2(i)} - \frac{\mu}{2\phi\sigma_\theta^2(i)}$ is increasing in the incentive to have an accurate posterior, ϕ. If the selection effect formalized by Assumption 3 is sufficiently strong, the rational-inattention Bayesian weight $\alpha\left(i\right)$ is increasing in the precision of the prior. In this case, the model can simultaneously match the three empirical findings that individuals with a higher precision of the prior report a higher willingness to pay for information, exert more effort on processing the displayed information, and put more weight on the displayed information. In addition, the model matches the finding that the weight on the displayed information is increasing in numeracy – see equation (4).

[34] This is consistent with column (3) of Table 2, although note also that the average weights put on the signal did not differ significantly across the high- and low-reward treatments (Figure 3.d).

Information acquisition and dispersion of expectations. The model matches the finding that in the high-cost treatment fewer individuals acquire information. The model also matches the finding that posterior beliefs diverge across the three information groups, because agents put weight on different pieces of information, and that a reduction in the cost of information can increase dispersion in beliefs because of the divergence of beliefs across information groups.[35, 36]

In Appendix D, we replace the assumption of heterogeneous priors over precisions by the assumption of a common prior over precisions but agents also facing the cost of information processing at the stage where they rank pieces of information. Heterogeneity in beliefs over precisions arises *ex post* due to idiosyncratic information-processing mistakes at the stage where agents rank pieces of information. In that model, additional features of the experimental data arise endogenously: the highest-precision information source is endogenously the modal choice, and high-numeracy individuals are endogenously more likely to select the highest-precision information source (Figure 1). This in turn has the reinforcing effect that high-numeracy individuals feel more confident about having selected the highest-precision information source and thus put even more weight on the displayed information.

5 Discussion and Conclusion

Using an innovative experimental setup measuring agents' choice, valuation and use of information, this paper attempts to understand the role of information frictions in explaining the heterogeneity in consumers' expectations. Our main empirical result is that with endogenous information

[35]Furthermore, 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, $\psi(i)$, and the fact that individuals put heterogeneous weights on the signal, $\alpha(i)$. Hence, the model can also match the experimental finding that dispersion in beliefs remains high in all groups, falls within some groups, and increases within other groups once information is displayed.

[36]Note that the model does not make general predictions for disagreement for different menus without adding additional assumptions. The effect of reducing the number of options from, say, three to two on disagreement depends on the realization of the third signal that is eliminated (for example, if an intermediate signal is eliminated, disagreement could increase) and on the perceived precision of the remaining two signals compared to the perceived precision of the eliminated signal. With two options instead of three options, the cross-sectional distribution of beliefs will tend to have two modes instead of three modes, but where these modes are located depends on the pieces of information, and disagreement also exists among those individuals who acquire the same information.

acquisition, lowering the cost of information does not cause lower cross-sectional dispersion in expectations, because respondents choose to look at different pieces of information. This result is not an artifact of our design where respondents can view at most one piece of information – we corroborate it in our supplementary survey where respondents may receive multiple signals.

Our findings have a number of implications for modeling expectation formation. First, given that dispersion in expectations is such a pervasive feature of survey data, it seems important to incorporate this dispersion in economic models, but probably the focus should shift somewhat from heterogeneity in the timing of updating to heterogeneity in the type of information that agents look at. A second important dimension of heterogeneity in posterior beliefs is heterogeneity among agents who see the same information. The fact that precision of the prior and numeracy increase the extent to which information is incorporated into posterior beliefs suggests the importance of an endogenous information-processing margin. Third, heterogeneity in the cost of processing information also seems central to understand cross-sectional differences in information acquisition.

These implications are well aligned with existing findings of numeracy and cognitive abilities being associated with the accuracy of (inflation) expectations and the link between expectations and behavior (Armantier et al., 2015; D'Acunto et al., 2019). We are able to provide direct insights into these relationships, and find that numeracy matters at all stages of belief formation: (1) whether to acquire 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 low-numeracy individuals have a higher cost of processing information, the theoretical model we outline in Section 4 can rationalize these empirical findings.

Future work that tries to understand the dynamics of information acquisition and expectation heterogeneity would benefit from high-frequency data on the information sources that consumers are paying attention to. It would also be important to shed further light on *why* different consumers choose to look at different information sources. Finally, in future work, it would also be interesting to estimate the model using the experimental data.[37]

[37]Gaglianone et al. (2019) estimate a model where professional forecasters face a fixed cost of updating their forecasts and a variable cost of processing information, as in Sims (2003). They exploit a dataset of Brazilian forecasters

Besides their implications for modeling expectation formation, 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; 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 weigh and interpret the various pieces of information.

References

[1] 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.

[2] 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.

[3] Angeletos, G.-M. and Lian, C. (2018). "Forward Guidance without Common Knowledge," American Economic Review, Vol. 108(9), pp. 2477–2512.

[4] Angeletos, G.-M. and Pavan, A. (2009). "Policy with Dispersed Information," Journal of the European Economic Association, Vol. 7(1), pp. 11–60.

[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.

who update their forecasts at days of their choice and participate in a monthly forecasting contest. Their paper, like ours, finds that the cost and benefit of attention matters for the allocation of attention.

[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. (2019). "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment," Review of Economic Studies, Vol. 86(4), pp. 1371-1410.

[9] 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, Vol. 126(6), pp. 2224-2276.

[10] 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.

[11] Bernanke, B.S. (2007). "Inflation Expectations and Inflation Forecasting," Speech at the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, MA.

[12] Binder, C., and Rodrigue, A. (2018). "Household Informedness and Long-Run Inflation Expectations: Experimental Evidence," Southern Economic Journal, Vol. 65(5), pp. 1059-1095.

[13] Brock, W.A., and Hommes, C.H. (1997). "A Rational Route to Randomness," Econometrica, Vol. 65(5), pp. 1059-1095.

[14] 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.

[15] Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2019). "Overreaction in Macroeconomic Expectations," American Economic Review, forthcoming.

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

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

[21] Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). “How Do Firms Form Their Expectations? New Survey Evidence,” American Economic Review, Vol. 108(9), pp. 2671-2713.

[22] Coibion, O., Gorodnichenko, Y., and Ropele, T. (2019). “Inflation Expectations and Firm Decisions: New Causal Evidence,” Quarterly Journal of Economics, forthcoming.

[23] Coibion, O., Gorodnichenko, Y., and Weber, M. (2019). “Monetary Policy Communications and their Effects on Household Inflation Expectations,” NBER Working Paper No. 25482.

[24] Conlon, J.J., Pilossoph, L., Wiswall, M., and Zafar, B. (2018). “Labor Market Search With Imperfect Information and Learning,” NBER Working Paper No. 24988.

[25] 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.

[26] Cullen, Z., and Perez-Truglia, R. (2017). “How Much Does Your Boss Make? The Effects of Salary Comparisons,” NBER Working Paper No. 24841

[27] D’Acunto, F. (2020). “Identity and Choice Under Risk,” Working Paper, Boston University.

[28] D’Acunto, F., Hoang, D., Paloviita, M., and Weber, M. (2019). “IQ, Expectations, and Choice,” NBER Working Paper No. 25496.

[29] de Quidt, J., Haushofer, J., and Roth, C. (2018). “Measuring and Bounding Experimenter Demand,” American Economic Review, 108 (11), 3266–3302.

[30] 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.

[31] Ericson, K.M.M., and Fuster, A. (2011). “Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments,” Quarterly Journal of Economics, Vol. 126(4), pp. 1879-1907.

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

[33] 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.

[34] Gaglianone, W.P., Giacomini, R., Issler, J.V., and Skreta, V. (2019). "Incentive-driven Inattention," CEPR Discussion Paper 13619.

[35] Gennaioli, N., and Shleifer, A. (2018). "A Crisis of Beliefs: Investor Psychology and Financial Fragility," Princeton University Press.

[36] Glaeser, E., and Nathanson, C. (2015). "Housing Bubbles," in Handbook of Regional and Urban Economics, Volume 5, edited by Duranton G., Henderson V., and Strange W., pp. 701-751.

[37] Giglio, S.; Maggiori, M.; Stroebel, J. and Utkus, S. (2019), "Five Facts about Beliefs and Portfolios," NBER Working Paper No. 25744.

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

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

[40] Guren, A. (2018). "House Price Momentum and Strategic Complementarity," Journal of Political Economy, Vol. 126(3), pp. 1172-1218.

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

[42] 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.

[43] Kaplan, G., Mitman, K. and Violante, G.L. (2019), "The Housing Boom and Bust: Model Meets Evidence," Journal of Political Economy, forthcoming.

[44] 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.

[45] Lucas, R.E. (1972). "Expectations and the Neutrality of Money," Journal of Economic Theory, Vol. 4, pp. 103-124.

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

[47] 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.

[48] 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.

[49] Meeuwis, M.; Parker, J.A.; Schoar, A. and Simester, D.I. (2018), "Belief Disagreement and Portfolio Choice," NBER Working Paper No. 25108.

[50] Mummolo, J. and Peterson, E. 2019. "Demand Effects in Survey Experiments: An Empirical Assessment," American Political Science Review, 113(2), 517-529.

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

[52] Roth, C., and Wohlfart, J. (2019). "How Do Expectations About the Macroeconomy Affect Personal Expectations and Behavior?" Review of Economics and Statistics, forthcoming.

[53] Shiller, R. (2005). Irrational Exuberance. Princeton: Princeton University Press.

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

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

[56] Van Nieuwerburgh, S., and Veldkamp, L. (2009). "Information Immobility and the Home Bias Puzzle," Journal of Finance, Vol. 64(3), pp. 1187-1215.

[57] 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.

[58] 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.

Figure 1: Type of Information Most Preferred

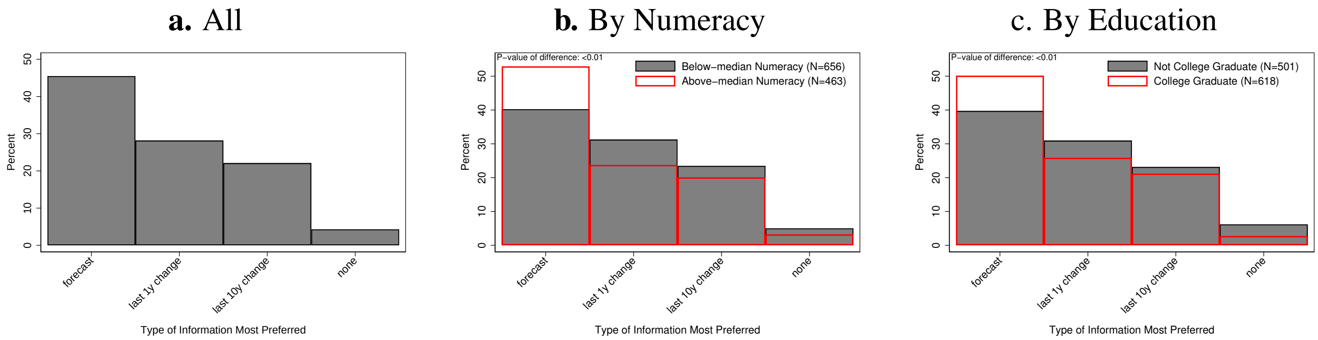

Notes: Panel (a) shows the distribution of the type of information most preferred by individuals to help them with forecasting future year-ahead U.S. home prices. Panel (b) provides the same information according to the level of numeracy, and panel (c) according to the level of education. P-value of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

Figure 2: Willingness to Pay for Favorite Information

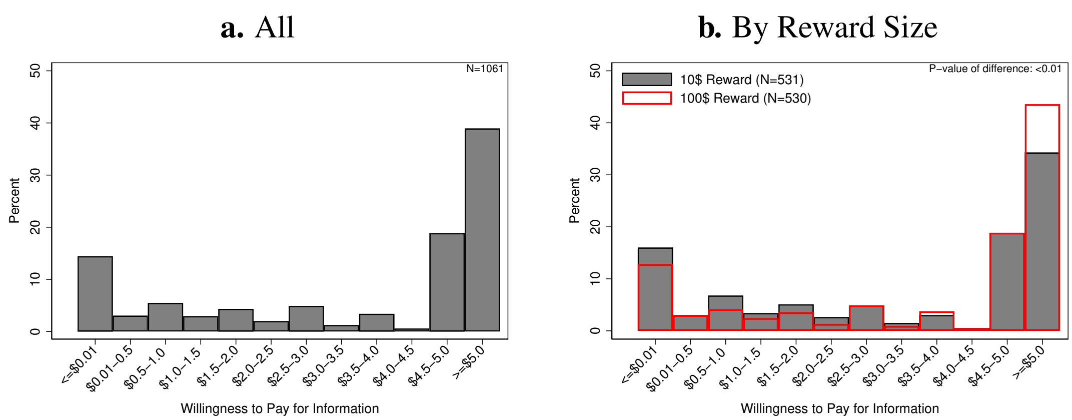

Notes: This sample does not include respondents with non-monotonic choices across the BDM scenarios. Panel (a) shows the distribution of willingness to pay (WTP) for favorite information in the whole sample. Panel (b) shows the distribution of WTP for information according the size of the reward. P-value of difference refers to the Mann-Whitney-Wilcoxon test of the equality of two distributions.

Figure 3: Learning Rates from Information

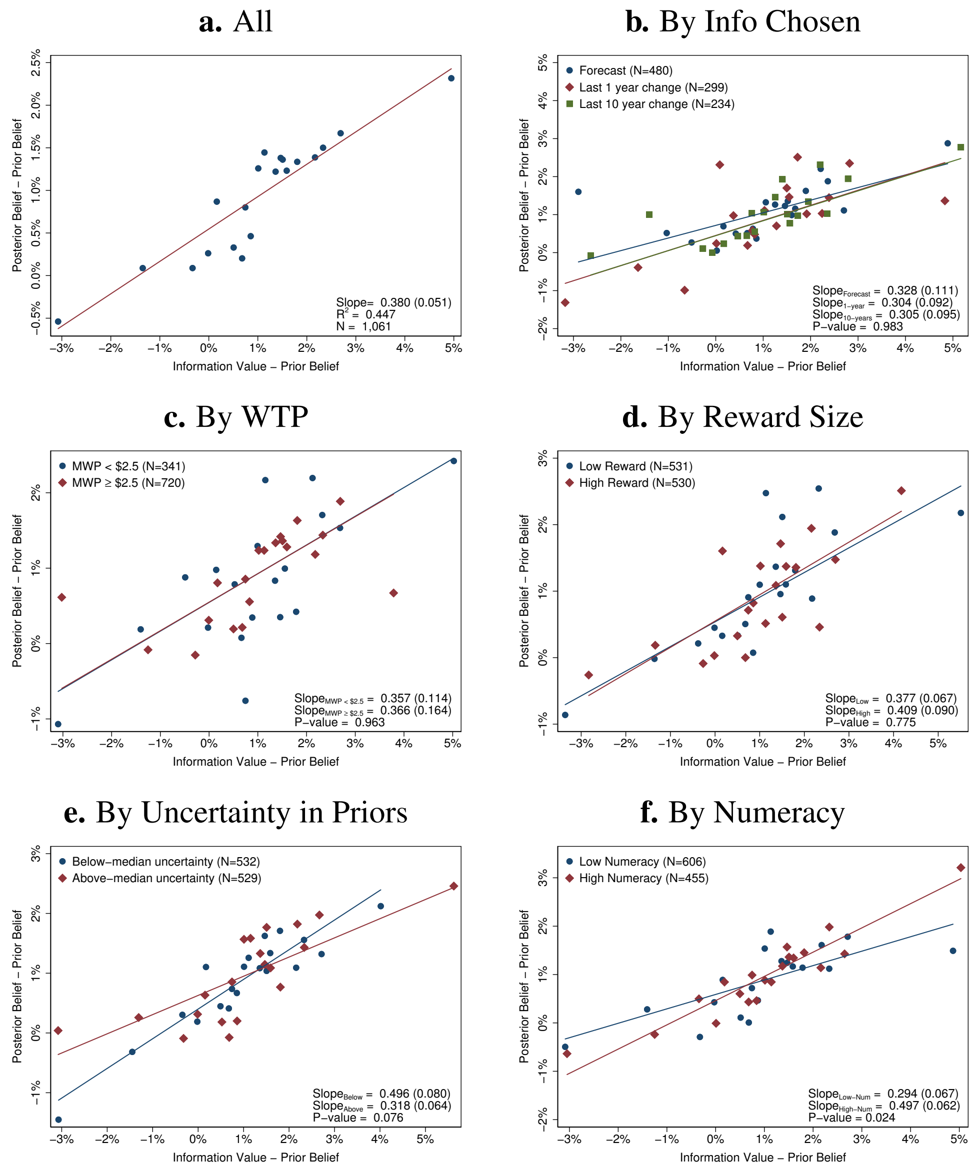

Notes: Learning rates are estimated using equation (1) from Section 3.3.2. The graphs show a binned-scatter plot based on 20 bins. Slopes, robust standard errors (in parentheses) and R^2 are based on a linear regression of the belief update (i.e., posterior belief minus the prior belief) on the signal 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 full sample of the main survey. The remaining panels estimate separate learning rates depending on (b) the chosen information (i.e., forecast, last 1-year change, and last 10-year change); (c) WTP (i.e., above and below the median WTP); (d) reward size (high or low reward); (e) uncertainty in prior beliefs (i.e., above and below the median uncertainty); and (f) numeracy (i.e., above or below median numeracy)

Figure 4: Distribution of Prior and Posterior Beliefs: Individuals Who Were Shown Information vs. Individuals Who Were Not

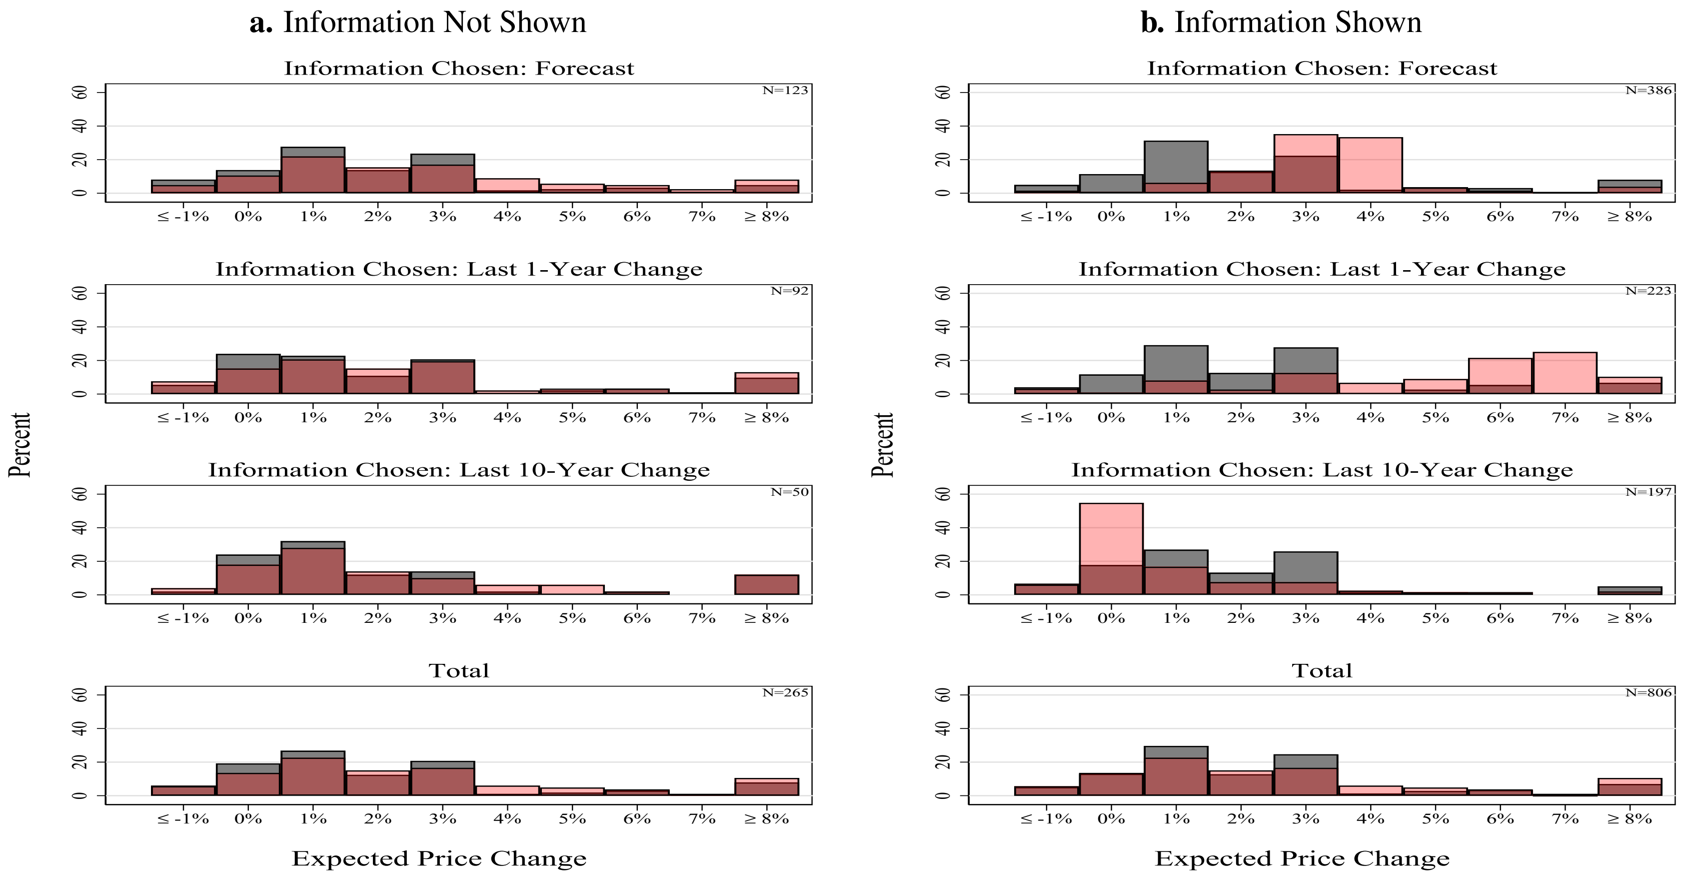

Notes: Distribution of prior and posterior beliefs according to the type of information most preferred (this sample does not include respondents who chose "None" as their most favorite information source). Panel (a) displays the distributions for individuals who were not shown the information. Panel (b) displays the distributions for individuals who were shown the information. The solid gray bars show individuals' prior beliefs, while the translucid red bars showindividuals' posterior beliefs (i.e., after they received or did not receive information).

Table 1: 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 bivariate 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 2: Factors Associated with WTP and Response Times

	Willingness To Pay (1)	Log Min Posterior Belief (2)	Log Min Posterior Belief | See Info (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. Each coefficient comes from a separate bivariate regression. 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 3: 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-values 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 4: Effect of Information-Acquisition on the Distribution of Expectations

		Prior	Posterior
		(1)	(2)
Information Shown			
All	Mean	2.27 (0.106)	3.28 (0.107)
N=806	MAD	2.04 (0.077)	2.05 (0.078)
	Uncertainty	4.02 (0.117)	2.95 (0.104)
	Disagreem. (%)	11.59 (0.83)	20.77 (1.10)
Forecast	Mean	2.41 (0.164)	3.38 (0.124)
N=386	MAD	2.19 (0.121)	1.14 (0.109)
	Uncertainty	3.99 (0.167)	2.97 (0.149)
	Disagreem. (%)	11.46 (1.17)	7.84 (1.05)
1 Year Change	Mean	2.42 (0.198)	5.17 (0.209)
N=223	MAD	2.01 (0.145)	2.25 (0.145)
	Uncertainty	3.85 (0.239)	3.48 (0.234)
	Disagreem. (%)	15.25 (1.89)	18.33 (2.09)
10 Year Change	Mean	1.82 (0.179)	0.92 (0.164)
N=197	MAD	1.79 (0.125)	1.35 (0.132)
	Uncertainty	4.27 (0.226)	2.28 (0.162)
	Disagreem. (%)	7.98 (1.29)	12.01 (1.73)
Information Not Shown			
All	Mean	2.07 (0.185)	2.66 (0.225)
N=313	MAD	2.17 (0.139)	2.64 (0.168)
	Uncertainty	4.32 (0.211)	3.78 (0.205)
	Disagreem. (%)	9.61 (1.15)	18.64 (1.75)
Forecast	Mean	1.97 (0.247)	2.99 (0.311)
N=123	MAD	1.93 (0.175)	2.38 (0.225)
	Uncertainty	4.41 (0.336)	3.49 (0.295)
	Disagreem. (%)	9.64 (1.76)	18.26 (2.75)
1 Year Change	Mean	2.32 (0.403)	2.56 (0.475)
N=92	MAD	2.61 (0.296)	2.97 (0.358)
	Uncertainty	5.23 (0.424)	4.84 (0.468)
	Disagreem. (%)	9.65 (2.23)	18.37 (3.18)
10 Year Change	Mean	2.29 (0.549)	2.60 (0.484)
N=50	MAD	2.55 (0.411)	2.48 (0.331)
	Uncertainty	3.47 (0.469)	3.39 (0.419)
	Disagreem. (%)	8.33 (2.55)	17.88 (4.14)

Notes: The average level, the dispersion, the uncertainty, and the fraction of disagreements within group is presented for the prior and posterior 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. See notes to Table 3 for additional details on definitions of various measures. Numbers in parentheses in each cell are standard errors.